Maximising the mutual interoperability of an MRI scanner and a cancer therapy particle accelerator

Brendan Whelan

This thesis is submitted in fulfilment of the requirements for the degree of Doctor of Philosophy.

University of Sydney

Faculty of Medicine

2016

It's my thesis, and I'll cry if I want to

Statement of originality

I, Brendan Whelan hereby declare that this thesis is my own work and has not been submitted in any form for another degree or diploma at any university or other institute of tertiary education. Results obtained with the assistance of other researchers are acknowledged in Appendix 2: Statements of authorship contributions. Where previously published information is referred to or used, it is properly acknowledged with references which are listed at the end of each chapter. Some chapters have been accepted or submitted for publication in academic journals; these are acknowledged below.

Acknowledgments

Completing a research doctorate is often thought of a horrendous ordeal, but personally I've had quite a good time! It is an enormous privilege to be in a position to complete a PhD, and along with a healthy dose of luck, the fact I am in this position is nearly entirely due to my family. In particular, I would like to thank my parents Barry and Anne for all the years of housing, food and clothing, and making the time to drive me to music practice, scouts, tennis etc. etc! I would also like to thank my sister Claire for her friendship over the last 27 years. I know that I would be a very different person had I not had such a supportive family who always encouraged me to explore my interests.

I would like to thank old and new friends that made living in a new city such a blast. Pat Edsall, Sophie Thrower and Sasha Saunders made me feel incredibly welcome when I first arrived in Sydney having never lived out of home before. Long live the goon sunrise and the Allman sausage fest! Emma Colvill and Sean Pollock were integral participants at many important career development meetings (which coincidentally were often held during the happy hour at corridor). Others offered important international perspectives; special mention has to go to Chiara Paganelli, Stefan Kolling and Svenja Ipsen. One of the best outcomes of this PhD is the friends I would never have made had I not undertaken it.

I would like to thank past and present members of the Radiation Physics Lab at Sydney University, many of whom had a big influence on my development as a researcher. Similarly, I would like to acknowledge the members the medical physics department and Ingham research institute in Liverpool (Particularly Michael Barton, Gary Liney, and Michael Jameson), the Fahrig lab at Stanford, and accelerator division at SLAC (particularly John Schmerge). A huge thankyou to the Gold family and to Waldo Hinshaw and Susan Kopiwoda for generously welcoming me into their homes when I was in Stanford.

Finally, I have to thank my supervisors. I say have to (1) because they might fail me if I don't but (2) because I absolutely mean it. There are two official supervisors listed on this thesis: Paul Keall and Lois Holloway, but there are also two people I consider as honorary supervisors: Rebecca Fahrig and Steven Gierman. If I had to give one piece of advice to an aspiring researcher, it would be this: work with good people. I managed to achieve this more through dumb luck than any applied strategy, but I can't imagine that I could have done much better! I learnt a lot from all of them; Lois had a major influence very early in my research career when she supervised my master's degree and strongly encouraged me to take on a PhD. She always tries to approach problems from a different angle and loves to think outside the box. Paul works harder and with more focus than anyone I have met. He has an amazing ability to make other people feel positive and empowered about their work and leads by example. BUT, he is not very good at squash! Rebecca questions everything around her with tireless energy, and this very often seems to reveal hidden and important details that make things happen! I also have to thank her and her family for their hospitality when I was in Stanford. Steve has an incredible metaphysical framework for the physical problems he approaches as part of his job and thinks about physics on a very deep level. Often when I was confused about an aspect of accelerator physics which no one else I knew had even heard of,

I was able to discuss it with Steve over lunch and come away with a much better understanding, or at least a reassurance that I wasn't completely daft for finding it confusing! After a few dinners at the fish market in Palo Alto, I also suspect he could have an alternative career as a political scientist should he so desire!

CONTENTS

Publications and presentations arising from this thesis

Chapter 3: Gridded electron gun

This chapter was accepted for publication in Medical Physics on the 11th of September 2016.

In addition, the following conference presentations resulted from this work:

ESTRO 2016, E-poster: Whelan, B., D. Constantin, R. Fahrig, P. Keall, L. Holloway, and B. Oborn, *EP-1929: Characterisation of a gridded electron gun in magnetic fields: implications for MRI-Linac therapy.* Radiotherapy and Oncology, 2016. **119**: p. S915.

AAPM 2016, oral presentation: Whelan, B., M. Bazalova-Carter, B. Oborn, D. Constantin, L. Holloway, R. Fahrig, and P. Keall, *TU-H-BRA-03: Performance of a Clinical Gridded Electron Gun in Magnetic Fields: Implications for MRI-Linac Therapy.* Medical Physics, 2016. **43**(6): p. 3769-3769.

Chapter 4: Passive Magnetic Shielding

This work has not yet been published; however it is planned that a version of this work will be submitted to either medical physics or PMB in the near future. One conference presentation was given on a much earlier incarnation of this work:

AAPM 2014, oral presentation: Whelan, B., D. Constantin, L. Holloway, S. Kolling, B. Oborn, R. Fahrig, and P. Keall, *WE-G-17A-08: Electron Gun Operation for in Line MRI-Linac Configurations: An Assessment of Beam Fidelity and Recovery Techniques for Different SIDs and Magnetic Field Strengths.* Medical Physics, 2014. **41**(6): p. 525-526.

Chapter 5: A novel electron accelerator for MRI-Linac radiotherapy

This work was published in medical physics on 17th February 2016. In addition, the following conference presentations resulted from this work:

AAPM 2015: oral presentation: Whelan, B., S. Gierman, J. Schmerge, L. Holloway, P. Keall, and R. Fahrig, *WE-G-BRD-09: Novel MRI Compatible Electron Accelerator for MRI-Linac Radiotherapy.* Medical physics, 2015. **42**(6): p. 3690-3690.

Chapter 6: Experimental characterisation of an RF gun in magnetic fields

This work has not yet been published; however it is planned that a version of this work will be submitted to either medical physics or PMB upon completion of the experimental data collection. One conference presentation was given on an earlier version of this work:

AAPM 2016: oral presentation: Whelan, B., P. Keall, L. Holloway, S. Gierman, J. Schmerge, S. Tantawi, A. Tremaine, A. Trautwein, B. Scott, and R. Fahrig, *TU-H-BRA-07: Design, Construction, and Installation of An Experimental Beam Line for the Development of MRI-Linac Compatible Electron Accelerator.* Medical Physics, 2016. **43**(6): p. 3770-3770.

Chapter 7: Patient reported outcomes of slow, single arc rotation

This work was submitted to the international journal of Radiotherapy and Oncology on the 23rd of August 2016. At the time of writing, it was under peer review. In addition, the following conference presentations resulted from this work:

ESTRO 2016, E-poster: Whelan, B., M. Welgampola, L. McGarvie, K. Makhija, I. Feain, L. Holloway, M. Berry, M. Barton, R. Turner, and M. Jackson, *EP-1930: Cancer patient experience of slow, single arc rotation to simplify radiation therapy delivery*. Radiotherapy and Oncology, 2016. **119**: p. S916.

Chapter 8: MRI compatible patient rotation system

This work was accepted for publication in Medical Physics on the 1st of December 2016.

Appendix 1: AAPM research database

This work was accepted for publication in medical physics on the 29th of December 2016. In addition the following invited presentation was given on this work.

AAPM 2016: Invited presentation: Whelan, B., *TU-FG-207B-01: Report On Development of a Database of NIH Funding of Medical Physicists*. Medical Physics, 2016. **43**(6): p. 3763-3764.

Abstract

The work described in this PhD thesis was undertaken as part of a much larger research project: *The Australian MRI-Linac program*. The goal of this program is to merge two existing medical technologies – an MRI scanner and a Linear Accelerator (Linac) – thereby creating an advanced form of cancer treatment incorporating cutting edge anatomical and physiological imaging techniques. An overview of the background information necessary to understand the work presented in this thesis is provided in chapters 1 (overview of radiotherapy) and 2 (overview of electromagnetism and accelerator physics). The work in the remainder of this thesis can be split into two distinct sections, corresponding to the two quite different (but ultimately related) projects I worked on throughout this thesis: modelling the impact of external magnetic fields on electron beam transport within the linear accelerator, and the implementation of patient rotation in radiotherapy.

The former project is the focus of Chapters 3-6. In Chapter 3 a finite element model of a clinical gridded electron gun is developed based on 3D laser scanning and electrical measurements, and the sensitivity of this gun in magnetic fields characterised. The results complement the existing literature in showing that conventional linear accelerator components are very sensitive to external magnetic fields – in fact this gun is over twice as sensitive to axial magnetic fields than the less realistic models existing in the literature.

A first order approach to overcoming this sensitivity is to use magnetic shielding – however magnetic shielding of the linear accelerator can negatively impact on the performance of the MRI scanner. This magnetic shielding problem is explored in Chapter 4, where the fundamental principles of passive magnetic shielding are explored, and magnetic shields are implemented for the two possible MRI-linac configurations (in-line and perpendicular) for the 1.0 Tesla MRI magnet used in the Australian MRI Linac program. The efficacy of the shielding and the impact on the MRI is quantified, with the conclusion that passive shielding could be successfully implemented to allow acceptable operation of the linac without overly degrading the magnet performance of the MRI scanner.

An alternative approach to magnetic shielding which would not have any impact on the magnet is to redesign the linear accelerator such that it functions robustly in an MRI environment without the need for shielding. This approach is explored in chapter 5, where a novel electron accelerator concept based on an RF-electron gun configuration is detailed. It is shown via particle in cell simulations that such a design would be able to operate in a wide range of axial magnetic fields with minimal current loss. In chapter 6, an experimental beam line based on this concept was constructed at Stanford Linear Accelerator Center (SLAC). This project is ongoing but progress so far is described in Chapter 6.

In the second part of this thesis, a completely different project is explored, patient rotation. Patient rotation would be very beneficial for MRI-Linac systems as it would eliminate the complicated engineering that is used in conventional systems to rotate the beam around the patient, and the MRI could be used to adapt in real time for the resultant anatomic deformation. Patient rotation would also minimise some of the sources of electromagnetic interference explored in chapters 3-7. The two major obstacles to patient rotation are (1)

patient tolerance to rotation, and (2) anatomical deformation due to rotation. To quantify patient rotation, a clinical study of 15 patients was carried out and is detailed in chapter 7. The results of this study suggest that patient tolerance to rotation may not be a major issue, although this result needs to be verified in larger patient cohorts. In chapter 8, the design and construction of an MRI-compatible patient rotation device is detailed. This device is the first of its kind, and will allow data on anatomic deformation under rotation to be collected, enabling strategies to adapt for this motion to be developed. Thus far, MRI compatibility has been assessed and a volunteer imaging study undertaken, in which pelvic images were acquired under rotation angles of 360^0 every 45^0.

In summary: In chapters 3-5, the impact of magnetic fields on conventional accelerator components was quantified; and two independent approaches to compensating for these effects (magnetic shielding and bespoke accelerator design) were explored. In chapter 6, an experimental beam is constructed to verify and support the findings of chapter 6. In chapter 7, a clinical study was undertaken quantifying patient tolerance of slow, single arc rotation. Finally, in chapter 8 a unique medical device was designed, constructed and tested, and through this device MRI images of anatomical distortion under lying rotation were collected and quantified.

Chapter 1
Overview of radiotherapy

1. Cancer & Radiotherapy

Cancer describes a large number of diseases characterized by malignant and uncontrolled cellular growth. Cancer is a leading cause of death and a heavy social health burden – in Australia around 3 in every 10 deaths are due to cancer, costing more than $3.8 billion annually. Globally, the economic impact of cancer approaches 1 trillion dollars annually [6, 7].

The major treatment options for cancer are surgery, chemotherapy, and radiotherapy, with the latter discipline being the focus of this thesis. Radiotherapy refers to the treatment of cancer with ionising radiation – that is, radiation which is energetic enough to break apart chemical bonds[*]. In best practice cancer care, radiotherapy would be included in the treatment plan of around half of all cancer patients [8]. In Australia, currently around 38% of all cancer patients receive radiotherapy [9].

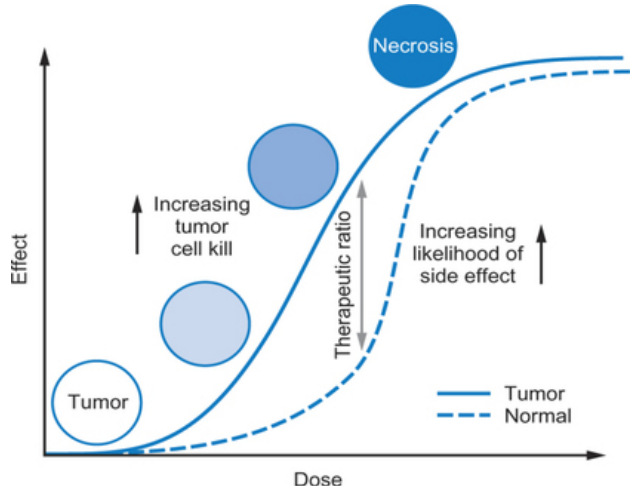

Figure 1: The probability of both cure and complication increases as the dose is increased. For a given dose, there is some ratio between these two outcomes, called the therapeutic ratio. Radiotherapy aims to maximise this goal. Image credit: http://creogreview.wikifoundry.com/page/Q261+Therapeutic+ratio

The basic mechanism of cancer treatment by radiation is as follows: when a beam of radiation traverses matter, interactions take place between that matter and the radiation which result in energy being transferred from the radiation beam to the matter. When the matter in question is biological, this energy deposition can break the chemical bonds that hold cells together, resulting in cell damage or death. There is an obvious and fundamental problem with utilising this process for treatment: radiation is not a discerning therapeutic agent. Whether cancerous or healthy, tissues will sustain damage when irradiated. As such, practitioners of radiotherapy are faced with a delicate quandary; for whilst higher doses of radiation correlate with a higher probability of curing the cancer, they correlate also with a higher probability of causing irreparable harm or even death to the patient. This balance between cure and trauma is known as the therapeutic ratio (Figure 1), and increasing this ratio is, in one way or another, at the core of almost every piece of radiotherapy research ever undertaken (or at least it should be). The concept of the therapeutic ratio can be expounded to formulate a fundamental tenet of curative radiotherapy:

[*] Radiation is in fact a very general term, simply describing the transport of energy by either waves or particles – however, in this thesis the term 'radiation' refers to ionising radiation.

To deliver sufficient radiation dose to cancerous tissues to maximise the chance of therapeutic benefit, whilst simultaneously minimizing the chance of serious side effects.

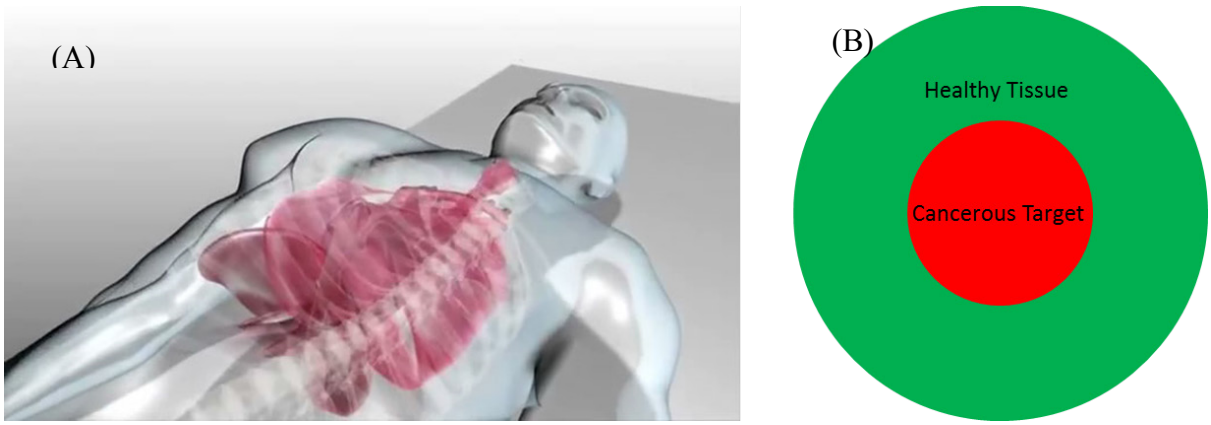

Figure 2: (A) Patients are (unfortunately) comprised of geometrically complex, moving, and heterogeneous tissues. For the purpose of deriving the key elements of high quality radiotherapy, I will instead consider (B) the idealised 'spherical patient'.

To determine how one would go about achieving this, let us consider the tasks which must be carried out to deliver high quality radiotherapy:

1.1. Tasks which must be completed to deliver high quality radiotherapy

I. Identify and localise all tissue types

It is axiomatic that in order to maximise the chance of hitting a target, one must first determine where that target is. In this instance, as well as delivering a curative quantity of dose to the cancerous target, we want to make sure that we *minimise* the chance of side effects to all healthy tissues. Clearly, in order to maximise the therapeutic ratio we will need to accurately determine the location of every tissue type likely to be exposed to radiation.

In clinical practice, doctors undertake this task with the aid of 3D medical imaging (typically Computed Tomography, or CT, imaging) which they use to help delineate the location of the therapeutic target and surrounding healthy organs in space. An example of a CT image with organs delineated is shown in Figure 3 (A) – note that only one 2D slice of the entire 3D image is shown.

II. Determine the optimal feasible dose distribution

Once we have determined where our target is, we must decide on the optimal dose distribution to treat it with. To guide this decision, we will need to know something about the dose response of the tissues involved - What dose *is* sufficient to maximise the chance of cure? At what dose *will* side effects occur in healthy tissue?

Of course, on the face of things this may seem an over complication, for surely the optimal dose distribution is quite clear: an infinite dose in the cancerous cells, and zero dose everywhere else. Unfortunately, there is no way to deliver such a dose distribution, and that is why I have emphasised that our task at this stage is to find the optimal *feasible* dose distribution. To do this, we must consider also the exact mechanism with which we will deliver dose (Task 3).

In clinical practice, this task is undertaken with the help of the 3D image from Task 1. Computational simulation of dose delivery is carried out, using this image as a surrogate for the patient, and various delivery parameters are tweaked to ensure a curative dose is delivered to the tumour whilst the surrounding organs receive an acceptably low dose. An example of a treatment plan is shown in Figure 3 (B).

III. Accurately deliver the planned dose

Once we have identified all the different tissues, their dose responses, and (somehow) decided upon the optimal feasible dose distribution, all that remains is to deliver this treatment. If careful attention has been payed to designing a *feasible* dose distribution, then this last step should be straightforward, since we will already have incorporated the delivery mechanism into our decision making process of Task 2.

A variety of delivery devices exist in radiotherapy, however the focus of this thesis is on photon External Beam Radiotherapy delivered with a linear accelerator (Linac), shown in Figure 3(C).

Returning now to what I have asserted is the fundamental tenet of curative radiotherapy (To deliver sufficient radiation dose to a cancerous tissues to maximise the chance of cure, whilst simultaneously minimising any serious side effects), it can be seen that carrying out Tasks 1-3 to a high (or at least acceptable) standard will result in high (or at least acceptable) quality radiotherapy. Of course, I have given a very brief and simplified overview of a radiotherapy workflow – but the derived principles hold even in more complex real world scenarios.

2. Major challenges to delivering high quality radiotherapy

45 years ago, US president Richard Nixon declared ‘war on cancer’, aimed at eliminating death and suffering from cancer [10]. Similarly, in 2003 Director of the USA’s National Cancer Institute Andrew von Eschenbach suggested that medical researchers should aim to ‘rid the world of suffering and death from cancer by 2015’ [11]. It is the year 2016 at the time of writing, and the fact that cancer remains such a massive social health burden (Section 1) implies (correctly) that effective cancer treatment is extremely difficult to achieve. All current treatments for cancer must overcome massive challenges if treatment efficacy is to be improved; radiotherapy is no different. In this section, I will outline what I believe are the four biggest challenges to delivering better quality radiotherapy.

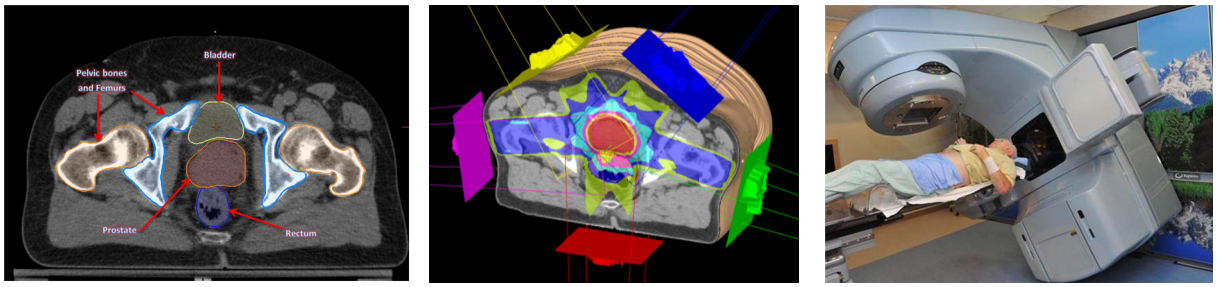

A) Localise all tissues **B) Determine optimal dose distribution** **C) Accurately deliver planned dose**

Figure 3: Delivering high quality radiotherapy can be split into three tasks: A) Localisation of the treatment volume and surrounding organs at risk. Shown here is a one slice of a CT image of a prostate cancer patient B) Determination of the optimal feasible dose distribution through computerised simulation [5], and C) Accurate delivery of the planned dose. Shown here is a linear accelerator – the workhorse of external beam radiation therapy.

I. Which tissues are cancerous and which tissues are not?

As outlined in Section 1, the first step to delivering high quality treatment is to determine which tissues need to be treated – that is, which tissues are cancerous and which are healthy? To determine this, doctors use 3D medical images to delineate the volume they believe requires a curative dose, as well as any surrounding organs at risk to which dose must be minimised. This is a very challenging task, as medical images are difficult to interpret and generally do not resolve cancerous from healthy tissues particularly well, leaving clinicians to make an 'educated guess' as to which is which. As such, a large amount of uncertainty is introduced into the treatment at this stage. Ensuring that the delineation process is as accurate as possible is obviously crucial to achieving an optimal treatment outcome. If the treatment volume is too large, then we are needlessly delivering high radiation dose to healthy tissues, whilst if it is too small, we increase the risk that not all the cancerous cells are killed and that treatment fails. Figure 4 shows an example of contouring variability between different clinicians for a lung cancer patient [4].

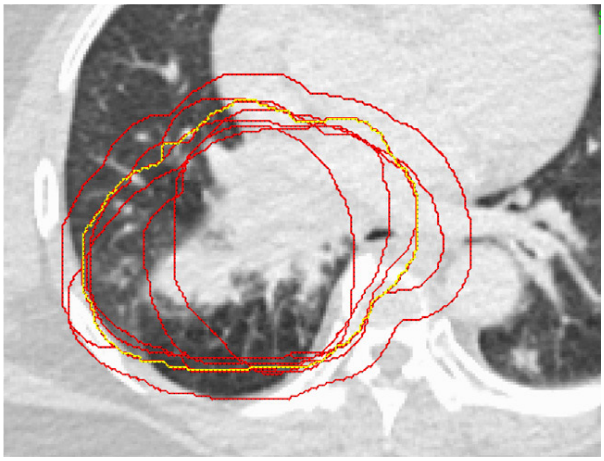

Figure 4: Contours drawn by different clinicians on a lung tumour case. The gold contour indicates the volume of best agreement based on all clinicians, whilst the red contours show the different volumes. It should be noted that disagreement between different clinicians is not always this stark! [4].

Target delineation variation has been measured at up to 10 mm standard deviation between contours [12], up to eight times difference in total volume [13], and estimated to account for variation in tumour control probability of up to 22% in lung cancer patients [4]. Large delineation uncertainty can also exist around normal anatomy, however this has not been as well studied [14-16]. It has been argued that target delineation uncertainty is the 'weakest link' in the chain of accurate radiotherapy treatment [17], or, more conservatively; "Interobserver variability is a major – for some tumour locations probably the largest – factor contributing to geometric inaccuracy" [13].

II. Different patients respond differently to treatment

Figure 5 shows a Kaplan-Meier survival curve from a recent clinical trial into a new radiotherapy treatment regime[†] [3]. Two years after the treatment, 50% of patients have survived disease free, and 50% are either dead or have recurrent disease – in other words, treatment has not been successful for half of these patients[‡]. This figure is from a controlled clinical trial where patients were all recruited with the same tumour site and staging and all given the same treatment – and yet the outcomes show large divergence. What's more, we are not good at predicting which group of patients are which – in another recent study on lung cancer patients, doctors were asked to predict which of their patients would still be alive in two years. Their predictions were correct around 50% of the time – in other words, the same accuracy as might be expected from flipping a coin [18].

Radiation response is difficult to predict not only for cancerous cells, but also for normal tissues [19]. In the case of normal tissues, quantifying dose response is not so simple – whereas typically the entire treatment volume receives a homogenous and prescribed dose, normal tissues typically receive a heterogeneous dose which is simply kept as low as possible. Also, it is not only the dose received by an organ which impact on response, but also the *volume* of the organ which receives a given dose [20]. Figure 6(A) shows the risk of developing pneumonitis, a serious side effect amongst lung cancer patients [20], and Figure 6(B) shows a summary of the various published data on pneumonitis [21]. Again, the difficulty of predicting the likelihood of an individual patient developing complications from radiotherapy treatment can be appreciated.

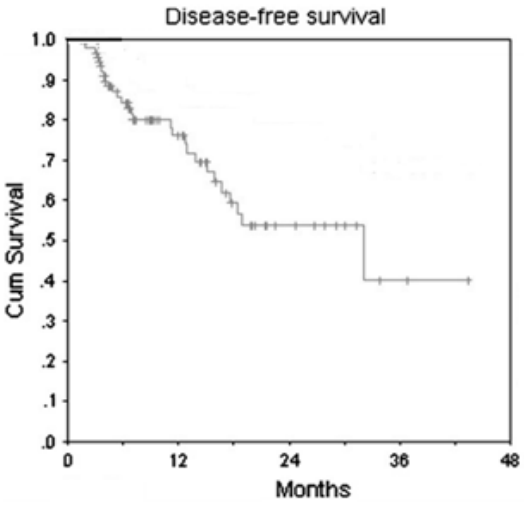

Figure 5: A Kaplan-Meier survival curve from a clinical trial of stage 1 non-small cell lung cancer patients, showing that two years after treatment, 50% of patients have survived disease free and 50% have not.[3]

[†] There is nothing especially pertinent or remarkable about this particular trial or survival curve – it is used as a representative example of the uncertainty in radiotherapy treatment outcomes.

[‡] This is a loose definition of 'successful' – there is still a high chance that treatment increased the length and/or quality of these patients lives.

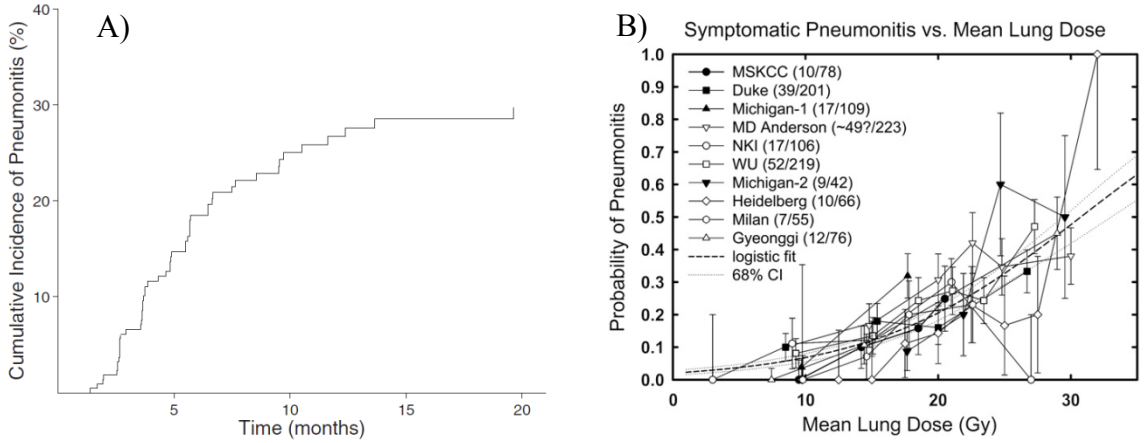

Figure 7: A) shows incidence of the normal tissue injury pneumonitis from a retrospective analysis of lung cancer patients. 20 months after treatment, around 30% of patients have developed some form of pneumonitis. B) Meta-analysis of several studies showing published values for mean lung dose versus pneumonitis probability. It can be see that whilst a loose correlation between mean lung dose and probability of pneumonitis exists, there is also a large degree of variability.

The variability in treatment outcomes is due to a myriad of complex and interwoven factors. The most obvious explanation is that there is a large variation in inherent radiosensitivity amongst patients – that is; just as some people get sunburnt much more easily than others, some people are much more sensitive to radiation than others. Whilst this does appear to be at least part of the explanation for treatment response variation, it is difficult to separate radio sensitivity effects from a vast array of complicating factors. Factors which have been nominated as contributing in some way to Radiation sensitivity include patient age, patient gender, smoking status, comorbidities (Primary and secondary malignancies, diabetes etc.), and intrinsic radio sensitivity [19, 22]. Overall, variation in radiation response is still not particularly well understood, and as such it is very difficult to predict how individual patients are likely to respond to treatment – both in terms of tumour control probability and normal tissue complication probability. Whilst clinicians to their best to develop a treatment plan which is optimal for each patient, this is very difficult given the unpredictability of patient response.

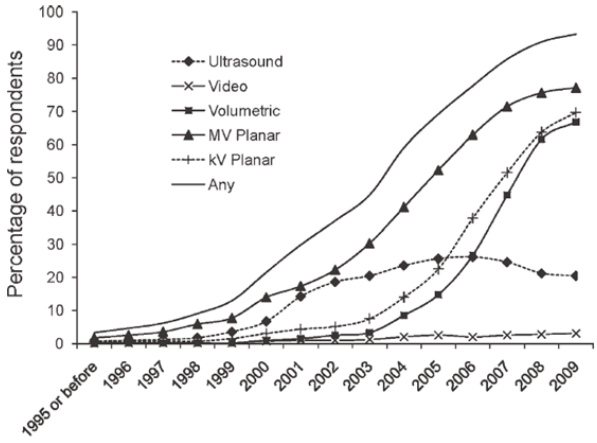

Figure 6: Percentage of clinics using various image guidance techniques as a function of year. [1].

III. Anatomy and physiology change over the course of treatment

If the uncertainties described in Sections I and II did not exist, then could we deliver a truly optimal plan to every patient and be confident of predicting their treatment outcomes?

Unfortunately: no. Human anatomy is in a constant state of geometric and physiological flux as our hearts beat, lungs inhale, and our digestive systems process food. This means that even

if we localise the relevant organs with 100% accuracy, the shape and location of these organs may have changed by the time we deliver therapy. In particular in the thoracic and abdominal regions, our anatomy is in constant motion due to respiration.

Of the challenges that I have listed thus far, there is no doubt that anatomic motion is the best studied and understood – probably because it is much easier to measure and understand than the problems described in Sections I and II. Respiration can cause motion of up to 50 mm, and in addition to the lungs, notably affects the liver, breast, kidney, heart, and pancreas [2, 23] (Figure 8). Rectal and bladder filling can cause motion of up to 40 mm in cervix cancer and prostate cancer cases [24, 25]. In addition to these physiologically based sources of motion, there are sources such as voluntary patient movement (fidgeting etc.) and set up errors (where the patient is not accurately positioned in the same way from one treatment to the next) [26].

In addition to the sources of geometric changes mentioned above, the treatment itself introduces a source of anatomical change. Throughout the course of treatment the tumour will (we hope) reduce in the size or (sometimes) grow. Tumour volumes have been reported as changing by up to 96 % in cervix cancer patients [24], 62% in lung cancer patients [27] (Figure 7) and 92% in head and neck cancer patients [28]. Changes in tumour volume can also cause changes in the surrounding healthy tissues.

In order to compensate for motion, various in-room techniques exist to monitor and adapt for motion throughout the course of treatment. The most commonly used technique is cone-beam CT imaging – however, solutions utilizing optical imaging, ultrasound imaging and electromagnetic transponders also exist. [29]. Collectively these techniques are called Image Guided Radiotherapy (IGRT). The use of IGRT in the clinic grew dramatically from 1995 to 2009 as shown in Figure 7**.** [1]. The use of image guidance helps to compensate for many of the motion effects outlined above. However, the current standard of care image guidance techniques suffer from a number of shortcomings – the quality of the images is often poor, additional internal markers need to be surgically implanted, the information returned is generally limited in either the temporal or spatial domains, and the extra radiation imaging dose received by healthy tissues can often be significant [30]. As such, motion management in radiotherapy is still a significant issue and the subject of ongoing research.

		Breathing mode	
Site	Observer	Shallow	Deep
Pancreas	Suramo (Ref. 57)	20 (10–30)	43 (20–80)
	Bryan (Ref. 59)	20 (0–35)	-
Liver	Weiss (Ref. 66)	13±5	-
	Harauz (Ref. 67)	14	-
	Suramo (Ref. 57)	25 (10–40)	55 (30–80)
	Davies (Ref. 58)	10 (5–17)	37 (21–57)
Kidney	Suramo (Ref. 57)	19 (10–40)	40 (20–70)
	Davies (Ref. 58)	11 (5–16)	-
Diaphragm	Wade (Ref. 68)	17	101
	Korin (Ref. 64)	13	39
	Davies (Ref. 58)	12 (7–28)	43 (25–57)
	Weiss (Ref. 66)	13±5	-
	Giraud (Ref. 63)	-	35 (3–95)
	Ford (Ref. 76)	20 (13–31)	-

Figure 8: Summary of various studies looking at respiratory induced motion during radiotherapy. Taken from [2]. All motion reported in mm.

IV. Physical limitations on dose conformity

The goal of radiotherapy is to deliver a curative dose to the tumour without causing any serious side effects in the surrounding healthy tissues. Typically, this requires that a higher

radiation dose is delivered to the tumour than to the tissues surrounding it. The ratio of dose inside the treatment region to outside the treatment region which can be achieved is (loosely) termed dose conformity. Dose conformity quantifies the extent to which we can deposit dose where we *do*[§] want it and avoid where we *don't* want it. There are many ways to increase dose conformity. For instance, modern radiotherapy uses multiple beams which are modulated in space and time to sculpt dose to a treatment region. Increasingly complex modulation techniques can lead to increased dose conformity [31]. As can be seen in Figure 9 [32], it is also possible to increase dose conformity by changing the type of radiation beam used to deliver dose. Protons and other heavy particles are intrinsically more conformal than photons due to differences in the way that they deposit dose.

The search for increased dose conformity in radiotherapy has been a controversial topic for the last two decades (at least) [33-37], and is likely to remain so for the foreseeable future. Critics argue that there is a lack of evidence that increased dose conformity improves outcomes – or, that any improvement in outcomes is not enough to justify the increased cost. For instance, proton therapy is estimated to cost around three times standard of care photon therapy [38], yet outcome differences are often minimal. Meanwhile, proponents for newer and more conformal techniques point to the improved dose distributions offered by new techniques as intrinsic proof of their efficacy.

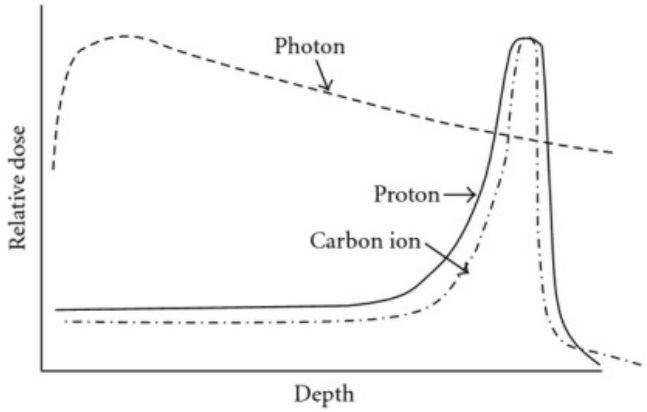

Figure 9: Different types of radiation deposit dose differently as they pass through matter. Protons, carbon ions, and other heavy particles can be used to achieve more conformal dose distributions.

It may seem that settling this debate should be a trivial matter – can't we simply collect the data of patients treated with different techniques and compare the treatment outcomes? Unfortunately, collecting such data in radiotherapy (and medicine generally) is time consuming, difficult, costly, and sometimes ethically fraught [39]. There is reasonably good evidence that advanced conformal photon treatments do result in lower side effects to patients [40, 41]. There is less evidence that they improve tumour control – however, this is not entirely surprising given that typically the same dose has been given to the treatment volume. There has been one study which showed dose escalation using more conformal techniques could improve tumour control without increasing complication rates [41, 42]. For therapy with heavy particles, the situation is less clear. Whilst these treatments have generally been shown to be effective and well tolerated, there is very little or no clinical evidence for increased efficacy compared to photon treatments (depending on treatment site and what one accepts as evidence) [43-45]. As Bentzen notes; "Lack of evidence (for an effect) is not evidence for a lack (of that effect)" [39]. This is true, but it is also true that the larger an

[§] Actually, where we *think* we want it…which may not be the same as where we really want it.

effect is, the easier it should be to find evidence pointing towards it – and in the case of proton therapy, it is not easy [46]. The importance of dose conformity depends on the dose response curves and relative location of the tumour and the critical tissues surrounding it. At some point, increasing dose conformity will yield diminishing returns in terms of treatment outcomes – exactly where this point is – or whether current advanced photon techniques approach it – is not clear.

3. MRI: Magic Radiotherapy Imaging?

In Section 2, I outlined the major challenges to improving the efficacy of radiotherapy. In this section, I will outline the potential of Magnetic Resonance Imaging (MRI, or as coined by Vincent Khoo; Magic Radiotherapy Imaging [47]) to positively impact multiple facets of the radiotherapy workflow.

Exactly what MRI is will be briefly explained in more detail in chapter 2. For now, it is sufficient to state that MRI is a medical imaging modality capable of forming 3D images of internal anatomy with exquisite soft tissue contrast. MRI is capable of achieving excellent spatial and temporal resolution, and is capable of imaging physiological processes such as blood flow and perfusion, and can even be used to directly detect cancer biomarkers (spectroscopic imaging) [48]. In addition to this, there are no known long term biological side effects from imaging with MRI, making it ideal for repeated scans and extended imaging times. There are a number of ways in which MRI could be superior to the X-ray CT imaging more conventionally used in radiotherapy – these will be briefly reviewed below.

I. Tumour delineation and staging

As outlined in Section 1.2I, there is considerable uncertainty in the initial process of localising the tumour and surrounding tissues. At least part of the problem here is that the images used to inform this task do not show adequate contrast between different tissue types. MRI images have inherently greater soft tissue contrast than CT images, and given tumours and most of the organs surrounding them *are* soft tissue, it follows that separating the different tissue types should be easier on MRI than on CT.

If this statement is true, then one would expect that inter-observer uncertainty (the difference between the contours drawn by different physicians, Figure 4) should be less on MRI than CT. Gratifyingly, this logical inference has largely been supported by published data – MRI has been shown to reduce inter-observer variability in several treatment sites, notably prostate, Head and Neck, base of skull, partial breast, and rectal cancer [49-51]. However, in truth, the existing data must be termed encouraging rather than conclusive, since the statistical and clinical significance of the reductions in inter-observer variation are not yet clear (at least to this reader). There are also some reports in the literature that MRI did *not* significantly improve inter-observer variability, notably a 20 patient study on pharyngo-laryngeal cancers where inter-observer variability was already quite low on CT images [52, 53]. Of course, inter-observer variability is only a surrogate metric to assess the quality of delineation – the true metric is something we do not have access to, which is how closely the delineated volume matches the cancerous volume. In addition to this, it should be noted that

the above studies are inherently biased against MRI, since whilst most clinicians have little experience with MRIs versus CT. Although the scientific evidence is not crystal clear, in general there is belief in the radiotherapy community that MRI based delineation is superior to CT, and this is leading an increasing number of centres to pursue MRI based workflows.

II. Physiological and Anatomical imaging and response monitoring

As outlined in section 1.2.II, a significant problem with the current practice of radiotherapy is that we are often unable to accurately predict how different patients will respond to radiotherapy. MRI has the potential to improve this situation by utilising functional imaging to extract physiological as well as anatomical information. Functional imaging is an umbrella term referring to any imaging technique which returns some information on physiological function, and along with MRI includes PET, SPECT, ultrasound and CT techniques [54]. There are a number of functional MRI techniques which fall under this definition; some of the most important are diffusion weighted imaging (DWI), dynamic contrast enhanced imaging (DCE), hyperpolarised MRI, and imaging to assess the oxygenation status of tumours (multiple techniques) [55]. These techniques can be used to increase prognostic confidence both by pre-treatment imaging and response monitoring throughout treatment. Considering the uncertainties outlined in 1.2.II, it is clear that an increased ability to triage patients into response groups offers the potential for large improvements in therapeutic outcomes. Although it is early days, there are several promising studies which suggest functional MRI has an important role to play. Recently, it was shown that DCE MRI was able to separate good responders from poor responders in rectal cancer patients [56]. DWI MRI can be used to assess tumour recurrence and treatment response in prostate, head and neck, liver, and breast cancers [57-59]. The potential for MRI to enable functional lung avoidance has been shown and is currently the subject of a randomised clinical trial [60, 61], and spectroscopic images have been shown to have strong diagnostic and response potential in brain, prostate and breast [62].

III. Real time tumour tracking and adaption

As outlined in Section 2.III, changes in tumour size and location throughout the course of treatment introduce significant uncertainty into the radiotherapy treatment process. Again, MRI has significant potential to address this uncertainty. Firstly, in light of the above information it can be appreciated that MRI is an ideal tool to assess tumour regression and movement throughout treatment, and if necessary, adapt the treatment. MRI is also suitable for monitoring tumour motion resulting from respiration and even cardiac motion. MRI can be used to develop motion models of the tumour position in pre-treatment imaging, which can inform the initial treatment plan. MRI is also capable of real time imaging, which opens up the exciting possibility of real time treatment adaption to tumour motion during treatment. This would involve integration of an MRI scanner with the radiotherapy treatment device - the problem tackled in this thesis.

4. Conclusions

Radiotherapy is one of the most important and cost effective treatment modalities for cancer. In Australia, it is currently used to treat around 38% of cancer patients, whilst evidence suggests it should optimally be used for closer to 50%. There are a number of ways in which current radiotherapy could be improved; better tumour delineation, adaptation to anatomical changes throughout the course of treatment, and greater understanding of the physiological responses of individual patients. Magnetic Resonance Imaging (MRI) has the potential to address all these issues; as such, there is an increasing focus on incorporating MRI into radiotherapy. This includes developing integrated MRI-Linac systems, which is the problem addressed by this thesis.

References

1. Simpson, D.R., J.D. Lawson, S.K. Nath, B.S. Rose, A.J. Mundt, and L.K. Mell, *A survey of image-guided radiation therapy use in the United States.* Cancer, 2010. **116**(16): p. 3953-3960.
2. Keall, P.J., G.S. Mageras, J.M. Balter, R.S. Emery, K.M. Forster, S.B. Jiang, J.M. Kapatoes, D.A. Low, M.J. Murphy, and B.R. Murray, *The management of respiratory motion in radiation oncology report of AAPM Task Group 76a).* Medical physics, 2006. **33**(10): p. 3874-3900.
3. Lagerwaard, F.J., C.J. Haasbeek, E.F. Smit, B.J. Slotman, and S. Senan, *Outcomes of risk-adapted fractionated stereotactic radiotherapy for stage I non–small-cell lung cancer.* International Journal of Radiation Oncology* Biology* Physics, 2008. **70**(3): p. 685-692.
4. Jameson, M.G., S. Kumar, S.K. Vinod, P.E. Metcalfe, and L.C. Holloway, *Correlation of contouring variation with modeled outcome for conformal non-small cell lung cancer radiotherapy.* Radiotherapy and Oncology, 2014. **112**(3): p. 332-336.
5. Varadhan, R., G. Karangelis, K. Krishnan, and S. Hui, *A framework for deformable image registration validation in radiotherapy clinical applications.* Journal of applied clinical medical physics/American College of Medical Physics, 2013. **14**(1): p. 4066.
6. Cancer Council Australia. ***Facts and Figures***. 2015; Available from: http://www.cancer.org.au/about-cancer/what-is-cancer/facts-and-figures.html.
7. American Cancer Society, *The Global Economic Cost of Cancer*. 2010: Atlanta: American Cancer Society.
8. Barton, M.B., S. Jacob, J. Shafiq, K. Wong, S.R. Thompson, T.P. Hanna, and G.P. Delaney, *Estimating the demand for radiotherapy from the evidence: a review of changes from 2003 to 2012.* Radiotherapy and Oncology, 2014. **112**(1): p. 140-144.
9. Morgan, G., M. Barton, C. Atkinson, J. Millar, N.K. Gogna, and E. Yeoh, *'GAP'in radiotherapy services in Australia and New Zealand in 2009.* Journal of medical imaging and radiation oncology, 2010. **54**(3): p. 287-297.
10. Spector, R., *The war on cancer. A progress report for skeptics.* Skeptical Inquirer, 2010. **34**(1): p. 25-31.
11. von Eschenbach, A.C., *A vision for the National Cancer Program in the United States.* Nature Reviews Cancer, 2004. **4**(10): p. 820-828.
12. Steenbakkers, R.J., J.C. Duppen, I. Fitton, K.E. Deurloo, L.J. Zijp, E.F. Comans, A.L. Uitterhoeve, P.T. Rodrigus, G.W. Kramer, and J. Bussink, *Reduction of observer variation using matched CT-PET for lung cancer delineation: a three-dimensional analysis.* International Journal of Radiation Oncology* Biology* Physics, 2006. **64**(2): p. 435-448.
13. Weiss, E. and C.F. Hess, *The impact of gross tumor volume (GTV) and clinical target volume (CTV) definition on the total accuracy in radiotherapy.* Strahlentherapie und Onkologie, 2003. **179**(1): p. 21-30.

14. Saarnak, A.E., M. Boersma, B.N. van Bunningen, R. Wolterink, and M.J. Steggerda, *Inter-observer variation in delineation of bladder and rectum contours for brachytherapy of cervical cancer.* Radiotherapy and Oncology, 2000. **56**(1): p. 37-42.
15. Mukesh, M., R. Benson, R. Jena, A. Hoole, T. Roques, C. Scrase, C. Martin, G. Whitfield, J. Gemmill, and S. Jefferies, *Interobserver variation in clinical target volume and organs at risk segmentation in post-parotidectomy radiotherapy: can segmentation protocols help?* 2014.
16. Martin, S., C. Johnson, M. Brophy, D.A. Palma, J.L. Barron, S.S. Beauchemin, A.V. Louie, E. Yu, B. Yaremko, and B. Ahmad, *Impact of target volume segmentation accuracy and variability on treatment planning for 4D-CT-based non-small cell lung cancer radiotherapy.* Acta Oncologica, 2014(0): p. 1-11.
17. Njeh, C., *Tumor delineation: The weakest link in the search for accuracy in radiotherapy.* Journal of medical physics/Association of Medical Physicists of India, 2008. **33**(4): p. 136.
18. Oberije, C., G. Nalbantov, A. Dekker, L. Boersma, J. Borger, B. Reymen, A. van Baardwijk, R. Wanders, D. De Ruysscher, and E. Steyerberg, *A prospective study comparing the predictions of doctors versus models for treatment outcome of lung cancer patients: A step toward individualized care and shared decision making.* Radiotherapy and Oncology, 2014. **112**(1): p. 37-43.
19. Bentzen, S.M. and J. Overgaard. *Patient-to-patient variability in the expression of radiation-induced normal tissue injury*. in *Seminars in radiation oncology*. 1994. Elsevier.
20. Hope, A.J., P.E. Lindsay, I. El Naqa, J.R. Alaly, M. Vicic, J.D. Bradley, and J.O. Deasy, *Modeling radiation pneumonitis risk with clinical, dosimetric, and spatial parameters.* International Journal of Radiation Oncology* Biology* Physics, 2006. **65**(1): p. 112-124.
21. Marks, L.B., S.M. Bentzen, J.O. Deasy, J.D. Bradley, I.S. Vogelius, I. El Naqa, J.L. Hubbs, J.V. Lebesque, R.D. Timmerman, and M.K. Martel, *Radiation dose–volume effects in the lung.* International Journal of Radiation Oncology* Biology* Physics, 2010. **76**(3): p. S70-S76.
22. Bentzen, S.M., *Preventing or reducing late side effects of radiation therapy: radiobiology meets molecular pathology.* Nature Reviews Cancer, 2006. **6**(9): p. 702-713.
23. Korreman, S.S., *Motion in radiotherapy: photon therapy.* Physics in medicine and biology, 2012. **57**(23): p. R161.
24. Lim, K., V. Kelly, J. Stewart, J. Xie, Y.-B. Cho, J. Moseley, K. Brock, A. Fyles, A. Lundin, and H. Rehbinder, *Pelvic radiotherapy for cancer of the cervix: is what you plan actually what you deliver?* International Journal of Radiation Oncology* Biology* Physics, 2009. **74**(1): p. 304-312.
25. Kupelian, P., T. Willoughby, A. Mahadevan, T. Djemil, G. Weinstein, S. Jani, C. Enke, T. Solberg, N. Flores, and D. Liu, *Multi-institutional clinical experience with the Calypso System in localization and continuous, real-time monitoring of the prostate gland during external radiotherapy.* International Journal of Radiation Oncology* Biology* Physics, 2007. **67**(4): p. 1088-1098.
26. Hurkmans, C.W., P. Remeijer, J.V. Lebesque, and B.J. Mijnheer, *Set-up verification using portal imaging; review of current clinical practice.* Radiotherapy and oncology, 2001. **58**(2): p. 105-120.
27. Fox, J., E. Ford, K. Redmond, J. Zhou, J. Wong, and D.Y. Song, *Quantification of tumor volume changes during radiotherapy for non–small-cell lung cancer.* International Journal of Radiation Oncology* Biology* Physics, 2009. **74**(2): p. 341-348.
28. Barker, J.L., A.S. Garden, K.K. Ang, J.C. O'Daniel, H. Wang, L.E. Court, W.H. Morrison, D.I. Rosenthal, K.C. Chao, and S.L. Tucker, *Quantification of volumetric and geometric changes occurring during fractionated radiotherapy for head-and-neck cancer using an integrated CT/linear accelerator system.* International Journal of Radiation Oncology* Biology* Physics, 2004. **59**(4): p. 960-970.
29. Verellen, D., M. De Ridder, N. Linthout, K. Tournel, G. Soete, and G. Storme, *Innovations in image-guided radiotherapy.* Nature Reviews Cancer, 2007. **7**(12): p. 949-960.
30. Smith-Bindman, R., J. Lipson, R. Marcus, K.-P. Kim, M. Mahesh, R. Gould, A.B. de González, and D.L. Miglioretti, *Radiation dose associated with common computed*

tomography examinations and the associated lifetime attributable risk of cancer. Archives of internal medicine, 2009. **169**(22): p. 2078-2086.
31. Palma, D.A., W.F. Verbakel, K. Otto, and S. Senan, *New developments in arc radiation therapy: a review.* Cancer treatment reviews, 2010. **36**(5): p. 393-399.
32. Fukumitsu, N., *Particle beam therapy for cancer of the skull base, nasal cavity, and paranasal sinus.* ISRN otolaryngology, 2012. **2012**.
33. Tait, D.M., A.E. Nahum, L.C. Meyer, M. Law, D.P. Dearnaley, A. Horwich, W.P. Mayles, and J.R. Yarnold, *Acute toxicity in pelvic radiotherapy; a randomised trial of conformal versus conventional treatment.* Radiotherapy and oncology, 1997. **42**(2): p. 121-135.
34. Brada, M., M. Pijls-Johannesma, and D. De Ruysscher, *Proton therapy in clinical practice: current clinical evidence.* Journal of clinical oncology, 2007. **25**(8): p. 965-970.
35. Suit, H., H. Kooy, A. Trofimov, J. Farr, J. Munzenrider, T. DeLaney, J. Loeffler, B. Clasie, S. Safai, and H. Paganetti, *Should positive phase III clinical trial data be required before proton beam therapy is more widely adopted? No.* Radiotherapy and Oncology, 2008. **86**(2): p. 148-153.
36. Miller, R.C., M. Lodge, M.H. Murad, and B. Jones. *Controversies in clinical trials in proton radiotherapy: The present and the future*. in *Seminars in radiation oncology*. 2013. Elsevier.
37. De Neve, W., W. De Gersem, and I. Madani. *Rational use of intensity-modulated radiation therapy: the importance of clinical outcome*. in *Seminars in radiation oncology*. 2012. Elsevier.
38. Peeters, A., J.P. Grutters, M. Pijls-Johannesma, S. Reimoser, D. De Ruysscher, J.L. Severens, M.A. Joore, and P. Lambin, *How costly is particle therapy? Cost analysis of external beam radiotherapy with carbon-ions, protons and photons.* Radiotherapy and Oncology, 2010. **95**(1): p. 45-53.
39. Bentzen, S.M., *Randomized controlled trials in health technology assessment: Overkill or overdue?* Radiotherapy and Oncology, 2008. **86**(2): p. 142-147.
40. Veldeman, L., I. Madani, F. Hulstaert, G. De Meerleer, M. Mareel, and W. De Neve, *Evidence behind use of intensity-modulated radiotherapy: a systematic review of comparative clinical studies.* The lancet oncology, 2008. **9**(4): p. 367-375.
41. Staffurth, J., *A review of the clinical evidence for intensity-modulated radiotherapy.* Clinical oncology, 2010. **22**(8): p. 643-657.
42. Vora, S.A., W.W. Wong, S.E. Schild, G.A. Ezzell, and M.Y. Halyard, *Analysis of biochemical control and prognostic factors in patients treated with either low-dose three-dimensional conformal radiation therapy or high-dose intensity-modulated radiotherapy for localized prostate cancer.* International Journal of Radiation Oncology* Biology* Physics, 2007. **68**(4): p. 1053-1058.
43. De Ruysscher, D., M.M. Lodge, B. Jones, M. Brada, A. Munro, T. Jefferson, and M. Pijls-Johannesma, *Charged particles in radiotherapy: a 5-year update of a systematic review.* Radiotherapy and Oncology, 2012. **103**(1): p. 5-7.
44. Allen, A.M., T. Pawlicki, L. Dong, E. Fourkal, M. Buyyounouski, K. Cengel, J. Plastaras, M.K. Bucci, T.I. Yock, and L. Bonilla, *An evidence based review of proton beam therapy: the report of ASTRO's emerging technology committee.* Radiotherapy and Oncology, 2012. **103**(1): p. 8-11.
45. Fujii, O., Y. Demizu, N. Hashimoto, M. Araya, M. Takagi, K. Terashima, M. Mima, H. Iwata, Y. Niwa, and D. Jin, *A retrospective comparison of proton therapy and carbon ion therapy for stage I non-small cell lung cancer.* Radiotherapy and Oncology, 2013. **109**(1): p. 32-37.
46. Sheets, N.C., G.H. Goldin, A.-M. Meyer, Y. Wu, Y. Chang, T. Stürmer, J.A. Holmes, B.B. Reeve, P.A. Godley, and W.R. Carpenter, *Intensity-modulated radiation therapy, proton therapy, or conformal radiation therapy and morbidity and disease control in localized prostate cancer.* Jama, 2012. **307**(15): p. 1611-1620.
47. Khoo, V., *MRI--" magic radiotherapy imaging" for treatment planning?* The British journal of radiology, 2000. **73**(867): p. 229-233.
48. Dirix, P., K. Haustermans, and V. Vandecaveye. *The Value of Magnetic Resonance Imaging for Radiotherapy Planning*. in *Seminars in radiation oncology*. 2014. Elsevier.

49. Rasch, C., R. Steenbakkers, and M. van Herk. *Target definition in prostate, head, and neck*. in *Seminars in radiation oncology*. 2005. Elsevier.
50. Villeirs, G.M., K. Van Vaerenbergh, L. Vakaet, S. Bral, F. Claus, W.J. De Neve, K.L. Verstraete, and G.O. De Meerleer, *Interobserver Delineation Variation Using CT versus Combined CT+ MRI in Intensity–Modulated Radiotherapy for Prostate Cancer.* Strahlentherapie und Onkologie, 2005. **181**(7): p. 424-430.
51. Khoo, V.S., E.J. Adams, F. Saran, J.L. Bedford, J.R. Perks, A.P. Warrington, and M. Brada, *A comparison of clinical target volumes determined by CT and MRI for the radiotherapy planning of base of skull meningiomas.* International Journal of Radiation Oncology* Biology* Physics, 2000. **46**(5): p. 1309-1317.
52. Weltens, C., J. Menten, M. Feron, E. Bellon, P. Demaerel, F. Maes, W. Van den Bogaert, and E. van der Schueren, *Interobserver variations in gross tumor volume delineation of brain tumors on computed tomography and impact of magnetic resonance imaging.* Radiotherapy and Oncology, 2001. **60**(1): p. 49-59.
53. Geets, X., J.-F. Daisne, S. Arcangeli, E. Coche, M. De Poel, T. Duprez, G. Nardella, and V. Grégoire, *Inter-observer variability in the delineation of pharyngo-laryngeal tumor, parotid glands and cervical spinal cord: comparison between CT-scan and MRI.* Radiotherapy and Oncology, 2005. **77**(1): p. 25-31.
54. Desar, I., C. van Herpen, H. van Laarhoven, J. Barentsz, W. Oyen, and W. van der Graaf, *Beyond RECIST: molecular and functional imaging techniques for evaluation of response to targeted therapy.* Cancer treatment reviews, 2009. **35**(4): p. 309-321.
55. Metcalfe, P., G. Liney, L. Holloway, A. Walker, M. Barton, G. Delaney, S. Vinod, and W. Tome, *The potential for an enhanced role for MRI in radiation-therapy treatment planning.* Technology in cancer research & treatment, 2013. **12**(5): p. 429-446.
56. Martens, M.H., S. Subhani, L.A. Heijnen, D.M. Lambregts, J. Buijsen, M. Maas, R.G. Riedl, C.R. Jeukens, G.L. Beets, and E. Kluza, *Can perfusion MRI predict response to preoperative treatment in rectal cancer?* Radiotherapy and Oncology, 2014.
57. Song, I., C.K. Kim, B.K. Park, and W. Park, *Assessment of response to radiotherapy for prostate cancer: value of diffusion-weighted MRI at 3 T.* American Journal of Roentgenology, 2010. **194**(6): p. W477-W482.
58. Vandecaveye, V., F. De Keyzer, S. Nuyts, K. Deraedt, P. Dirix, P. Hamaekers, V. Vander Poorten, P. Delaere, and R. Hermans, *Detection of head and neck squamous cell carcinoma with diffusion weighted MRI after (chemo) radiotherapy: correlation between radiologic and histopathologic findings.* International Journal of Radiation Oncology* Biology* Physics, 2007. **67**(4): p. 960-971.
59. Thoeny, H.C. and B.D. Ross, *Predicting and monitoring cancer treatment response with diffusion-weighted MRI.* Journal of Magnetic Resonance Imaging, 2010. **32**(1): p. 2-16.
60. Hoover, D.A., D.P. Capaldi, K. Sheikh, D.A. Palma, G.B. Rodrigues, A.R. Dar, E. Yu, B. Dingle, M. Landis, and W. Kocha, *Functional lung avoidance for individualized radiotherapy (FLAIR): study protocol for a randomized, double-blind clinical trial.* BMC cancer, 2014. **14**(1): p. 934.
61. Sheikh, K., D.P. Capaldi, D.A. Hoover, D.A. Palma, B.P. Yaremko, and G. Parraga, *Magnetic resonance imaging biomarkers of chronic obstructive pulmonary disease prior to radiation therapy for non-small cell lung cancer.* European Journal of Radiology Open, 2015.
62. Kwock, L., J.K. Smith, M. Castillo, M.G. Ewend, F. Collichio, D.E. Morris, T.W. Bouldin, and S. Cush, *Clinical role of proton magnetic resonance spectroscopy in oncology: brain, breast, and prostate cancer.* The lancet oncology, 2006. **7**(10): p. 859-868.

Chapter 2
Physics of MRI-Linac systems

1. Introduction

In chapter 1, an overview of obstacles to increasing the quality of radiation therapy treatments was given, and a case made that MRI, and in particular integrated MRI-Linacs, have the potential to address many of these issues. Large portions of the remainder of this thesis will deal with quantifying and mitigating the physical interactions that can occur when linac and MRI are placed in close proximity. In this chapter, a brief overview of the fundamental physical principles which govern these interactions is given. The purpose here is not to provide a comprehensive overview of the topics covered here (electrodynamics, MRI physics, and accelerator physics). Rather, the intention is to provide a very top level introduction into the physics which is relevant to the next three chapters of this thesis, and to compile a list of references such that further information can be found if so desired.

2. Electromagnetism

I. Introduction to electromagnetism

According to modern physical theory, every single physical interaction in the universe can ultimately be traced back to one of four fundamental forces: gravity, the strong nuclear force, the weak nuclear force, and the *electromagnetic force*. The latter is the mechanism via which MRI scanners and linear accelerators become coupled when placed in close proximity. Fundamentally, electromagnetism is the study of charged particles, and the way they interact with each other – so the first question which arises is: what is charge? Charge is a fundamental property held by some sub atomic and elemental particles, the electron and proton being the most common examples. Particles which have charge are subject to electromagnetic forces; of course this is a circular definition since I said the electromagnetic force is defined by how charged particles interact! Nevertheless, it is the best I can do – such is life when dealing with fundamental quantities! Charge is measured in Coulombs (C) and can be either negative or positive. An electron has a charge of $-1.602*10^{-19}$ C, and a proton $+1.602*10^{-19}$ C. Two positive charges repulse each other, whilst positive and negative charges attract each other. The mechanism by which they do this is called an *electromagnetic field.* A field is loosely defined as a quantity which has a value at every point in space in time; for instance one may also speak of a temperature field or a gravitational field. Electromagnetic fields are a special type of field called a *vector field*; this means that at every point, they not only have a *value* (like a temperature field) they have a *direction* (like a gravitational field). There are two fields of primary importance in electrodynamics; the *electric field* and the *magnetic field*. (If you delve deeply enough into physical theory, you will learn that these are in fact two facets of the one field – but practically, they act quite differently and are treated accordingly). All the material covered in this section is quite fundamental; as such multiple authors have developed extensive references on it. Two I recommend are Griffiths [2] and Jackson [3].

An electric field is created by electric charge. Left to its own devices, the field would extend from the electric charge to infinity; the field strength at distance r from the charge of q is given by:

$$\mathbf{E} = \frac{q}{4\pi\varepsilon_0 r^2} \qquad \textbf{Equation 1}$$

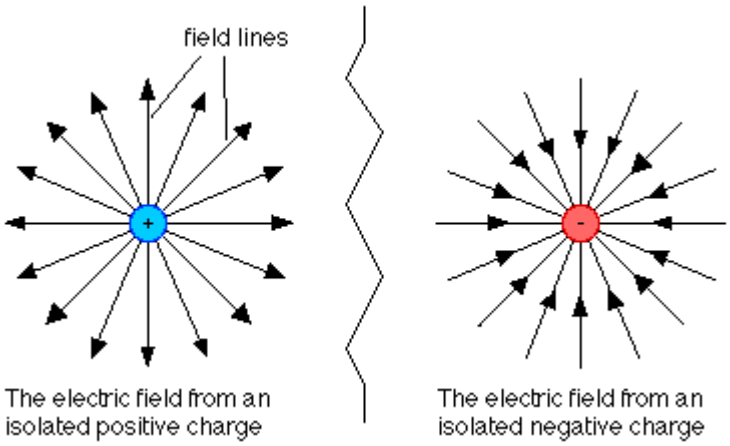

Figure 10: The electric field around a positive and negative charged. Image credit: http://physics.bu.edu/~duffy/PY106/Electricfield.html

(ε_0 is the vacuum permittivity constant). Electric fields point away from positive charge, and towards negative charge – although this definition is traditional rather than fundamental, the important thing to note is that there are both sources and sinks for the electromagnetic field. Ok, so that's what the field *is* – but what does it *do*?? Well, say we put another charge within this field; this charge will then experience a force:

$$\mathbf{F_{electric}} = q\mathbf{E} \qquad \textbf{Equation 2}$$

Of course, the new charge creates a field of its own, and so our initial charge will experience an equal force in the opposite direction (which would make Newton very happy).

The second field which must be defined is the magnetic field. Where electric fields are caused by any charged particle – magnetic fields only come into existence when particles *move*. Of course, a moving charge is simply a current (Current=Coulombs/second=Amperes). Rather than consider the magnetic field arising from a single moving charge (which is actually quite complicated!), let us consider the case of an infinite straight wire carrying a steady current. The magnetic field at a distance r from the wire is:

$$\mathbf{B} = \frac{\mu_0 \mathbf{I}}{2\pi r} \qquad \textbf{Equation 3}$$

In equation 3, **I** is the current through the loop and μ_0 is the vacuum permeability constant. There is an important difference between magnetic fields and electric fields; unlike electric fields, magnetic fields do not 'begin' and 'end' somewhere but rather form continuous closed loops. This is a consequence of the fact that the magnetic field is always produced by currents rather than fundamental particles.

If we were to place another charge at a distance r from the wire, then it would experience a force of…nothing**! What trickery is this?! Well, it turns out that just as magnetic forces are created when charge is moving, they also act only *on* charge that is moving. The equation which governs the force a particle of charge q moving with speed **v** would experience in a magnetic field **B** is:

$$\mathbf{F}_{magnetic} = q(\mathbf{v} \times \mathbf{B}) \qquad \textbf{Equation 4}$$

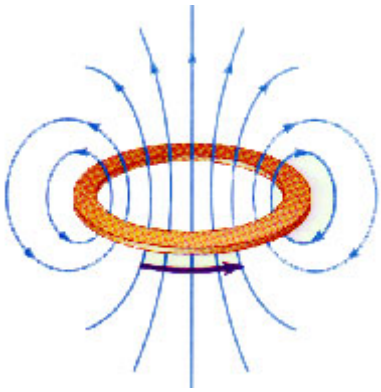

Figure 11: magnetic field from a current loop. Image credit: http://www.ux1.eiu.edu/~cfadd/1360/30MagSrcs/HmwkSol.html

Of course, in general we are dealing with situations much more complex than a single point of charge, or an infinite line of current. As such the simple formulas laid out above are not adequate. Luckily, there is a set of generalized formulas governing the relationships between currents, charges, and fields. These formulas are called Maxwell's equations:

$$\nabla . \mathbf{E} = \frac{1}{\varepsilon_0}\rho \qquad \textbf{Equation 5: Gauss's Law}$$

$$\nabla \times \mathbf{E} = -\partial \mathbf{B}/\partial t \qquad \textbf{Equation 6: Faraday's Law}$$

$$\nabla . \mathbf{B} = 0 \qquad \textbf{Equation 7: No name}$$

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{J} + \mu_0 \varepsilon_0 \, \partial \mathbf{E}/\partial t \qquad \textbf{Equation 8: Ampere's Law}$$

In these equations, **E** is the electric field, **B** is the magnetic field, **J** is a current distribution, μ_0 & ε_0 are the constants equal to the permeability and permittivity of free space, t is time, and ρ is a distribution of charge (The notation used here is that of vector; see e.g. Griffiths for further explanation [2]). There is a similarly succinct law governing the force experienced by a charged particle placed into an electromagnetic field: the Lorentz force law, which is just the combination of the force laws defined in equations 2 and 4.

$$\mathbf{F} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}) \qquad \textbf{Equation 9: The Lorentz force law}$$

** This assumes that whilst there is charge moving through the wire, the wire is also electrically neutral so there is no electric force. This situation sounds a little contrived, but in fact is the normal state of affairs in a current carrying wire.

II. Magnetic materials

Thus far I have explained magnetism purely in terms of moving charges and changing electric fields, but this seems to have little connection with the more commonly experienced phenomenon of magnetism, which of course lies in the objects you use to attach the takeaway pizza menu to your fridge. Fridge magnets are examples of magnetic materials, which are materials that generate their own magnetic field. How do materials do this? Surely they're not all filled up with tiny current loops, which each generates its own magnetic field…well, essentially – yes! We call these atoms! In most materials, the electrons of the atom combine in such a way that the net current[††] is 0, and these materials do not exhibit magnetic effects[‡‡]. However, other materials (normally with an odd number of electrons) have some net current, and as such generate a tiny magnetic field. In general, all these tiny magnetic fields are orientated at random – but when an external magnetic field is applied, each tiny magnetic field tends to align with this, in which case we say that the material is *magnetized*. Materials which behave like this are called paramagnetic materials. Materials like fridge magnets however, have an additional trick – the material stays magnetized, even after the magnetic field has been removed! These materials are called *ferromagnetic.* Only three elemental materials exhibit ferromagnetism; iron, nickel, and cobalt. The reason ferromagnetic materials behave like this is that each of the little magnetic fields feels a force from its neighbors, and they 'like' to line up in the same direction [2].

III. Electromagnetic Waves (Radiofrequency radiation)

So far, we have discussed what happens when charge is motionless (electric fields) and when it's moving (magnetic fields). There is a third scenario we must discuss: what happens when charge is *accelerating*? As it turns out, something altogether different! When charged particles are accelerated, *electromagnetic waves,* also called *radiofrequency (RF) radiation* is created. Electromagnetic waves consist of oscillating electric and magnetic fields which carry energy through space (Figure 3); in fact RF radiation can easily travel far enough away from whatever charged particles produced it that we forget about them altogether and focus on the RF radiation as an independent entity. As it turns out, we are all very familiar with electromagnetic waves; this is what light is!

[††] Of course, we are dealing on a quantum scale here, and I am being a little loose with my definitions of current loops. In fact the picture I have painted is fundamentally incorrect, as magnetism is very much a quantum mechanical phenomena. Nevertheless, it is a useful picture, and adequate for our needs.

[‡‡] Again, this isn't quite true – sorry! Such materials still respond to magnetic fields via diamagnetism – but this is so weak compared to the either types, it can essentially be ignored in most cases.

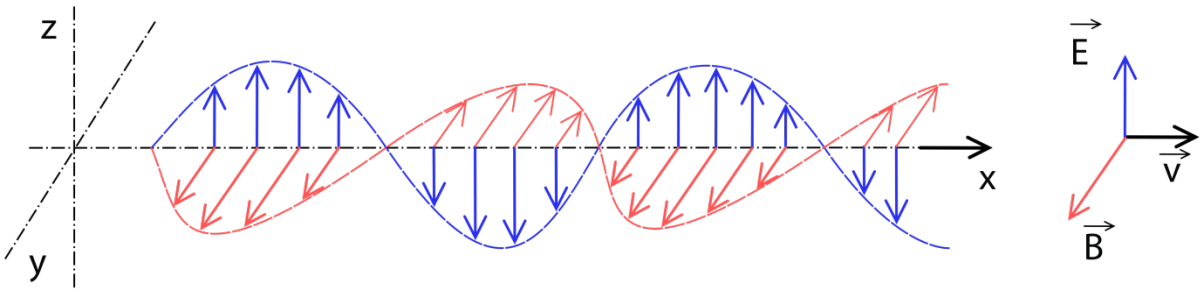

Figure 12: Electromagnetic radiation comprises electric and magnetic fields which are (in a vacuum) transverse to the direction of propagation.
Image credit:
https://commons.wikimedia.org/wiki/File:Onde_electromagnetique.svg

Electromagnetic waves are of importance to several parts of this thesis: the x-rays used to treat cancer are very high energy RF radiation, the accelerating waveguides which produce these x-rays are filled with RF energy, and MRI machines are basically just complicated RF detectors. Some important properties of RF radiation are listed below:

1. *RF radiation is principally defined by its frequency*

What separates different colours of light from each other, and light from X-rays, or X-rays from microwaves? The *frequency* of the radiation. Frequency simply refers to how often per second the field oscillates back and forth[§§], and is measured in Hertz. The frequency of RF radiation also defines its energy – for instance, ultraviolet radiation is a higher frequency than other electromagnetic radiation reaching earth's surface from the sun, and this is why it burns you. (The intention of this paragraph is not to imply that the frequency of RF radiation is the *only* thing which defines it. Other quantities of importance in various scenarios include polarisation, phase, and bandwidth)

[§§] If you think about it – it is quite amazing that our eyes are able to so simply detect the difference between RF radiation oscillating at 450 trillion times a second (blue) and 550 trillion times a second (green).

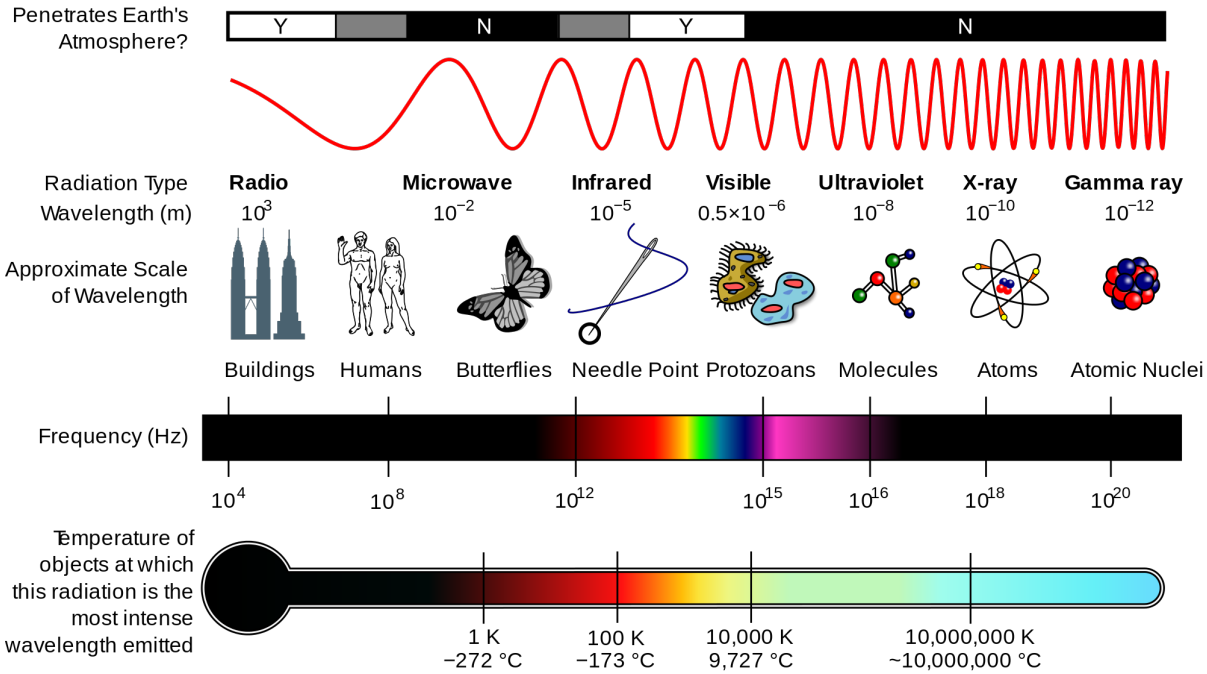

Figure 13: The electromagnetic spectrum. Different frequencies of electromagnetic radiation can behave quite differently and are used for different purposes. Image credit: https://commons.wikimedia.org/wiki/File:EM_Spectrum_Properties_edit.svg

2. *RF radiation can be treated as a particle or a wave*

In fact according to quantum theory, everything (including you) can be treated as a particle or a wave! However (unlike you!) this duality is of particular importance in RF radiation. The extent to which the wave or particle approach is appropriate really depends on the wavelength of the radiation, and the physical scale you are interested in working with. For instance, in Figure 4, we would nearly always treat very high frequency radiation such as gamma rays as particles and low frequency radiation such as radio waves as waves, whilst intermediate frequencies such as visible light are approached depending on the application at hand! In this work, we treat the RF in the accelerator structure as a wave, since the wavelength of this radiation is 10cm, which is comparable to the dimensions of interest. On the other hand, X-rays in dosimetry are essentially always treated as particles - despite being fundamentally the same thing[***].

3. *RF radiation can be guided (and shielded)*

Finally, and of great importance to this thesis; RF radiation can be guided down specially designed tubes called waveguides. The geometry and composition of the waveguide depends on the frequency of radiation to be guided, and the waveguide must be carefully designed for the need at hand. Everyday examples of waveguides include optical fibres and coaxial cables (the thing that connects your TV to your antenna). Essentially, these waveguides are simply specially constructed mirrors which the RF 'bounces' down. Of course, the same physics which is used to keep RF radiation *in* one place can also be used to keep it *out* of another, and this can also be of crucial importance when it comes to shielding sensitive MRI scanners from external RF radiation.

3. Physics of MRI

In Chapter 1, I outlined the potential of Magnetic Resonance Imaging for improving a broad range of uncertainty in radiotherapy treatments - however I have not yet explained what it is or what it does (although I did claim it was magic!). In fact, when the mechanisms underlying MRI are examined in more depth, I think that 'magic' is not that farfetched! Here is the short story: our body is made up of billions upon billions of atoms - mostly carbon, oxygen and hydrogen[†††]. Each hydrogen atom consists of 1 electron and 1 proton – and via some rather fortuitous quirks in the fundamental laws of quantum mechanics, the lone proton, and hence nucleus as a whole, has a magnetic field arising from the unpaired spin of the proton. When our body is placed in a strong magnetic field, these magnetic fields become preferentially aligned – that is, we become magnetised! (It is however important to note that when I previously talked about magnetisation, I talked about atomic magnetisation, whereas now we are talking of *nuclear* magnetisation, i.e. it is the protons rather than the electrons which

[***] The reasons for all this confusion are historic – really, we should not say that RF radiation is 'both' a wave and a particle; we should say it is neither! It is a phenomena which under certain circumstances exhibits behaviours consistent with historically defined concepts of (1) waves and (2) particles, and so we are stuck with the convention of saying it is two things at once, which of course is complete rubbish! It is one thing, and we don't have a proper word for it.

[†††] Actually about 7×10^{27} atoms, of which 99% are oxygen, carbon, or hydrogen.

become magnetised[‡‡‡]). MRI works by extremely clever manipulation and measurement of this nuclear magnetisation to develop an image. The contrast in MRI images is due to the fact that different tissue types have slightly different magnetic properties, and hence the nuclear spins which are being measured behave slightly differently. The fact that we can measure signal from nuclear magnetisation is rather incredible, and is the basis for the well-known analytic technique of nuclear magnetic resonance analysis (NMR) which has been in use since the 1940's. However, the ground-breaking work which made MRI possible was the ability to *localise* where this signal was coming from, developed in the 1970's. This breakthrough was awarded the 2003 Nobel Prize in physiology and medicine [4]. I will not be giving a detailed explanation of how MRI works, since (1) it has been covered in many excellent textbooks already [5-7], and (2) it is actually not of direct relevance to anything else in this thesis. In order to understand the motivation for the work undertaken here, it is necessary only to know the following three facts:

1. *MRI requires a very strong magnetic field*

Typically clinical MRI scanners operate between 0.3 and 3 Tesla. To put this in context, the magnetic field from a strong bar magnet is 0.01 Tesla. The magnetic fields in MRI magnets are (normally) generated via current loops. Because of the huge amounts of current required to generate magnetic fields this strong, if left to their own device the coils would quickly become extremely hot and melt! In order to avoid this, the coils are typically superconducting, which means they have no resistance to current. The magnetic field falls in strength as the distance from the scanner increases. Typically around 3 meters away from the scanner the field has become negligible, but of course this is magnet specific.

2. *MRI requires a very homogenous magnetic field*

In addition to being very strong, the magnetic field used for MRI imaging must be extremely homogenous over the region being imaged in order to generate accurate images. This is because any change magnetic field yields to a change in the magnetisation of tissue, which results in image artifacts and geometric distortion. Typically, the maximum variation over a 30 cm diameter sphere is on the order of a few parts per million. Achieving this level of homogeneity is very difficult, and a huge amount of engineering is required. Of

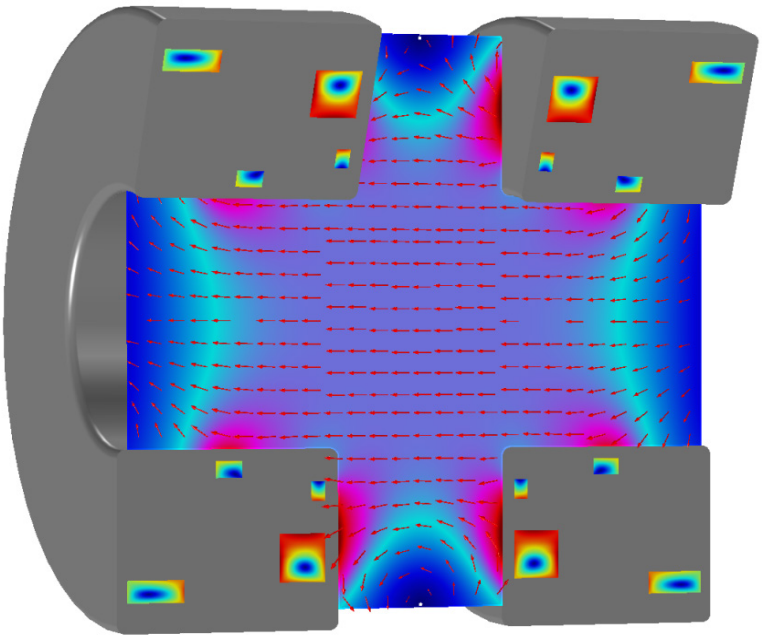

Figure 14: MRI magnets typically use superconducting loops of wire to produce strong and homogenous magnetic fields at the centre of the scanner.

nce imaging. However, the nuclear was dropped from the name in recognition of the fact that the general public are quite suspicious of anything ıfer than X-ray imaging.

course, any other magnetic fields which are introduced to the scanner will degrade the homogeneity of the magnet, so MRI scanners are normally operated in magnetically shielded rooms, and of course are extremely sensitive to ferromagnetic materials (iron, nickel, cobalt) being inside these room.

3. *MRI scanners form image based on the detection of RF signals.*

The manner in which an MRI scanner gathers information about the nuclear magnetisation of the object being scanned is RF radiation. A set of three gradient coils is used to spatially encode the RF signal in terms of frequency, phase, and axial position. In essence, the scanner repeatedly sends some RF signals into the body, measures the spatially encoded signals which come out, and is eventually able to form an image. This means that MRI scanners are extremely sensitive to external sources of RF radiation. As such, MRI scanners are typically operated inside a room which is shielded against RF radiation – but again, this is of no use if a source of RF radiation (say, a linear accelerator) is actually inside this room!

These three statements entail essentially all you need to know about the physics of an MRI scanner in order to understand this thesis - which of course is vastly less than there *is* to know! For further reading, I encourage perusal of the supplied references [5-7].

4. Accelerator physics and beam dynamics

Particle accelerators use the electromagnetic force to accelerate particles to high energy. Particle accelerators are used in high energy physics, nuclear physics, high resolution imaging fundamental science, semiconductor fabrication, cargo and explosives scanning, the creation of medical isotopes for nuclear medicine, and of course in radiation oncology (Chapter 1). [8]. In radiation oncology, the most common type of accelerator which is utilized is a linear electron accelerator (Linac). Most of the time the electrons are not directly used to irradiate tissue; instead they are collided with a heavy metal target, thereby producing x-rays (Remember in section 2.II I said that when charged particles are accelerating they produce RF energy? That is exactly how the X-rays are produced!). The electron energies which are most commonly used to give high quality results in radiation therapy range from ~1 MeV to ~18 MeV. 1 MeV is the amount kinetic energy possessed by an electron which has been accelerated through an electric potential of 1 million volts. To put this value in context, consider that the potential difference between the ionosphere and the ground generated by a lightning storm is on the order of 0.5 MeV[§§§]!

The physics and engineering governing the operation of charged particle accelerators can be loosely split into two sections; the first is to consider the kinematics of particles in time varying electromagnetic fields (I will call this beam dynamics); the second is the process of creating and shaping the electromagnetic fields which are used to accelerate the charged particles (I will call this accelerator physics). To design a particle accelerator of any quality, these two aspects must be considered in concert – however, for the purpose of explanation it is adequate to consider them separately, which is exactly what I will do. Because this physics

[§§§] According to the internet.

is of more obvious applicability to latter chapters of this thesis, I describe it in greater detail than the MRI physics section.

4.1. Beam dynamics

I. How are electrons generated?

Before we can accelerate electrons, we need some means of sourcing them. The most common way to generate electrons in accelerator physics is to extract them from a metallic cathode. To do this, energy must be supplied - the amount of energy required to extract a single electron from a cathode is called the *work function*. This energy can be in the form of a laser pulse (photocathode), a very strong electric field (field emitter) or heat (thermionic cathode) [9]. Thermionic cathodes are by far the most common cathodes in radiation therapy, and this includes all models used in this thesis. The process of extracting electrons from a thermionic cathode can be likened to 'boiling' electrons off the surface. The most commonly used thermionic cathodes for medical electron accelerators are *dispenser cathodes*, which consist of a matrix of tungsten impregnated with a Barium oxide mixture. The tungsten allows the cathode to become very hot without melting, whilst the barium mixture substantially lowers the work function [10].

II. What is a beam?

What exactly do we mean when we talk about beams? In fact, there is no clear cut definition of when a cloud of particles can be reasonably called a beam, but broadly speaking; a beam is a group of particles are all moving in the same direction, having similar energy, and low spread of transverse (i.e. perpendicular to the average motion) velocities.

Having (sort of) decided on what a beam is, what are the quantities which differentiate good beams from bad? Of course, this is application specific – but for medical applications, the primary factors are actually quite simple; mean beam energy, energy spread, and beam current [11]. Secondary considerations include the size and intensity profile of the beam, and the *emittance* of the beam. Emittance is quantity which essentially quantifies how chaotic a beam is – an ideal beam in which all particles were travelling parallel to each other with no transverse velocity would have an emittance of zero [12]. For time varying beams, the length of timing of each bunch in the beam (aka longitudinal emittance) is also important.

III. Relativistic kinematics and space charge

In order to design the electromagnetic fields which will be used for particle acceleration, one has to understand (1) the desired beam requirements, and (2) the dynamics of how electrons will move through the field. From a top level perspective, this is fairly easy – after all, we have the Lorentz equation (Equation 9), which gives us the force experienced by a particle in a given electromagnetic field – so if we know the starting position of the electron, and the electromagnetic field at all moments in space and time, we can fairly easily calculate the motion path of that electron. This is true, but practically, there are other considerations. Firstly, a beam of identically charged particles will repulse each other. This repulsive force is called *space charge*. Incorporating space charge into simple equations of motion is not at all

simple. Secondly, because energetic electrons very quickly approach the speed of light, one must incorporate the laws of special relativity into the motion calculations (relativistic kinematics). In addition to this, one is normally dealing with millions or billions of electrons, and each begins with a slightly different position and initial velocity – so for an exact solution, you would have to calculate the motion trajectory for each electron individually, whilst simultaneously considering the impacts of space charge from every other electron!

In the past, the complexity was dealt with by taking a statistical approach, and by making a number of simplifying assumptions, resulting techniques such as the paraxial approximation [13]. However, today it is more common to use computational codes to simulate the motion of charged particles through electromagnetic fields, and this is what is done in this thesis.

4.2. Accelerator physics and radiofrequency acceleration

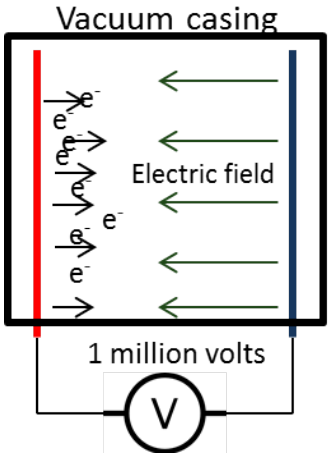

Figure 15: A very simple electron accelerator would simply apply the desired voltage between two electric plates. However, such a design is unsafe and inefficient for high energy beams. Note that force experienced by directions is the opposite direction to the electric field produced.

Suppose you want to produce electrons with energy of 1 MeV. How would you achieve this? The first thing you should note is that the *only* field you can use to do this is the electric field. This is not an obvious conclusion – based on the Lorentz equation (Equation 9), one might conclude that since both electric or magnetic fields apply a force on charged particles, either could be used for acceleration. However, whilst magnetic fields do apply a force, it is a centripetal force – it will change the direction, but not the kinetic energy of the particles. In fact, it is fundamental to electromagnetic theory that *magnetic forces do no work*. So, with this in mind – conceptually the simplest way we can accelerate our electrons to 1 MeV would be something like a parallel plate capacitor with a voltage gap of a million volts. (Figure 6). Electrons accelerated through this gap will have energy of 1 MeV – voila!

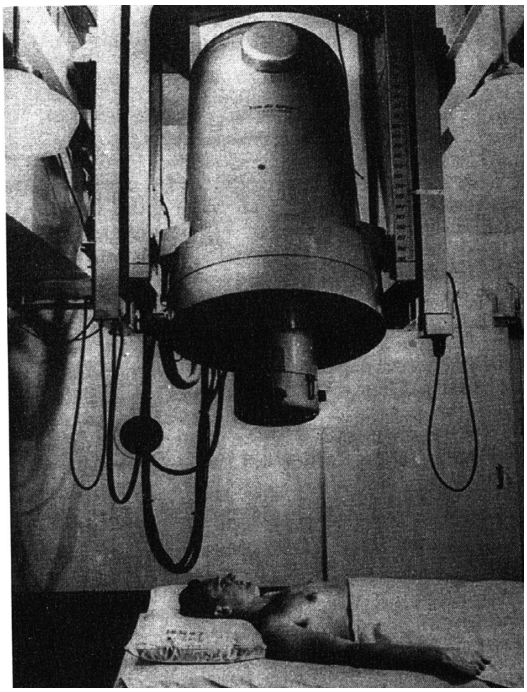

Figure 16: The Van De Graaf accelerator was one of the earliest megavoltage accelerators, and used a direction acceleration concept (Figure 3) to produce electrons of ~1 MeV [1].

The problem with this approach is that 1 million volts is an awful lot (remember the lightning example??). As such, our conceptual accelerator would be very difficult to build and quite dangerous. In fact, this is exactly how the earliest MV electron accelerators worked [1, 14] such as Figure 4, which shows a Van De Graaf

generator[****] – the scale of such a design is massive[††††]! Today, this direct acceleration approach is still used for relatively low energy (kilovoltage) electron beams such as diagnostic x-ray tubes (~100 KeV) [15]. However, for higher energies, a different approach is used – RF acceleration. The basic principle of RF acceleration is to use the electric field component of electromagnetic radiation (Figure 3) to push electrons forward. There are a number of advantages to using RF instead of DC acceleration. Firstly, electrical hazards risks associated with creating strong electric fields are minimised, as RF is quite happy propagating through space and does not seek out ground like a DC voltage. Secondly, one does not ever have to create a voltage gap of more than a few hundred kilovolts to produce fields of up to GeV energies. Thirdly, at higher frequencies, the electric breakdown limit increases, meaning an RF cavity supports a higher electric field than a DC cavity [16]. Finally, by using RF of a suitable frequency, accelerators can be made much smaller than their DC counterparts. The fundamental principles necessary to grasp RF acceleration are laid out below.

I. Resonance and resonant microwave cavities

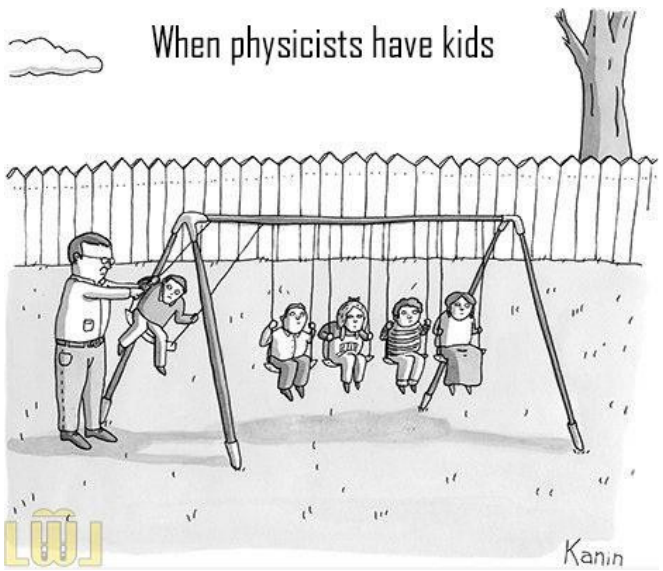

Figure 17: A weakly coupled system of 5 resonators.

Crucial to the principle of RF acceleration is the concept of resonance and resonant cavities. *Resonance* is a very broad term, which essentially describes a state of affairs where energy is added to a system at such a frequency that the net energy of the system grows with each addition until reaching a steady state. A practical everyday example of this rather esoteric sounding definition is pushing a child on a swing; you must time the frequency of your pushes quite precisely in order to cause the net energy of the child/swing system to grow (and to avoid being hit in the face by an angry child). An everyday example of a *resonant cavity* is an empty bottle. When you blow across the rim of a bottle, it begins to make noise – it *resonates.* In this case, the cavity is being filled up with acoustic energy. The *frequency* with which the system resonates is dependent on the size and shape of the bottle[‡‡‡‡]. If you have strong enough lung capacity, you will find that by blowing hard enough, you can force the bottle to resonate at a higher frequency. This simple example illustrates an extremely important principle across a huge range of physical systems; for a

[****] Today, much smaller versions of this device are used to make students hair stand on end at high school science demonstrations.
[††††] Interestingly, this is a bit what medical proton and carbon accelerators look like today – we can only hope that future technologies will improve these as much as it has electron accelerators!
[‡‡‡‡] And the medium the bottle is filled with, but I'm assuming this is not a variable.

given set of boundary conditions, a system resonates at discrete frequencies[§§§§] [12, 17].

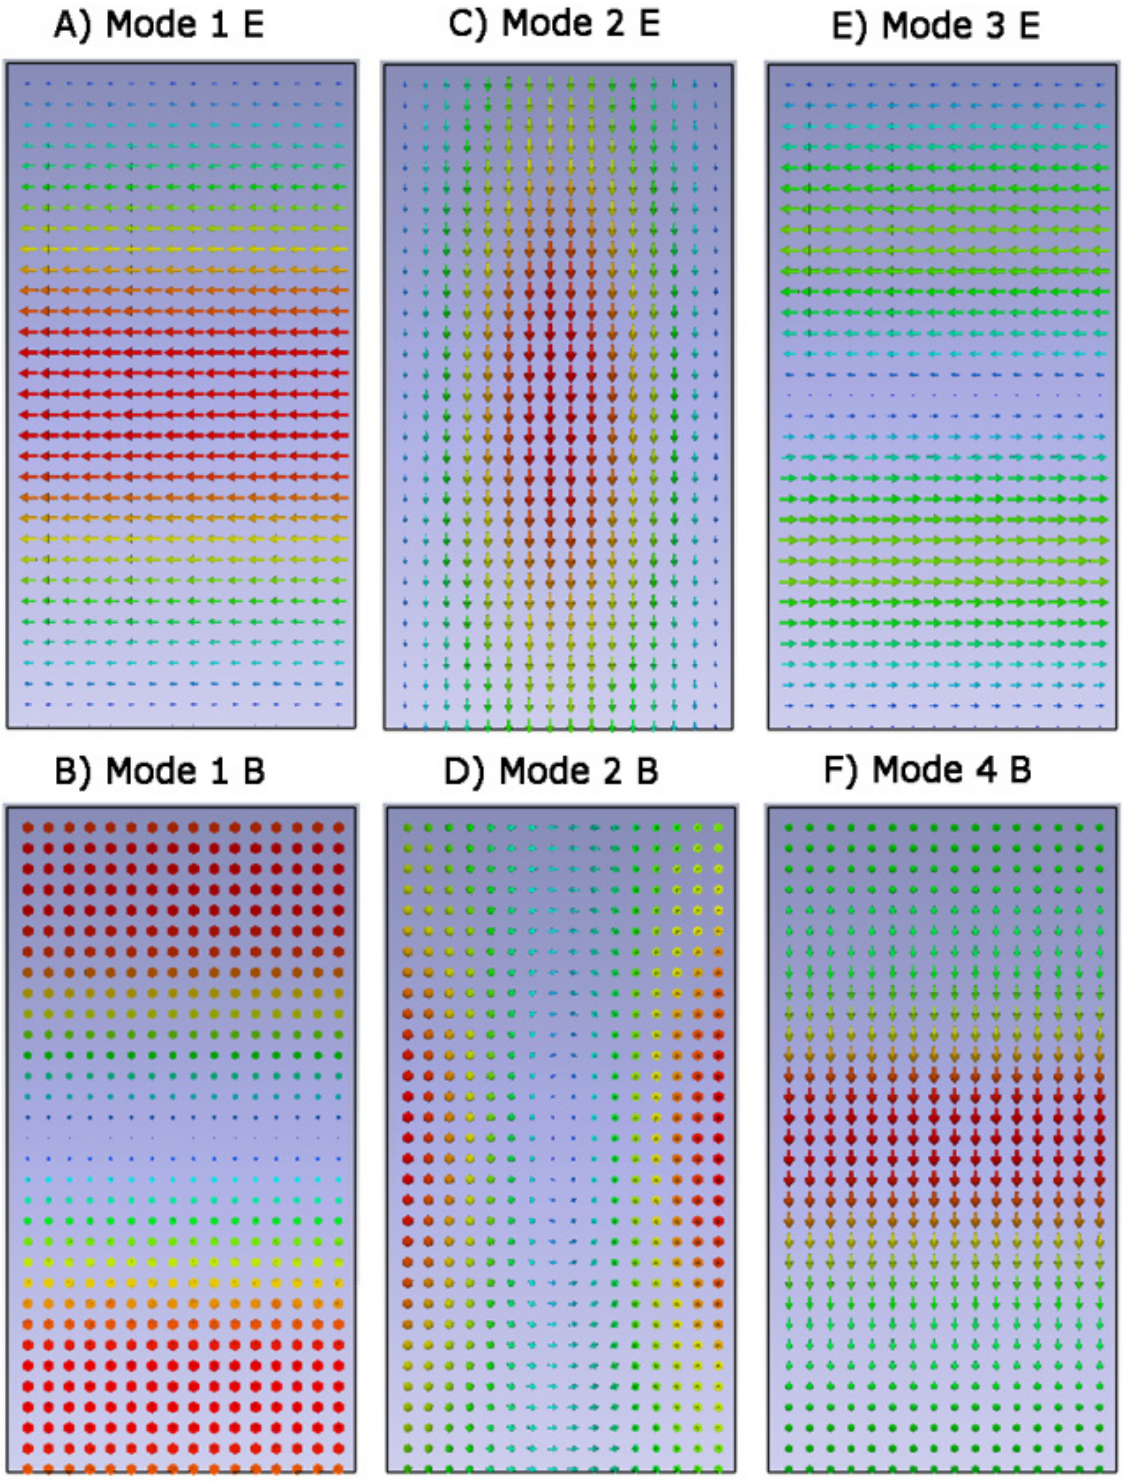

Figure 18: The electric (top row) and magnetic (bottom row) field distributions of the first three resonant modes of a simple cylindrical microwave resonant cavity. Typically for particle acceleration, we would use the first mode, as the particle can pass through a strong electric accelerating field and minimal magnetic field. (note that in some panels, i.e. B) the field is pointing in or out of the page).

[§§§§] Although I have never had any success at replicating this principle with a child on a swing.

An RF resonant cavity (also often called microwave resonant cavity) is direct analogy to this example. Instead of filling up a cavity with acoustic energy, the system fills up with RF energy. Just like the bottle, the frequencies at which the system resonates depend on the boundary conditions (the size, shape, and material properties of the cavity). The different frequencies at which a given cavities resonates are called the *modes*. Figure 9 shows the electromagnetic field distribution of the first three modes of a simple cylindrical cavity. Each mode has a different field distribution associated with it. As such, only some of the modes are useful for accelerating particles.

II. Using microwave resonance cavities for particle acceleration

Consider the electromagnetic field distributions shown in Figure 9. Of these three modes, the first is the most attractive for accelerating a charged particle – a particle passing through the horizontal axis of this cavity will experience an almost parallel electric field, and very little magnetic field (the technical name for this made is the TM_{010} mode [12]). In fact, an electron passing through the central axis of this cavity would experience a field environment quite like that of our simple parallel plate accelerator (Figure 6)! However, in an RF cavity there is an additional complication; the field is changing as the particle moves through it. Figure 10 shows the electric field distribution for the TM_{010} mode at various moments in time. It can be appreciated that if the field changes too fast (or the particle is too slow), then the field direction will reverse before the particle has passed through the cavity, meaning the electron will lose energy or even be accelerated in the wrong direction. From this simple example, it is easy to see that there must some optimal length for this cavity[*****]. Finding this length would be the first step in optimising the cavity for particle acceleration. After that, we might begin to examine the shape of the cavity such that the field distribution can be more optimally shaped; of course, we would also have to add some holes in either end for electrons to get in and out. Other considerations in the design of accelerating cavities include minimising electromagnetic heating (which can lead to thermal deformation) and electric field breakdown. A good way to avoid problems in both cases is to ensure that there are no sharp edges inside the cell, as these cause high local electromagnetic fields. Once all these factors are taken into account, accelerating cells normally end up looking a little different than the simple example shown in Figures 9 & 10. More typical accelerating cell geometry is shown in Figure 11[†††††]. Despite the difference in shape, if you understand the process of acceleration with the simple cylindrical cells, you have essentially already grasped the key

[*****] For electron accelerators, we normally just assume that electrons travel at the (constant) speed of light, which is close enough that we normally get away with it! The other implication to note at this stage is that RF accelerators will *always* produce bunches of accelerated particles rather than a steady stream.

[†††††] This is not to imply that they all look like this, although this a very common design.

concepts of using microwave resonance cavities for particle acceleration.

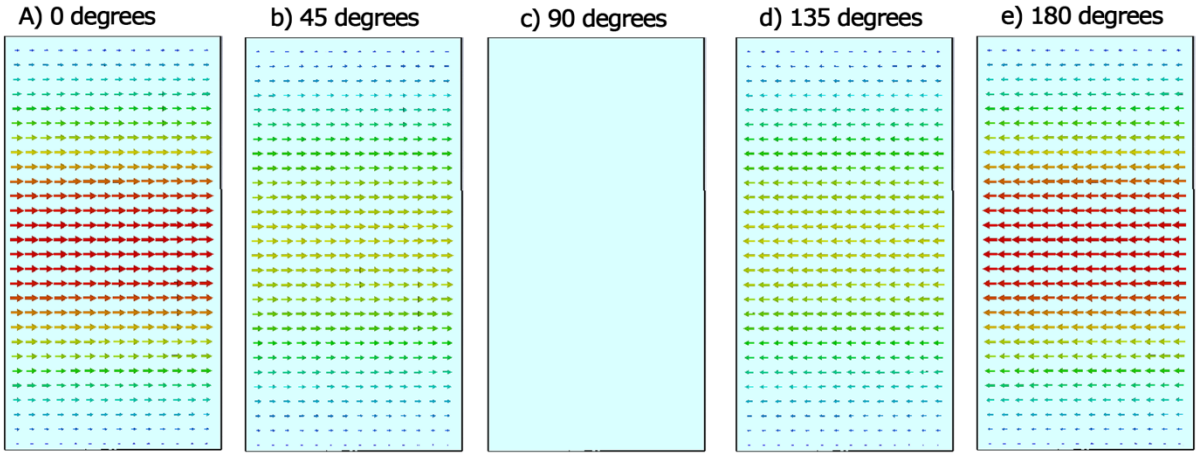

Figure 19: The electric field in the first resonant mode of a cylindrical cavity. The first pane; (A) is equivalent to panel A in Figure 9 – the next four panels represent the field distribution at various phases of the oscillation. After it has reached 180 degrees, the process reverses.

III. Supplying power to RF cavities and impedance matching

Thus far, I have talked about RF cavities being filled with energy, without actually explaining how it got there. In this section, I will (extremely briefly) review the major methods of generating RF power, transporting it to a cavity and filling that cavity up. There are two main devices which are used to generate RF power, called a *magnetron* and a *klystron* [12]. Although there are important theoretical and practical differences in the way these two devices operate, for the purpose of this section, the similarities outweigh the differences. Therefore, to avoid repeated references to 'klystrons or magnetrons' below, I am simply going to refer to klystrons - however the top level overview could be equally well applied to magnetrons.

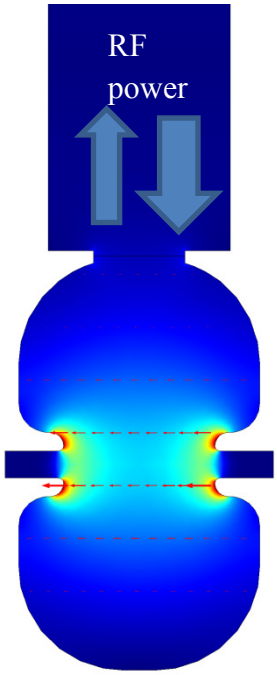

Figure 20: an input RF power waveguide is coupled to an accelerating cavity. RF power flows both in and out of the cavity; inside the cavity power is dissipated in the cavity walls (and beam if there is one).

A klystron can be thought of as a 'particle accelerator in reverse'. In an RF particle accelerator, energy is transferred from resonant cavities to a beam; in a klystron, energy is transferred from a beam of particles to resonant cavities. This sounds rather counterintuitive – our final aim is to accelerate particles, and yet we start by *decreasing* the energy of particles?? A useful everyday analogue to this process is a voltage transformer that transforms high current and low voltage, into low current and high voltage (or vice versa). A klystron does something similar – we start with a beam of fairly low energy and high current, and end up with a beam of lower current but higher energy. RF power must be transported using a *waveguide* – a simple wire won't do it. A waveguide is actually quite similar to an RF cavity; whereas in the RF cavity RF radiation is *stored*, in a transport waveguide it travels down the waveguide.

There is a real art (or more accurately science!) to coupling the

travelling waveguide to the resonant cavity. In very crude terms: we simply do this by having a hole in the cavity which is attached to the travelling waveguide (Figure 11). It can be easily appreciated that the size of this hole is very important – if it is too small, then no power can get into the cavity, whilst if it is too large the cavity will lose its ability to store energy (too much power will leak out). We refer to these scenarios as being *undercoupled* and *overcoupled* respectively. There will of course be some optimally sized hole where the power travelling into the waveguide and the power 'leaking' back out are well balanced – we call this situation *matched*[‡‡‡‡‡]. In accelerator physics, *the cavities should always be made overcoupled.* In the absence of any beam, this would be suboptimal – however, when a beam is passing through the cavity it absorbs some of the energy in that cavity thereby changing the coupling physics. This is why accelerator systems should always be designed to be overcoupled – in order to compensate for the presence of the beam. The process of matching the power source to the resonant cavity is called *impedance matching.* Obviously in real world operation, it is crucial to have an import RF port. However, for the purpose of explaining physics and first order simulations, we can often simply assume that the cavities have energy in them without considering how it got there.

IV. Coupled RF cavities

Well machined resonant RF cavities can support very high electric fields – much higher than could be safely produced over the same distance using DC acceleration. Even so, a single RF cavity cannot support a field strong enough to accelerate electrons to the energies desired for medical applications (and certainly not for high energy physics). The obvious solution to this problem is to use more than one RF cavity. There are two ways one could utilise multiple RF cavities; the most immediately obvious way is simply by simply lining a number of them up in a row. Whilst this approach is used in specialised cases, there are some problems with it. Firstly, introducing an RF power port into the structure will reduce the optimal shunt impedance of each cell and introduce asymmetries into the beam. Secondly, powering each cell individually like this increases the possibility that the phase between cells becomes asynchronous (i.e. that electrons can arrive at a cell at the wrong time). As such, in medical electron accelerators a different approach is taken; individual RF cavities are coupled to each other in such a way that the entire structure is resonant. Power is then supplied to only one cavity, and from this 'driven' cell, all the other cells eventually fill up with energy (the time constant governing this process is called the 'filling time'). The advantage of this approach is that the phase of the field in each cell is fundamentally required to be correct (assuming you designed the structure properly in the first place) and the non-driven cells can have a higher shunt impedance.

How exactly does one 'couple' two cells together? In fact, it's quite easy! Simply by creating a path for the electromagnetic field to travel through (I'll call this a slot). Note that this concept is very similar to coupling an input RF feed into a resonant cell. In order to

[‡‡‡‡‡] Impedance matching of RF cavities is quite a complex idea to get one's head around (I have 'dumbed it down' very considerably here!). Consider a mechanical example: imagine trying to hit a tennis ball with a baseball bat, or a baseball with a tennis racquet. In both instances, you would be quite unsuccessful. This illustrates the very general principle that when transferring power from a source to a load, they must be *matched*, and so it is with RF cavities.

understand the physics of coupled resonators, let us first consider a mechanical analogue – a system of three coupled pendulums (Figure 12). There are three ways in which this system can oscillate (1) all pendulums move together (2) the middle pendulum is stationary whilst the outer two pendulums move in opposite directions (3) each pendulum is moving in the opposite direction to its neighbour. Each of these cases will also have a slightly different resonant frequency – this phenomenon is called dispersion. If we consider coupled accelerating cells, it can be seen that choosing the correct mode of operation is very important, as this will affect the field an electron experiences upon arriving at a given cell. Although various modes can be used in accelerators, I will focus on case 2 (middle pendulum doesn't move). This is called the $\pi/2$ mode (as there is a 90 degree phase shift between each oscillator) and is one of the most commonly used modes in medical electron accelerators. The size, shape, and location of the slot determine the strength of the coupling between the two cells.

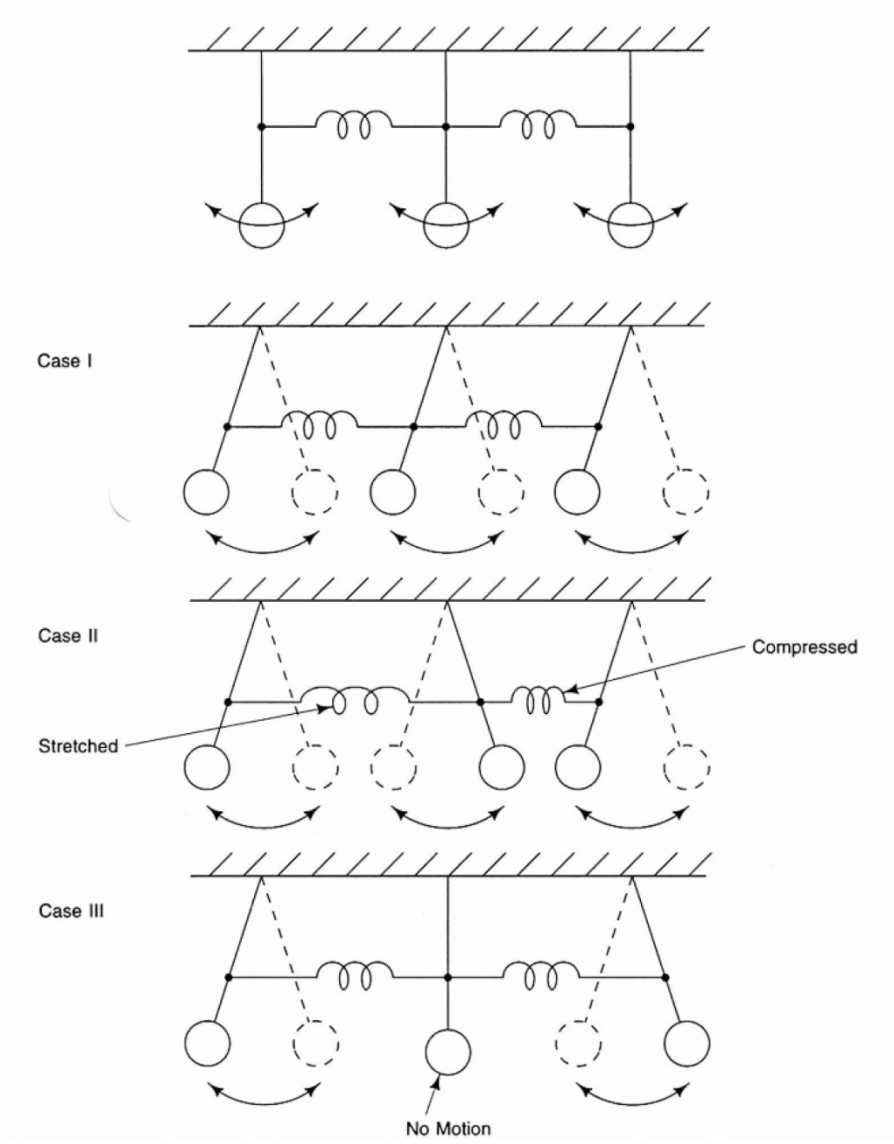

Figure 21: the behaviour of coupled microwave cavities can be understood by considering the analogy of three coupled pendulums. Whereas a single pendulum has a single resonant frequency[§§§§§], the coupled system has three different resonant modes, each with a slightly different resonant frequency. Figure taken from [12].

The strength of coupling between cells (called the *intracavity coupling constant,* and essentially determined by the size, shape, and location of the slot) is very important. It effects the filling time of the structure, and the frequencies of each mode on the dispersion curve (consider the spring analogy again – as the springs becomes stiffer, the separation of three dispersive frequencies increases). In general, one wants to minimise filling time and increase the mode separation (to avoid the accidental excitation of undesired modes). Both requirements are achieved by increasing the intra-cell coupling factor – but there is a catch! As this factor is increases, the shunt impedance of each cell will go down. As such, some consideration must be put into choosing an optimal value of the coupling strength between cells.

[§§§§§] As (apparently) first noted by a young Galileo when studying a swinging chandelier in a church.

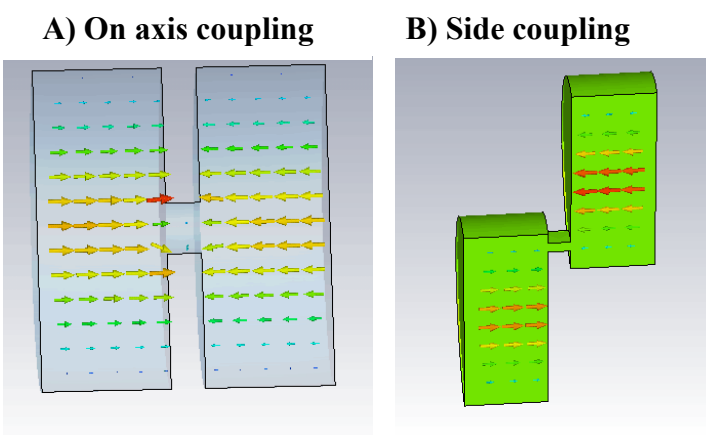

Figure 22: Two cylindrical cavities can be coupled together in different ways A) On axis coupling and B) side coupling. . The vector plot shows the electric field. Because on axis coupling places the slot in a region of high electric field, it can also be called electric coupling; similarly, side coupling can also be called magnetic coupling (see Figure 9).

If we desired to couple two pillbox cavities (Figures 9 & 10) together, there are two ways we could do it; by putting a coupling slot in a region of strong electric field (near the centre of the cavity) or by putting a coupling slot in a region of strong magnetic field (near the edges of the cavity). These two approaches are shown in Figure 13. The second method is particularly attractive for accelerator physics as a bunch of electrons has to travel through the centre of the cavity (and there must be a beam passage in and out of it). By placing the coupling slots near the side of the cavity, the central part of the cavity can be shaped to be optimal for electron acceleration, whereas if the slot is placed in the middle, then in order to allow adequate coupling the beam passage must be made larger than would otherwise be necessary, reducing the shunt impedance of the cell. By operating in the $\pi/2$ mode every second cell will hold on energy (e.g. the middle pendulum does not move in case 3 in Figure 12), and as such it does not matter that every second cavity is off the axis of the beam direction. In fact, this is a distinct advantage since it means the total length of the accelerator can be reduced compared to the on axis coupling! When one puts all these pieces together (optimal shape of the accelerating cell, off axis coupling cells, RF power to one cell) a final accelerator structure may look like that shown in Figure 14.

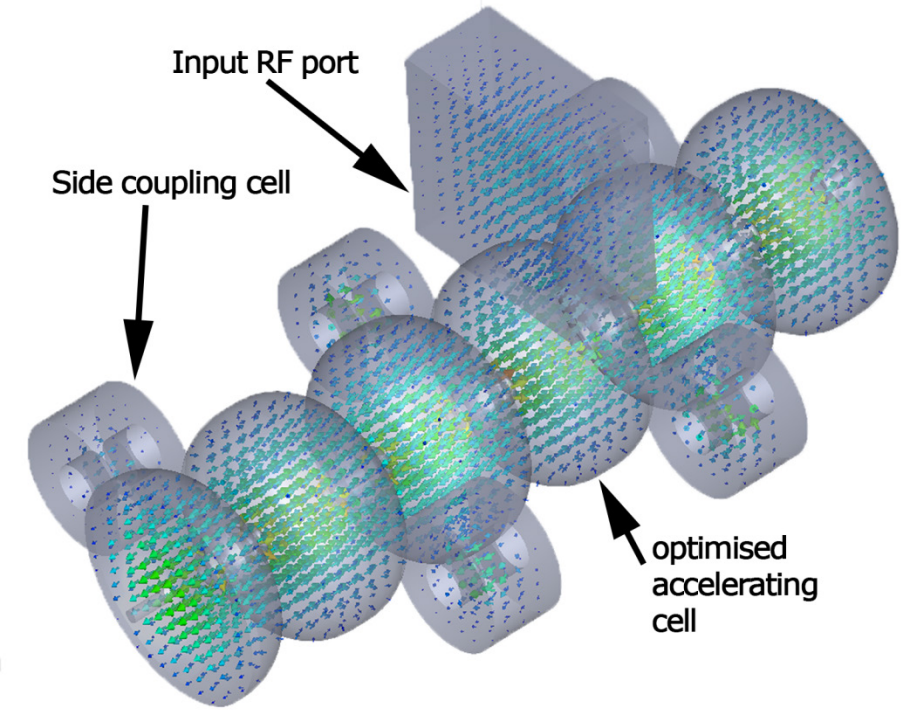

Figure 23: An example of a short medical accelerator structure. The coloured arrows show the electric field distribution.

5. Interactions between MRI and Linear accelerators

Having given a top level overview of the independent operation of MRI scanners and linear accelerators, we can shift our focus back to the main topic of this thesis; the integration of these two devices. The fundamental issue with integrating MRI-Linac systems is that when placed in close proximity, the two devices become electromagnetically coupled – that is, each machine begins to influence the operation of the other! One of the luminaries in this field, Gino Fallone, has an excellent way to describe this situation – he says the two machines are 'allergic' to each other. The exact ways in which the two machines can influence each other are outlined below.

I. Impact of magnetic fields on linear accelerator

When a linac is operated in the fringe field of an MRI scanner, it is subjected to external magnetic fields which produce a force on the accelerated electrons. The direction of this force is dependent on both the magnetic field and electron motion direction. As such, the relative orientation of the linear accelerator with MRI scanner becomes important. There are two sensible orientations – the in-line configuration, in which the electrons are accelerated in the same direction as the magnetic field, and the perpendicular configuration, in which they are accelerated perpendicular to this field (Figure 15). Both of these configurations are being developed by different groups, as outlined in section 6. Each configuration has unique advantages and disadvantages associated with it. However, if one considers particle acceleration in isolation, then the in-line configuration is indisputably the superior option. This is because magnetic force is minimised when charged particles are travelling in the same direction as the magnetic field lines (Equation 6). The effects of both in-line and perpendicular magnetic fields on linear accelerator operation have previously been studied via computational simulations. For the perpendicular case, total beam loss occurred at 14 G, and 45 % at 6 G [18]. This means that in order to produce a treatment beam for the perpendicular orientation, the linac must be operated in a near zero field environment. This can be achieved by modifying the magnet and magnetically shielding the linac [19, 20] – however, MRI magnet design (and redesign) is not a trivial task, and magnetic shielding causes distortion in the scanner (section 5. IV). For the inline case, maximum beam loss was 79 % at a field of 600 G [21]. It was also shown that the effect of magnetic fields on the electron accelerator is nearly entirely on the electron gun – that is, operation of the accelerating cavities is largely unaffected for the in-line case [22, 23]. Chapters 3 and 5 of this thesis examine electron transport within the fringe fields of MRI scanners.

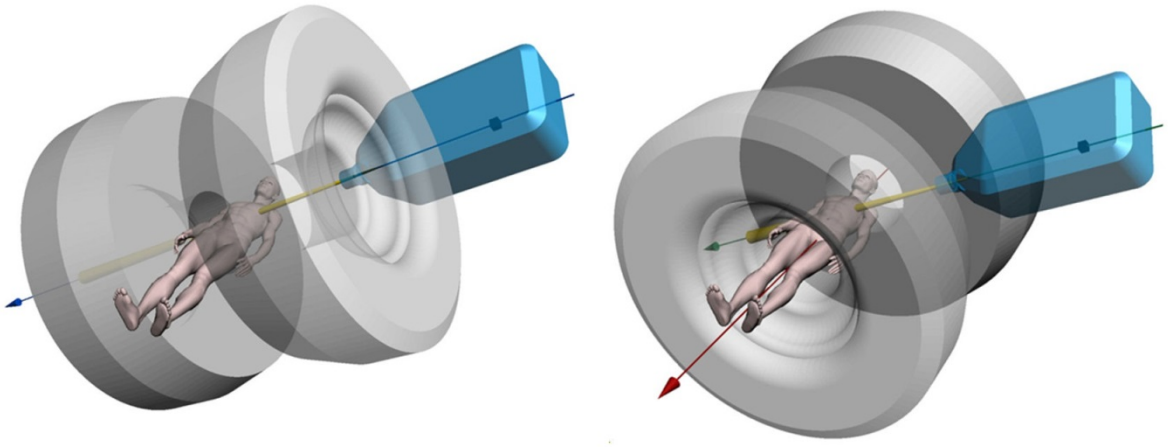

Figure 24: Conceptual layout of a linac (blue) with respect to an MRI scanner. There are two possible locations of the linac (A) the in-line orientation, in which the fringe field from the magnet points in the same direction as the electron being accelerated, and (B) the perpendicular configuration, in which the electron motion and fringe field are at right angles. Figure taken from [24].

II. Impact of magnetic fields on the multi leaf collimator (MLC)

The multileaf collimator is a crucial part of the treatment beam apparatus, used to dynamically shape and modulate the photon treatment beam such that a complex dose distribution can be delivered to the patient. The MLC consists of a number of leaves which move independently. In order to accurately control the position of each leaf, a magnetic encoder is used to verify that each controlling stepper motor has taken the correct amount of steps. The function of this magnetic encoder is compromised when it is placed in an external magnetic field. Motor performance in magnetic fields has been assessed using independent optical encoders. It was found that depending on the specific motor used, performance began to degrade from 450 Gauss (1 G = 10^{-4} T) [25]. Therefore, to operate a conventional MLC in an MRI-Linac system, the MLC must be moved far away from the MRI scanner that the field is less than 4.5e-2 T, or magnetic shielding must be utilised. Note however that magnetic shielding has its own issues (section 5. IV).

III. Impact of magnetic fields on dose transport and dosimetry

Since the dose in MRI-Linac systems is delivered by photons, which have no charge, it might be supposed that the treatment beam is not affected once it has collided with the target to produce the bremsstrahlung photon beam. However, this is not correct – even in photon therapy, it is still electrons delivering the dose! How does this work? Something like this: a beam of photons cause ionisation in the matter they traverse, resulting in *free radicals* (electrons and ions). The more mobile electrons, go onto to deposit dose, and are subject to the Lorentz force defined by equation 9. When the patient is placed in a strong magnetic field, the trajectories of the electrons are altered – whereas in zero field, the electrons undertake a 'random walk'[******] from their point of creation, and in the presence of magnetic fields, the direction they take is biased. Again, the exact behaviour depends on the relative orientation of the radiation beam and the magnetic field. When the beam and the magnetic

[******] But not entirely random; there will be a direction preference to satisfy conservation of momentum.

field lines are in (roughly) the same direction, the effect is to reduce lateral scattering – which is not necessarily a bad thing, since it allows for dose to be more focused along the photon beam. In fact, before anyone thought of building an MRI-Linac, the theoretical benefits of a strong longitudinal fields on dose transport within the patient had been studied in some detail [26]. Recent developments towards MRI guided radiotherapy have prompted further quantification of this effect [27-29]. In general, it seems that the presence of longitudinal magnetic fields has a positive impact on dose distributions around the treatment volume – however there is another problem. The same mechanism that enhances the dose to tumours also acts to focus dose to the patient surface, resulting in unwanted hot spots to the patient's skin. Hot spots of up to 1000% have been observed in Monte Carlo simulations [30]! However, these hotspots can be greatly reduced through the use of electron purging devices and optimisation of the magnetic fringe field [31, 32].

In contrast, when the magnetic fields are perpendicular to the beam direction, skin dose is not a problem as most contaminant electrons are swept away before reaching the patient. However, in this instance other issues emerge – the dose distribution inside the patient begins to become skewed, and at tissue air interfaces (e.g. lung) hot spots can appear due to the much storied 'electron return effect', which has been shown to result in dose differences of up to 30-40% [27, 33]. It has also been shown that there are various ways to compensate for this effect – for instance, by using an opposed beam to compensate for the affect, and by including the effect of the magnetic fields in the inverse plan optimisation [34, 35]. Such approaches do need to be tested in the presence of motion (remembering the fact that the target geometry may move is one of the main motivators for developing an MRI Linac in the first place!). Whilst some early promising work on this has been undertaken [34, 36, 37], further quantification of these effects in the presence of motion is needed.

IV. Impact of magnetic fields on the MRI scanner

MRI scanners require an extremely homogenous magnetic field in order to function correctly (section 3). Any magnetic field external to the scanner introduces distortion to this field, which can compromise the geometric accuracy and image quality of images returned by the scanner. There are two ways in which external magnetic fields could be generated within the imaging volume of the MRI scanner by a linear accelerator. The first is the same way in which the MRI scanner itself generates a magnetic field – through current carrying wires. Although there are plenty of current carrying wires (and currents in vacuums) in a linac, it turns out these are nowhere near strong enough to impact on normal imaging performance. The second way is the most familiar example of magnetism – magnetic materials (Section 2.II). When magnetic materials are placed in a magnetic field (such as exists in an MRI scanner), they become magnetised, and as such generate their own magnetic field. A Linac tends to use a lot of ferromagnetic materials – for instance bending magnets, MLC motors, ion pumps, magnetrons, ferrite recirculates etc. [12]. The impact of most of these components has not been studied or quantified in any detail yet, however since most of them are small, and can be moved quite far away from the accelerator, it is broadly assumed that impact of these components is negligible. The most concerning component is the MLC, since it is the closest component to the magnet, and it contains the largest mass of metal. As such, this is

one of the few components which has been studied in detail [38]. The results of this study showed that a standard MLC would have minimal impact on the imaging field, provided it was operated more than ~1 meter from the centre of the magnet.

Based on this information, it might seem that we don't have to worry too much about the Linac distorting the MRI field – but there's a catch. The easiest way to minimise the impact of the MRI fields on the linear accelerator (IV, V) is to magnetically shield the sensitive components. Shielding can be either passive (using magnetic materials to reduce the field in a volume of interest) or active (using current loops). But magnetic shielding, by definition, distorts magnetic fields! This is where a lot of the difficulty and compromise in MRI-Linac design stems from. We have to carefully trade off the improved linac performance magnetic shielding can give us versus the degradation in the magnetic field of the MRI scanner it causes. Magnetic shielding of an electron gun in the in-line configuration of a 0.5 Tesla magnet has previously been studied; it was found that it was quite straightforward to shield the electron gun without substantially distorting the MRI imaging volume [20]. However, the efficacy of shielding for other linac parts, other magnets, and other orientations has yet to be studied. Magnetic shielding is the subject of chapter 4 of this thesis.

V. Impact of magnetic fields on RF power flow

Magnetic fields can impact on the flow of RF power down a travelling waveguide (for example Figure 11) via the Faraday affect. The faraday affect refers to the rotation of polarized light via magnetic fields; since the travelling waveguides only support certain modes, rotation of this distribution can lead to decreasing RF power flow. As yet, there are no reports in the literature quantifying this effect in MRI linac systems, although it has been noted as a potential concern in conference abstracts [39].

VI. Impact of RF noise on the MRI scanner

As described in section 3, MRI scanners produce images by the measuring the RF signal of protons in the body. As such, MRI scanners are very sensitive to RF noise. A linear accelerator emits a wide spectrum of RF noise – traditionally not a problem but potentially a big one for MRI-Linac systems! The RF noise spectrum from three clinical linacs has been measured and found to be comparable in power and frequency to that measured in MRI, as such there is the potential of image artefacts occurring if correct shielding is not used [40]. However, multiple studies and experimental results have demonstrated that it is possible to decouple the RF signals from the two devices to the extent that simultaneous operation is not affected [41-43]. In addition to the study on RF noise from Linacs above, the noise from MLC's and modulators has been quantified. It was shown that it is straightforward to shield the RF noise from the MLC [44], and that the largest source of RF noise from the linac was the magnetron. As such, it was recommended that the magnetron be stored in a different RF cage to the MRI scanner. [45]. Whilst RF decoupling is an important design consideration, it is the *magnetic* decoupling which poses the greatest design challenges for MRI-Linac integration. In essence, this is because any efforts to change the magnetic field (i.e. magnetic shielding) impact on magnet homogeneity; this is not the case with RF shielding.

6. Proposed MRI Linac systems

Having heard at some length about the arguments for (and problems with) MRI-Linac integration, the reader is no doubt left wondering only where best to purchase such a device!

There are currently three academic groups worldwide working on MRI Linac devices (including our own): The University of Utrecht/ Phillips/ Elekta, the University of Alberta, and the University of Sydney. All of these systems were recently outlined in some detail in a recent edition of Symposiums in Radiation Oncology [43, 46, 47]. There is also a fourth group (ViewRay) which has developed and bought to market a solution based on a radioactive source [48]. Replacing the accelerator with a cobalt source does solve many of the problems arising from the magnetic coupling of the two devices. However, in general cobalt based radiation therapy is generally considered an inferior treatment for a wide range of reasons (to be truthful, perhaps a little unfairly in many cases! [49]). Recently, the same commercial group has developed a linac based solution from which their earlier device can be upgraded [50] – this device is currently awaiting FDA approval and will probably be first to market with an MRI-Linac system.

There are a variety of differences between the differences proposed systems. Probably the two most obvious are (1) field strength of the MRI (ranging from .5 T to 1.5 T), and the relative orientation of the linac and MRI field (touched on already in 1.1IV). Less obvious, but equally important is the fringe field of different MRI magnets. As discussed in section 5, this factor will dominate the extent of the electromagnetic coupling of the two devices. The differences between different systems are summarised in Table 1 whilst Table 2 lists the advantages and disadvantages of the in-line and perpendicular orientations (Figure 14).

Table 1: A summary of the MRI-Linac systems which are currently under development.

Group	MRI field strength	Orientation	Treatment source & energy	Website
Australia	1.0 Tesla	In-line	Varian Linatron, 6 MeV	http://sydney.edu.au/medicine/radiation-physics/research-projects/MRI-linac-program.php
Utrecht/ Elekta/ Phillips	1.5 Tesla	Perpendicular	Elekta Linac, 6 MeV	https://www.elekta.com/mr-linac
U. Alberta/ Magnetx	0.56 Tesla	In-line	Varian 600C, 6 MeV	http://www.mp.med.ualberta.ca/linac-mr/
Viewray	0.35 Tesla	Perpendicular	Currently available: cobalt, Future system: 6 MeV Linac	http://www.viewray.com/

Table 2: Advantages and disadvantages of in-line versus perpendicular configuration (see Figure 14). This table is taken from [47].

Advantages of the Inline Approach	Advantages of the Perpendicular Approach
No beam attenuation and Compton scatter to the patient from irradiation through the cryostat (if closed bore)	More similar design to mass-produced conventional MRI systems (if closed bore)
Less effect of the B field on electron gun operation	Lower constraints on magnet, gradient coil, and RF design, resulting in higher potential imaging performance and higher B field (if closed bore)
Less effect of the B field on waveguide operation	Lower skin dose
Less effect of the B field on electron transport within the patient: sharper penumbra and no electron return effect	No need to rotate the magnet or the patient
Lower exit dose	
Linac fixed with respect to the magnet. This reduces the need to manage eddy currents or dynamic shimming requirements, where the linac moves with respect to the magnet	

Abbreviation: RF, radiofrequency.

References

1. Suit, H.D. and J.S. Loeffler, *Radiation Oncology and the MGH 1896–1945*, in *Evolution of Radiation Oncology at Massachusetts General Hospital*. 2011, Springer. p. 13-27.
2. Griffiths, D.J. and R. College, *Introduction to electrodynamics*. Vol. 3. 1999: prentice Hall Upper Saddle River, NJ.
3. Jackson, J.D., *Classical electrodynamics*. 1999: Wiley.
4. Geva, T., *Magnetic resonance imaging: historical perspective.* Journal of Cardiovascular Magnetic Resonance, 2006. **8**(4): p. 573-580.
5. Liney, G., *MRI from A to Z: a definitive guide for medical professionals*. 2010: Springer Science & Business Media.
6. McRobbie, D.W., E.A. Moore, M.J. Graves, and M.R. Prince, *MRI from Picture to Proton*. 2006: Cambridge university press.
7. Bushberg, J.T. and J.M. Boone, *The essential physics of medical imaging*. 2011: Lippincott Williams & Wilkins.
8. Barbalat, O., *Applications of particle accelerators*. 1994, CERN.
9. Hernandez-Garcia, C., P.G. O Shea, and M.L. Stutzman, *Electron sources for accelerators.* Physics today, 2008. **61**(2): p. 44.
10. Cronin, J. *Modern dispenser cathodes*. in *IEE Proc*. 1981.
11. Sheikh-Bagheri, D. and D. Rogers, *Sensitivity of megavoltage photon beam Monte Carlo simulations to electron beam and other parameters.* Medical physics, 2002. **29**(3): p. 379-390.
12. Karzmark, C., C.S. Nunan, and E. Tanabe, *Medical electron accelerators*. 1993: McGraw-Hill.
13. Pierce, J.R., *Theory and design of electron beams*. 1954: van Nostrand.
14. Robison, R.F., *The Race For Megavoltage X-Rays Versus Telegamma.* Acta Oncologica, 1995. **34**(8): p. 1055-1074.
15. Sprawls, P., *Physical principles of medical imaging*. 1987: Aspen Publishers.
16. Loew, G.A. and J. Wang, *RF breakdown studies in room temperature electron linac structures*. 1988.
17. Serway, R.A., C.J. Moses, and C.A. Moyer, *Modern physics*. 2004: Cengage Learning.
18. St Aubin, J., S. Steciw, and B. Fallone, *Effect of transverse magnetic fields on a simulated in-line 6 MV linac.* Physics in medicine and biology, 2010. **55**(16): p. 4861.
19. Overweg, J., B. Raaymakers, J. Lagendijk, and K. Brown, *System for MRI guided radiotherapy.* Proc. Int. Society Magnetic Resonance in Medicine, 2009: p. 593.
20. Santos, D., J.S. Aubin, B. Fallone, and S. Steciw, *Magnetic shielding investigation for a 6 MV in-line linac within the parallel configuration of a linac-MR system.* Medical physics, 2012. **39**(2): p. 788-797.
21. St .Aubin, J., D. Santos, S. Steciw, and B. Fallone, *Effect of longitudinal magnetic fields on a simulated in-line 6 MV linac.* Medical physics, 2010. **37**(9): p. 4916-4923.
22. St. Aubin, J., D. Santos, S. Steciw, and B. Fallone, *Effect of longitudinal magnetic fields on a simulated in-line 6 MV linac.* Medical physics, 2010. **37**(9): p. 4916-4923.
23. Constantin, D.E., L. Holloway, P.J. Keall, and R. Fahrig, *A novel electron gun for inline MRI-linac configurations.* Medical physics, 2014. **41**(2): p. 022301.
24. Constantin, D.E., R. Fahrig, and P.J. Keall, *A study of the effect of in-line and perpendicular magnetic fields on beam characteristics of electron guns in medical linear accelerators.* Medical physics, 2011. **38**(7): p. 4174-4185.
25. Yun, J., J.S. Aubin, S. Rathee, and B. Fallone, *Brushed permanent magnet DC MLC motor operation in an external magnetic field.* Medical physics, 2010. **37**(5): p. 2131-2134.
26. Bielajew, A.F., *The effect of strong longitudinal magnetic fields on dose deposition from electron and photon beams.* Medical physics, 1993. **20**(4): p. 1171-1179.
27. Kirkby, C., T. Stanescu, S. Rathee, M. Carlone, B. Murray, and B. Fallone, *Patient dosimetry for hybrid MRI-radiotherapy systems.* Medical physics, 2008. **35**(3): p. 1019-1027.

28. Mahdavi, S., A. Esmaeeli, M. Pouladian, A. Monfared, D. Sardari, and S. Bagheri, *Breast dosimetry in transverse and longitudinal field MRI-Linac radiotherapy systems.* Medical physics, 2015. **42**(2): p. 925-936.
29. Kirkby, C., B. Murray, S. Rathee, and B. Fallone, *Lung dosimetry in a linac-MRI radiotherapy unit with a longitudinal magnetic field.* Medical physics, 2010. **37**(9): p. 4722-4732.
30. Oborn, B., P.E. Metcalfe, M. Butson, A.B. Rosenfeld, and P. Keall, *Electron contamination modeling and skin dose in 6 MV longitudinal field MRIgRT: Impact of the MRI and MRI fringe field.* Medical physics, 2012. **39**(2): p. 874-890.
31. Oborn, B., S. Kolling, P.E. Metcalfe, S. Crozier, D. Litzenberg, and P. Keall, *Electron contamination modeling and reduction in a 1 T open bore inline MRI-linac system.* Medical physics, 2014. **41**(5): p. 051708.
32. Keyvanloo, A., B. Burke, B. Warkentin, T. Tadic, S. Rathee, C. Kirkby, D. Santos, and B. Fallone, *Skin dose in longitudinal and transverse linac-MRIs using Monte Carlo and realistic 3D MRI field models.* Medical physics, 2012. **39**(10): p. 6509-6521.
33. Raaijmakers, A., B. Raaymakers, and J. Lagendijk, *Integrating a MRI scanner with a 6 MV radiotherapy accelerator: dose increase at tissue–air interfaces in a lateral magnetic field due to returning electrons.* Physics in medicine and biology, 2005. **50**(7): p. 1363.
34. Bol, G., S. Hissoiny, J. Lagendijk, and B. Raaymakers, *Fast online Monte Carlo-based IMRT planning for the MRI linear accelerator.* Physics in medicine and biology, 2012. **57**(5): p. 1375.
35. Raaijmakers, A., B. Hårdemark, B. Raaymakers, C. Raaijmakers, and J. Lagendijk, *Dose optimization for the MRI-accelerator: IMRT in the presence of a magnetic field.* Physics in medicine and biology, 2007. **52**(23): p. 7045.
36. Kontaxis, C., G. Bol, J. Lagendijk, and B. Raaymakers, *Towards adaptive IMRT sequencing for the MR-linac.* Physics in medicine and biology, 2015. **60**(6): p. 2493.
37. Kontaxis, C., G. Bol, J. Lagendijk, and B. Raaymakers, *A new methodology for inter-and intrafraction plan adaptation for the MR-linac.* Physics in medicine and biology, 2015. **60**(19): p. 7485.
38. Kolling, S., B. Oborn, and P. Keall, *Impact of the MLC on the MRI field distortion of a prototype MRI-linac.* Medical physics, 2013. **40**(12): p. 121705.
39. Low, D., S. Mutic, S. Shvartsman, T. Chmielewski, G. Fought, A. Sharma, and J. Dempsey, *TU-H-BRA-02: The Physics of Magnetic Field Isolation in a Novel Compact Linear Accelerator Based MRI-Guided Radiation Therapy System.* Medical Physics, 2016. **43**(6): p. 3768-3768.
40. Burke, B., M. Lamey, S. Rathee, B. Murray, and B. Fallone, *Radio frequency noise from clinical linear accelerators.* Physics in medicine and biology, 2009. **54**(8): p. 2483.
41. Lamey, M., B. Burke, E. Blosser, S. Rathee, N. De Zanche, and B. Fallone, *Radio frequency shielding for a linac-MRI system.* Physics in medicine and biology, 2010. **55**(4): p. 995.
42. Fallone, B., B. Murray, S. Rathee, T. Stanescu, S. Steciw, S. Vidakovic, E. Blosser, and D. Tymofichuk, *First MR images obtained during megavoltage photon irradiation from a prototype integrated linac-MR system.* Medical physics, 2009. **36**(6): p. 2084-2088.
43. Lagendijk, J.J., B.W. Raaymakers, and M. van Vulpen. *The Magnetic Resonance Imaging–Linac System*. in *Seminars in radiation oncology*. 2014. Elsevier.
44. Lamey, M., J. Yun, B. Burke, S. Rathee, and B. Fallone, *Radio frequency noise from an MLC: a feasibility study of the use of an MLC for linac-MR systems.* Physics in medicine and biology, 2010. **55**(4): p. 981.
45. Lamey, M., S. Rathee, L. Johnson, M. Carlone, E. Blosser, and B. Fallone, *Radio frequency noise from the modulator of a linac.* Electromagnetic Compatibility, IEEE Transactions on, 2010. **52**(3): p. 530-536.
46. Fallone, B.G. *The rotating biplanar linac–magnetic resonance imaging system*. in *Seminars in radiation oncology*. 2014. Elsevier.
47. Keall, P.J., M. Barton, and S. Crozier. *The Australian magnetic resonance imaging–linac program*. in *Seminars in radiation oncology*. 2014. Elsevier.

48. Mutic, S. and J.F. Dempsey. *The ViewRay System: Magnetic Resonance–Guided and Controlled Radiotherapy*. in *Seminars in radiation oncology*. 2014. Elsevier.
49. Wooten, H.O., O. Green, M. Yang, T. DeWees, R. Kashani, J. Olsen, J. Michalski, D. Yang, K. Tanderup, and Y. Hu, *Quality of Intensity Modulated Radiation Therapy Treatment Plans Using a 60 Co Magnetic Resonance Image Guidance Radiation Therapy System.* International Journal of Radiation Oncology* Biology* Physics, 2015. **92**(4): p. 771-778.
50. Mutic, S., D. Low, T. Chmielewski, G. Fought, M. Hernandez, I. Kawrakow, A. Sharma, S. Shvartsman, and J. Dempsey, *TU-H-BRA-08: The Design and Characteristics of a Novel Compact Linac-Based MRI Guided Radiation Therapy (MR-IGRT) System.* Medical Physics, 2016. **43**(6): p. 3770-3770.

Chapter 3
Gridded electron gun

Performance of a clinical gridded electron gun in magnetic fields: Implications for MRI-Linac therapy

Brendan Whelan[1,2,3]

Lois Holloway[2,3]

Dragos Constantin[4††††††]

Brad Oborn[5]

Magdalena Bazalova-Carter[6,7]

Rebecca Fahrig[4‡‡‡‡‡‡]

Paul Keall[1,2]

1. Radiation Physics Laboratory, University of Sydney, Sydney (NSW), 2006, Australia
2. Ingham Institute for Applied Medical Research, Liverpool (NSW), 2170, Australia
3. Liverpool Cancer Therapy Centre, Liverpool Hospital, Liverpool (NSW), 2170, Australia
4. Stanford University, Palo Alto (CA), 94305, USA
5. Illawarra Cancer Care Centre, Illawarra Hospital, Wollongong, 2500, Australia
6. Xcite Lab, University of Victoria, Victoria, Canada
7. Royal Jubilee Hospital, Victoria, Canada

†††††† Dragos Constantin is currently employed by Varian Medical Systems
‡‡‡‡‡‡ Rebcca Fahrig is currently employed by Siemens Healthcare

Abstract

Purpose: MRI-Linac therapy is a rapidly growing field, and requires that conventional linear accelerators are operated with the fringe field of MRI magnets. One of the most sensitive accelerator components is the electron gun, which serves as the source of the beam. The purpose of this work was to develop a validated Finite Element Model (FEM) model of a clinical triode (or gridded) electron gun, based on accurate geometric and electrical measurements, and to characterise the performance of this gun in magnetic fields.

Methods: The geometry of a Varian electron gun was measured using 3D laser scanning and digital callipers. The electric potentials and emission current of these guns was measured directly from six dose matched true beam linacs for the 6X, 10X, and 15X modes of operation. Based on these measurements, a Finite Element Model (FEM) of the gun was developed using the commercial software Opera/Scala. The performance of the FEM model in magnetic fields was characterised using parallel fields ranging from 0-200 G in the in-line direction, and 0-35 G in the perpendicular direction.

Results: The FEM model matched the average measured emission current to within 5% across all three modes of operation. Different High Voltage settings are used for the different modes; the 6X, 10X and 15X modes have an average high voltage setting of 15 kV, 10 kV, and 11 kV. Due to these differences, different operating modes show different sensitivity in magnetic fields. For in line fields, the first current loss occurs at 40 G, 20 G, and 30 G for each mode. This is a much greater sensitivity than has previously been observed. For perpendicular fields, first beam loss occurred at 7.5 G, 5 G and 5G and total beam loss at 27 G, 22 G, 20 G.

Conclusions: A validated FEM model of a clinical triode electron gun has been developed based on accurate geometric and electrical measurements. Three different operating modes were simulated, with a maximum mean error of 5%. This gun shows increased sensitivity to in-line magnetic fields than previously presented models, and different operating modes show different sensitivity.

1. Introduction

MRI-Linac guided radiotherapy is being actively developed and implemented by several groups as the logical next step in the evolution of radiation therapy. Since MRI-Linac radiotherapy requires that an electromagnetically sensitive electron accelerator functions within the magnetic fringe field of an MRI scanner, the emergence of MRI-Linac radiotherapy has driven increased efforts in linear accelerator simulation amongst the medical physics research community. Operation of the electron gun in external magnetic fields is one of the most sensitive and studied aspects of integrated MRI-Linac systems [1-5], and is the focus of this paper. An electron gun is used to inject a steady stream of electrons at kilovoltage (kV) energies into an accelerating waveguide, where they are accelerated to megavoltage (MV) energies. Assuming that the linear accelerator is of the straight ahead type (no bending magnet), the electron gun is the most sensitive part of the accelerator to external magnetic fields. This is because slow moving electrons are less magnetically stiff, and hence more sensitive to external magnetic fields.

Two classes of electron gun are used in medical accelerators; diode and triode (or gridded) [6]. Diode electron guns have only two separate electric potentials; the cathode and focusing electrode, which are set to some negative voltage (typically on the order of tens of kilovolts) and the anode, which is held at ground. In a triode gun, a control grid is added just above the surface of the cathode. The grid is held at a third potential, typically within around one hundred volts of the cathode potential. However because it is close to the cathode, this can create strong variations in the electric field at the cathode surface, and hence emitted current. The focusing electrode is held at the same electric potential as the grid in this configuration. A schematic of a triode electron gun is shown in Figure 1.

Whilst diode electron guns are suitable for use in simple, low energy machines, the majority of linear accelerator manufacturers utilise a triode electron gun [6, 7]. This is because a triode gun allows far more flexibility than a diode gun can achieve; for a given beam energy any beam current can be achieved, and for a given current any beam energy can be achieved (within physical limitations) [8]. This is not the case in diode guns, where the current and energy of the beam are intrinsically linked. It also allows a very simple and stable beam gating mechanism, as only a few hundred volts need to be switched to turn on or off the beam. On linear accelerators which use diode guns, beam gating must be carried out by radiofrequency (RF) pulse interruption methods, and this appears to result in quite long beam on/off latency [9, 10]. Triode guns are also important for treatment machines which provide multiple treatment energies, and both photon and electron treatment modes. This is because both the target heating and the dose rate are correlated with the product of beam current and beam energy. In order to keep these parameters within an acceptable range, it is often desirable to decrease beam current when beam energy is increased. For electron therapy, the beam current must be further decreased, as the dose conversion efficiency is orders of magnitude higher.

Although the MRI-Linac devices investigated in the published literature to date appear to utilise diode electron guns [11-13], recent publications have focused on the potential of multi-

mode MRI-Linac devices to improve achievable dose distributions for deep seated tumours [14, 15], and beam gating using MRI-Linacs [16, 17]. As such, it is clear that triode electron guns are likely to find useful application in the next generation of MRI-Linac systems. However, to date all published work in this area has focused on the performance of diode electron guns in magnetic fields. Further, no technical details on medical triode guns can be found in the literature. Therefore, the purpose of this work is twofold; (1) to present an accurate model of a medical triode electron gun such that it may be utilised by future researchers, and (2) to test the sensitivity of this gun in magnetic fields, in particular to assess whether different operating modes show different sensitivity to magnetic fields.

2. Materials and methods

2. 1. Theory and background

The geometric structure of the electron guns typically used in medical linacs can be broadly classified as a Pierce gun [18]. Electron guns in medical accelerators are operated in a steady state, and as such the physics of the system is essentially electrostatic and is described by Poisson's law (Equation 1):

$$\nabla^2 V = \frac{p}{\varepsilon_0} \qquad \text{Equation 10}$$

Where V is the electrostatic potential, p is the charge density (defined by the beam), and ε_0 is the electric vacuum permittivity. Although formulating the problem exclusively in terms of Poisson's equation does neglect the self-magnetic fields of the beam, it has previously been shown that these effects are small for these guns [3].

Accurate computational simulation of a DC electron gun comprises three important sub steps. The first is to calculate the electrostatic field arising from the applied electrode potentials and space charge of the beam; the second is to approximate thermionic electron emission in the presence of this electric field, and the third is to track the emitted electrons through this field. These steps are not independent of each other; therefore an iterative solving technique is required. Calculating the net electric field of the electrode and beam is a complicated problem, and for accurate results computational electrodynamics is required, in which complex geometry is discretised (meshed) into many subdomains. Examples of appropriate computational schemes include the finite element method and the finite difference technique.

Medical electron guns are typically operated in a space-charge limited mode of thermionic emission. This means that current is extracted from a hot (around 1200 K) cathode and is limited by the space charge associated with the cloud of electrons already extracted. The major advantage of operating in this regime is stability, as small fluctuations in temperature or cathode emissivity have minimal impact on emitted current. Space charge limited emission in one dimension is described by Child's law (Equation 2)

$$J_{emitted} = \frac{4\varepsilon_0}{9}\sqrt{\frac{2q}{m_0}}\frac{V^{3/2}}{d^2} \qquad \text{Equation 11}$$

Where $\boldsymbol{J_{emitted}}$ is the emission current, ε_0 is the permittivity of free space, q & m_0 are the electron charge and mass, and $\boldsymbol{V}$ is the potential difference at distance $\boldsymbol{d}$ from the cathode – a function of both the electrode potentials and the beam itself [6]. Modelling thermionic emission is one of the most challenging aspects of electron gun simulation. In 2D or 3D, the above formulation is only strictly valid for an infinite planar diode, and so for realistic geometries the cathode face is discretised into a number of smaller faces over which the 1D formulation is applied. In order to ensure that this approach is valid, *V* must be sampled at a specified distance from the discretised faces, typically less than 1mm. The presence of a grid adds further complexity to the simulation, as the sampling distance and meshing parameters become more sensitive. If the thermal Maxwell-Boltzmann distribution of the emitted electrons is considered, the situation becomes still more complex. A detailed discussion of the various thermionic emission effects is beyond the scope of this paper, but is provided in ref [19].

The overriding purpose of an electron gun in a therapeutic system is to inject a stream of electrons into the accelerating wave guide. An important figure of merit for the system performance is the capture efficiency – that is, the percentage of electrons injected into the Linac which actually reach the target. It is obvious that for optimal capture efficiency, the injected electrons should have low transverse velocities and be confined to a small region of space around the axis of the accelerator. Both requirements can be quantified by examining the distribution of the electron beam in phase space. Phase space is a six dimensional position-momentum space: [*x,y,z,px,py,pz*]. However, typically in the case of DC electron guns the longitudinal components are trivial, and the *x* and *y* components identical. Therefore, the beam can be sufficiently described in a two dimensional transverse phase space [*x,px*] along with information about the beam energy and current. It is often convenient to describe the phase space in terms of a divergence angle instead of transverse momentum: *x'=px/pz*. This is referred to as trace space.

In order to characterize a distribution in phase space or trace space, one can fit an ellipse to the particle distribution. The most common approach to this is to calculate Twiss parameters, also known as the Courant-Snyder parameters. These parameters can be used to describe an ellipse which will encircle the central part of the beam. Mathematically, these are described as follows:

$$\varepsilon_{x,RMS} = \sqrt{\langle \boldsymbol{x}^2 \rangle \langle \boldsymbol{x'}^2 \rangle - \langle \boldsymbol{x x'} \rangle^2} \qquad \text{Equation 12}$$

$$\boldsymbol{\alpha} = \langle \boldsymbol{x x'} \rangle / \boldsymbol{\varepsilon_{x,RMS}} \qquad \text{Equation 13}$$

$$\boldsymbol{\beta} = \langle \boldsymbol{x^2} \rangle / \boldsymbol{\varepsilon_{x,RMS}} \qquad \text{Equation 14}$$

The Twiss parameters allow one to quickly infer a lot of useful information about the beam. $\varepsilon_{x,RMS}$ (Equation 3) quantifies the area of the beam in phase space; larger emittance corresponds to less ordered electron beams. α (Equation 4) determines the orientation of the ellipse, which describes whether a beam is diverging or converging. Another interpretation of α is that it reflects the distance from a beam waist. At a beam waist, α is zero, and the average transverse electron velocities are zero – an opportune time to capture the electron beam in an

RF accelerator. The (*rms*) beam width in space is described by $\sqrt{\varepsilon_{x,RMS}.\beta}$. In the absence of space charge, accelerating fields, and non-linear transverse fields (conditions which are approximately met in a field free drift tube), emittance as defined in equation 3 is a constant, whilst the other parameters evolve as a function of the longitudinal component *z*. The *rms* emittance only encapsulates the central 37% of the beam. Some authors add a factor of 4 or π to this definition in order to better describe the full beam. Unfortunately, it is often unclear exactly which definitions is being used, which can make it difficult to compare different results. In this work, emittance is calculated according to equation 3 without any additional factors. Units of mm and mm rad are used.

2. 2. Gun geometry

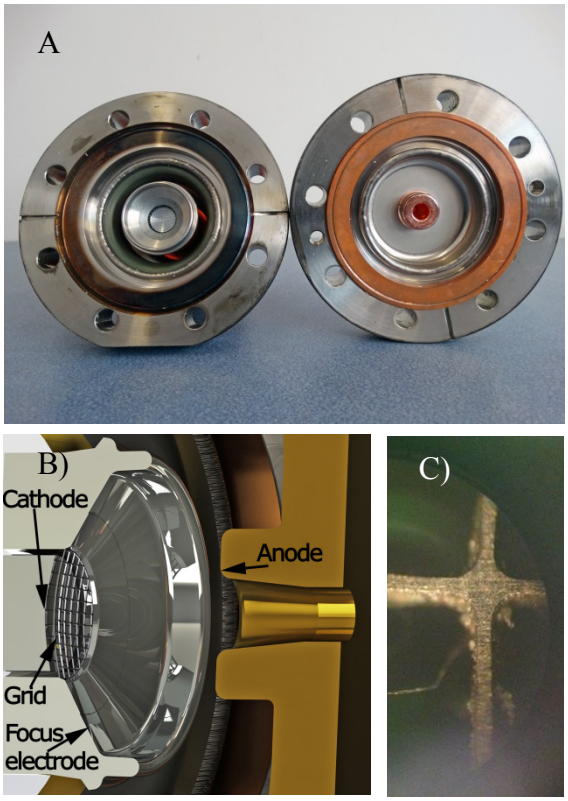

Figure 25: (A) Photograph of the electron gun modelled in this work. On the left side is the cathode section, and on the right side is the anode section. (B) A rendering of the gun geometry developed based on a laser scan of the electron gun. (C) A photograph of the grid taken through a microscope lens.

A commercial electron gun from a Varian high-energy linac was obtained through a spare parts supplier (Figure 1A). The geometry of the gun was measured using a 3D laser scanner operated by a third party engineering company[§§§§§§]. The scan was rated as being accurate to within 0.1 mm. The control grid was too fine to be resolved by the 3D scan; in order to measure curvature of the grid, a thin film was placed over the grid and measured in a separate 3D scan. The wire thickness and spacing of the grid was measured using digital callipers (error rating of 0.01 mm). A microscope image of the grid is shown in Figure 1C. Based on these measurements a geometrical model of the gun was developed. This geometry is detailed in the results section 3.1 and Figure 2.

2. 3. Finite Element Model of Gun

A Finite Element Model (FEM) of the gun geometry from section 2.2 was developed using the commercial software Opera/SCALA, version 16 (Kidlington, UK). Opera solves Poisson's equation in the presence of thermionic electron emission. Child's law was used to simulate thermionic space charge limited emission (Equation 2). The type 1 emitter option (virtual cathode) was used, which is the most accurate model according to the Opera user manual. The initial electron energy was set to the mean of the Maxwell-Boltzmann equation for a temperature of 1245 K (a typical dispenser cathode temperature). It is possible to sample additional energies from the Maxwell-Boltzmann distribution, but this significantly increases simulation time and does not strongly affect the solution. Opera requires that the user inputs a sample distance where the program calculates the electric field, which was set to 0.25 in this work. Importantly, this distance is less than 1/3 the cathode-grid spacing, as recommended in the Opera user manual. The emitted electron current is strongly dependent on the electric field at the cathode, which in turn is altered by emitted current. Therefore, an iterative simulation procedure is required for accurate electron gun simulation. In this study, the convergence tolerance was 0.001; the model is considered solved when the relative RMS change in potential is less than this. A maximum limit of 21 iterations was set; in practice, the models normally resolved within 10 iterations. Wherever possible, symmetry in the model was exploited and only a quarter of the geometry was solved, greatly reducing computational

[§§§§§§] 3D Scanning Australia, http://3dscanningaustralia.com.au/

resources. The quarter model was discretised into around three million tetrahedral elements. In electron gun simulations, a high quality mesh near the cathode is crucial, as the emitted current is critically dependant on an accurate solution of the electric field here. We used quadratic elements of maximum size 0.15 mm near the cathode, and 1 mm elsewhere. To check for mesh independence, we also created a model with doubled mesh density and verified that the solution did not change. In Opera, the number of macro particles representing the beam is a product of the mesh density, so this step also confirmed the number of macro particles was adequate to represent the beam. The emission current was comprised of ~15000 macro particles.

Next, parameter sweeps were carried out over the high voltage and grid voltage in order to infer the working region of the gun. The high voltage was varied from -5 kV to -30 kV in 5 kV steps, whilst the grid voltage was varied from 20V to 140V in 7.5 V steps. The grid voltage is always quoted relative to the high voltage, and the focusing electrode is set to the same boundary condition as the grid. The electron beam phase space and current was extracted from the solved models for axial planes spaced every 2 mm, 21 to 39 mm from the cathode. This information was imported into Matlab version 2014 (Nattick, MA) and in house code was used to calculate the transverse emittance and current. The current emitted from the cathode was also recorded for each simulation.

2. 4. Experimental measurement of gun operating modes

Although the above process allows one to derive approximate working regions of the gun for a given high voltage, it does not provide information about the conditions in which the gun is actually operated clinically. In order to acquire experimental data regarding the operating modes used clinically, the high voltage, grid voltage and emission current were read directly from the console of six dose matched Varian TrueBeam Linacs (Varian, Palo Alto) in service mode. Data was acquired for the 6X, 10X, and 15X photon modes, and 2D interpolation of the data from section 2. 3 was used to compare simulated emission current as function of high voltage and grid voltage with the experimentally measured values. Based on these results, three separate gun models were built to represent the 6X, 10X, and 15X modes of operation.

2. 5. Grid height determination

The most sensitive parameter in the model is the spacing between the grid and the cathode (grid height). Unfortunately, this is also the parameter with the largest uncertainty. This is because the grid had to be measured separately during the laser scan, and so in the worst case scenario the error could be 0.2mm – that is, 0.1mm from each scan. In order to obtain the best match to the experimental data (section 2.4), we repeated the simulations characterising beam parameters as a function of the electrode potentials (section 2.3) for three different grid heights – 0.84mm (as measured by the 3D scan), 0.9mm, and 1.0mm. The simulated emission current for each case was compared to the experimental data. The model with the grid height of 1.0 mm gave the best match to the experimental data, and is used throughout the results section (except for section 3.4 where data on the grid height determination is presented). As is also detailed in this section, the grid height does not have a major impact on the beam phase space other than beam current

2. 6. Gun sensitivity to magnetic fields

In order to test the sensitivity of the gun to magnetic fields, the Opera FEM models of the 6X, 10X and 15X photon beams based on Section 2. 4 were solved in constant parallel fields in-line and perpendicular to the beam direction. In-line field strengths of 0-200 G were tested in 10 G steps, and perpendicular field strengths from 0-35 G in 5 G steps. As the gun is small (~ 30 mm beam length), constant parallel fields provide a very close approximation to the averaged heterogeneous MRI fringe fields across the gun. For each simulation, the current and transverse phase space at the gun exit were recorded.

2. 7. Comparison of commercially available electron gun solvers

In addition to the Opera/SCALA solver, we have access to the CST 'tracking' solver (Darmstadt, German, version 2016) which is also used to simulate DC guns. The difference in the solution returned by these two commercial solutions is one measure of the inherent uncertainty in simulations. Therefore, exactly the same geometry was implemented in CST, and a mesh built using the same (or as close as possible) meshing parameters as in Opera. Like Opera, CST uses a tetrahedral mesh in conjunction with an FEM algorithm, although the CST model uses quadratic (i.e. higher order) elements everywhere instead of just near the cathode. The source was defined using exactly the same parameters as in Opera. A simulation with much higher mesh density was solved in order to check that mesh independence had been reached, as was done with the Opera model.

3. Results

3. 1. Gun geometry

Figure 2 shows a detailed sketch of the geometry used in all simulations.

The shaded part of the anode drift tube in Figure 2 is subject to additional uncertainty beyond the 0.1 mm of the 3D laser scan. This is because the 3D scan has a limited ability to 'see'

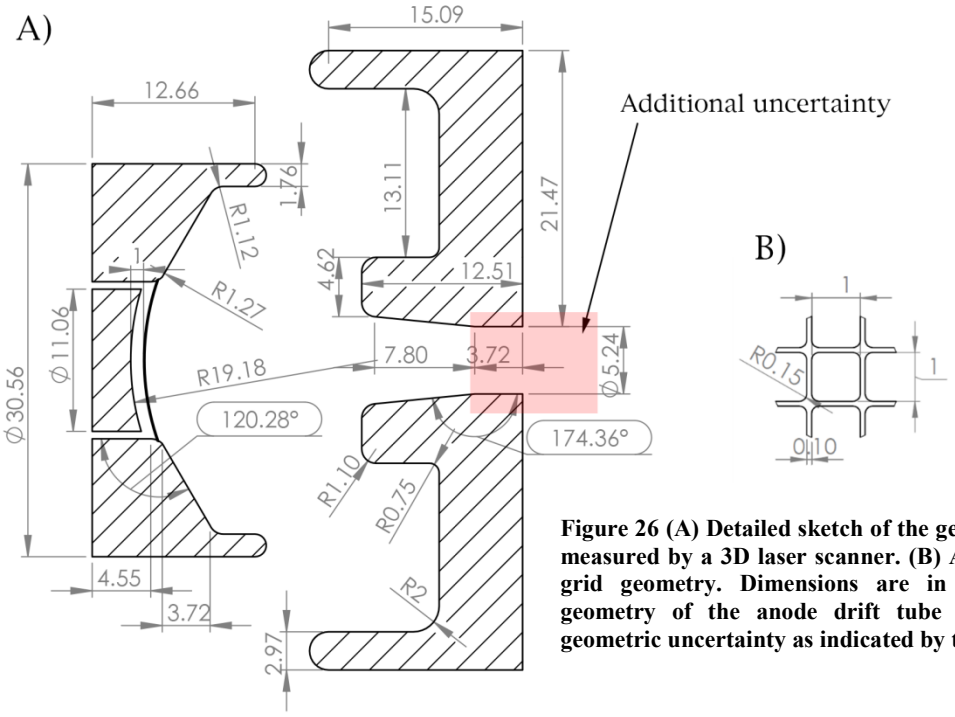

Figure 26 (A) Detailed sketch of the geometry used in this study measured by a 3D laser scanner. (B) A section of the measured grid geometry. Dimensions are in mm and degrees. The geometry of the anode drift tube is subject to additional geometric uncertainty as indicated by the shaded region.

down a narrow tube. The tapered section of the anode was measured with the 3D scan. After this, we extended the length of the anode drift tube with a straight cylinder, as this geometry best matches other available schematics of Varian triode guns [6] and is consistent with previously published models [2, 3]. However, it is also possible that the taper continues to some minimum. This part of the geometry can be very important for mitigating back bombardment [20], and can have a strong effect on the gun sensitivity to in-line magnetic fields. This is discussed in section 4.

3. 2. Finite Element Model of Gun

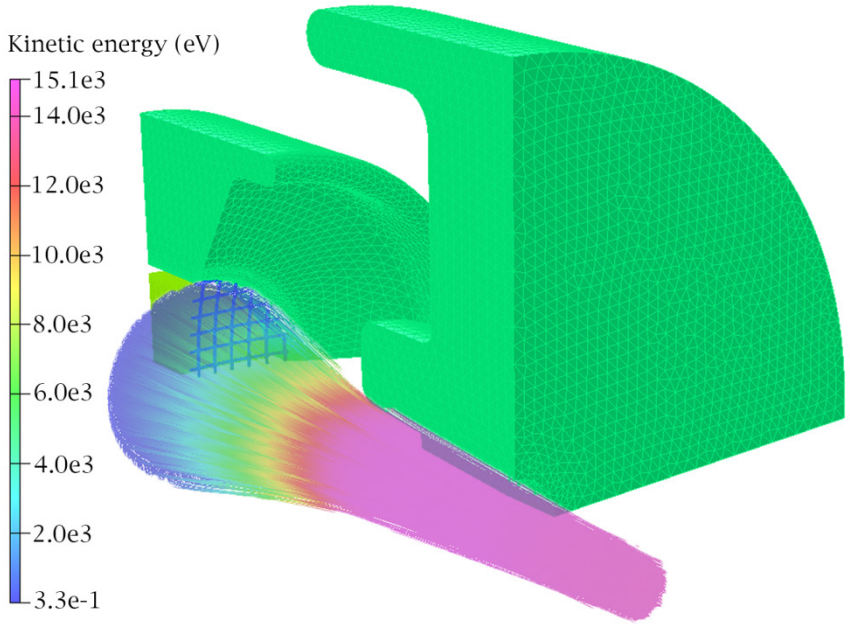

Figure 27: Gun solution at zero magnetic field. The 6X mode is shown and the particle trajectories are coloured according to their energy. The solid structures represent one quarter of the gun geometry with the surface FEM mesh displayed.

Figure 3 shows the particle trajectories of the solved model. Particles are color coded according to their kinetic energy. Figure 4 show various gun metrics plotted as a function of both HV and grid voltage; (A) emission current, (B) injection current (current injected into the waveguide) (C) efficiency (injection current as percentage of emission current) and (D) emittance in trace space. The injection current is always less than the emission current as some part of the beam is blocked by the grid. This fraction increases as the grid voltage is turned up, which make it more electrically attractive to electrons (Figure 4C). Also, if the current becomes too large, then the focusing electrode will no longer focus the beam, leading to further current loss from anode collisions. This effect can clearly be seen in the 5kV and 10kV series of Figure 4C. The emittance changes as total field (electrodes + space charge) changes. It can be seen that there is a certain amount of current which creates a minimum in emittance for a given HV setting. The data shown in Figure 4A was used to compare the models to the experimental data in section 3.3 and optimise the model, as outlined in section 3.4.

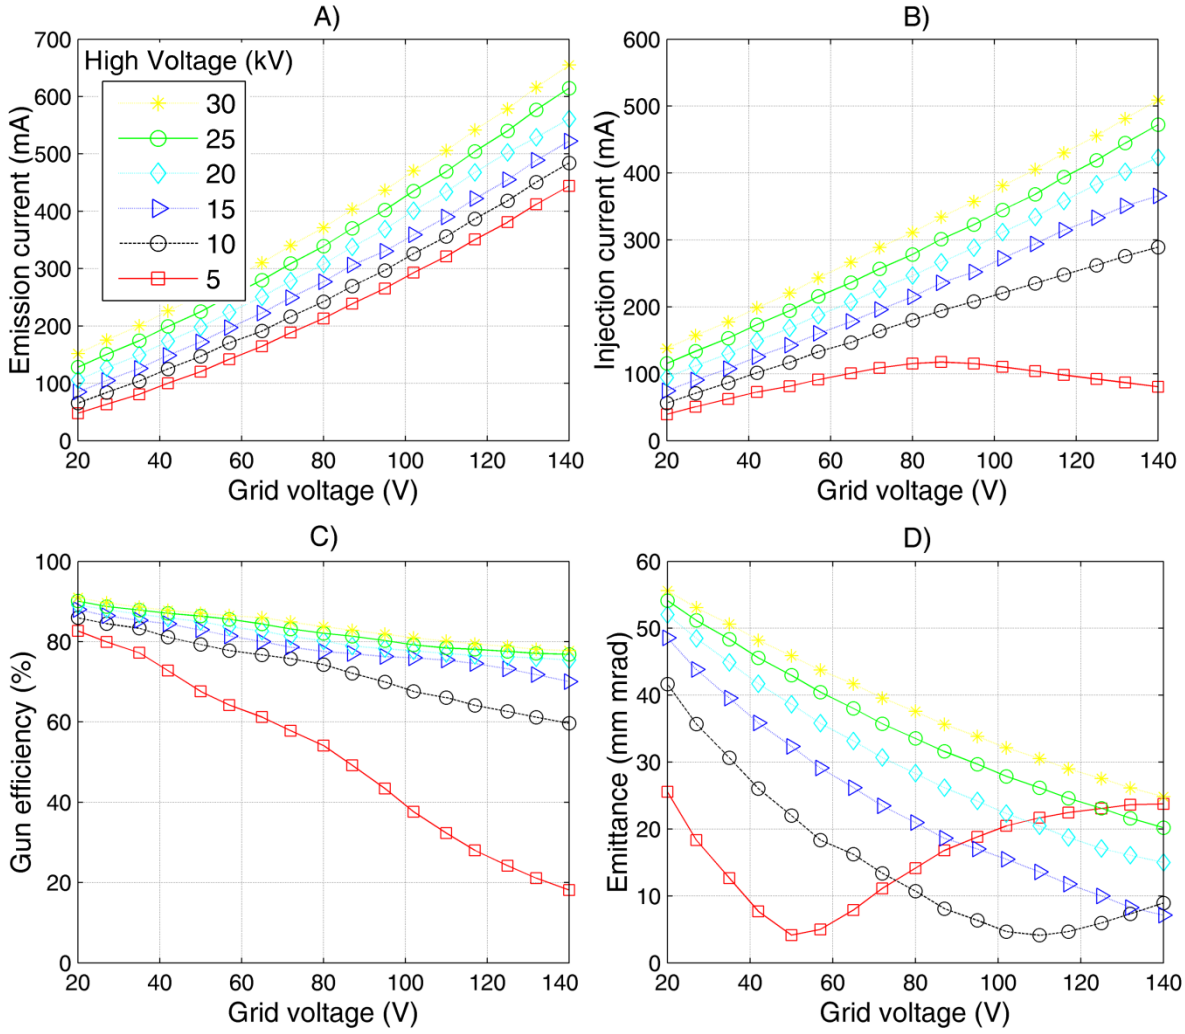

Figure 28: A) Emission Current, B) Injection Current (current injected into the waveguide), C) Gun Efficiency (percentage of the emission current which exits the gun), and D) Emittance as a function of grid voltage (x axis) and high voltage (different series)

3. 3. Experimental measurement of gun operating modes

The values for HV and grid voltages measured on the 6 TrueBeam Linacs are shown in Table 1, along with the interpolated emission current based on the data from section 3. 2. The mean absolute error in emission current between experimental and simulated data across the 6 Linacs is 3.2%, 5.1%, and 3.5% for the 6X, 10X and 15X modes. Note that one data point has been excluded from these averages in the 10X series, as the HV is clearly an outlier compared to the other linacs.

Table 3: Gun parameters (HV, Grid voltage, and gun current) measured from six dose-matched TrueBeam linacs. The shaded entry for 10X was not included in this study, as the HV value is clearly an outlier compared to the other five data points. Note that error reported in the 'mean' row for each mode is the absolute mean of all measured errors.

Mode	Measured High Voltage (kV)	Measured Grid Voltage (V)	Measured Emission Current (mA)	Simulated Emission Current (mA)	% Error
6X (6MV X rays)	14.998	105.6	373.4	372.8	-0.2
	14.971	87.7	323.9	308.2	-4.9
	14.965	95.3	341.6	331.4	-3.0
	15.023	87.0	320.9	306.1	-4.6
	14.989	92.8	306.3	323.8	5.7
	14.928	86.4	299.4	303.0	0.0
Mean	**14.979±0.03**	**92.5±7.4**	**327.6±26.8**	**322±27**	**3.3 (absolute)**
10X (10 MV X rays)	9.998	47.3	142.4	139.5	-2.0
	9.992	40.2	130.5	119.0	-8.8
	9.958	37.5	119.7	110.8	-7.4
	12.012	39.1	140.7	125.3	-11.1
	10.001	42.4	120.2	125.3	4.2
	9.946	40.6	123.5	119.9	-2.9
Mean	**9.979±0.025**	**41.6±3.6**	**127.3±9.5**	**122.9±10.6**	**5.1 (absolute)**
15X (15 MV X rays)	10.996	62.2	193.0	190.0	-1.6
	10.99	76.0	243.1	235.9	-2.9
	10.941	72.2	226.7	223.2	1.5
	11.02	62.8	197.8	191.6	-3.1
	10.993	72.6	207.6	224.9	8.3
	10.923	72.4	217.1	223.8	3.1
Mean	**10.98±0.03**	**69.7±5.7**	**214.2±18.8**	**214.0±19.0**	**3.4 (absolute)**

Based on the results in 3. 1 to 3. 3, three models were arrived at to represent the 6X, 10X, and 15X modes of operation. The beam current and energy for each mode is reported on in Table 2. This data differs from the data shown in Table 1 in that it is solved explicitly rather than being interpolated from section 3. 2. The beam waist occurred at 38mm, 33 mm, and 39mm from the centre of the cathode for the 6X, 10X and 15X mode respectively. The phase space at the beam waist (or nearest sample plane) for each mode is shown in Figure 5 (beam waist and phase space are defined in section 2.1). The impact of the grid can clearly be seen in the segmented nature of the phase space distribution.

Table 4: The parameters used for each of the clinical mode simulations.

Mode	High Voltage (kV)	Grid Voltage (V)	Emission Current (mA)	Injection Current (mA)	Emittance (mm mrad)
6X	15	93	323	246	17.8
10X	10	42	124	100	26.4
15X	11	70	214	164	16.8

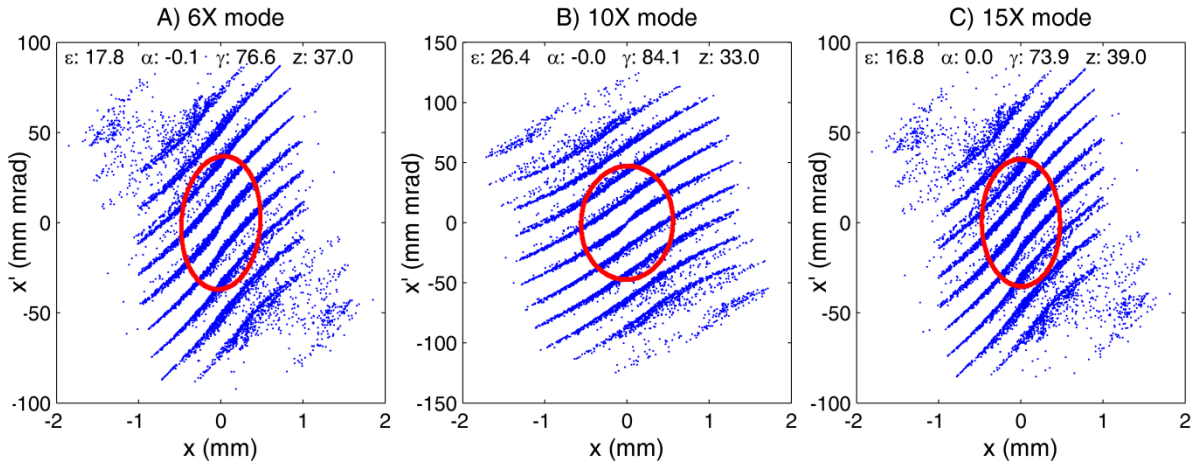

Figure 29: Phase space at the beam waist for the A) 6X, B) 10X, and C) 15X modes. Printed on each plot are the Twiss parameters and resultant RMS phase ellipse, as described in section 2. 1. The z parameter is the distance from the cathode in mm.

3. 4. Grid Height Determination

The results shown in all other sections of this paper are for a model with a cathode to grid height of 1mm, as detailed in Figure 2. However, the measured value was 0.84mm. In this section we detail how this value was optimised. Table 3 shows the mean error for each mode for grid heights of 0.84, 0.9, and 1.0 mm (note that all these values are within the experimental uncertainty of the measured geometry). The model with grid height of 1.0 mm returned the lowest mean error values when compared to the experimental data in Table 1. Figure 6 shows the impact of changing the grid height for the 6X mode. At first glance, it may appear that altering the grid height has a dramatic impact on the gun emittance – however, this is not quite correct. Altering the grid height has a strong impact on emission current, which does strongly impact emittance (also apparent in Figure 4D). However, for the range of grid heights investigated here, similar emittance is achieved when similar current is drawn from the cathode. As a representative example: For grid height 0.84mm, HV=15kV, and Grid=93V, 410 mA of current is extracted and the emittance is 10 mm mrad. For grid height 1.0mm, the same current is drawn with the grid set to 115, resulting in an emittance 12 mm mrad. This gives a better indication of the impact altering the grid height on the phase space distribution.

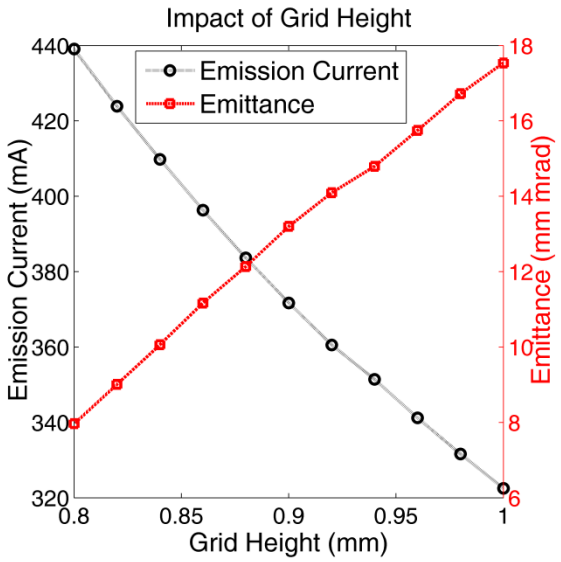

Figure 30: Impact of the grid height on the emission current and the emittance for the 6X mode. Note that the change in the emittance is primarily due to the changing current rather than the changing grid height.

Table 5: Mean error in simulated emission current for each mode for three different grid heights, compared to the experimental data in Table 1.

Mode	**Mean Error (%)**		
	Grid Height: **0.84 mm**	Grid Height: **0.9 mm**	Grid Height: **1.0 mm**
6X	34%	14%	-2%
10X	30%	12%	-3%
15X	37%	16%	0%

3. 5. *Gun sensitivity to magnetic fields*

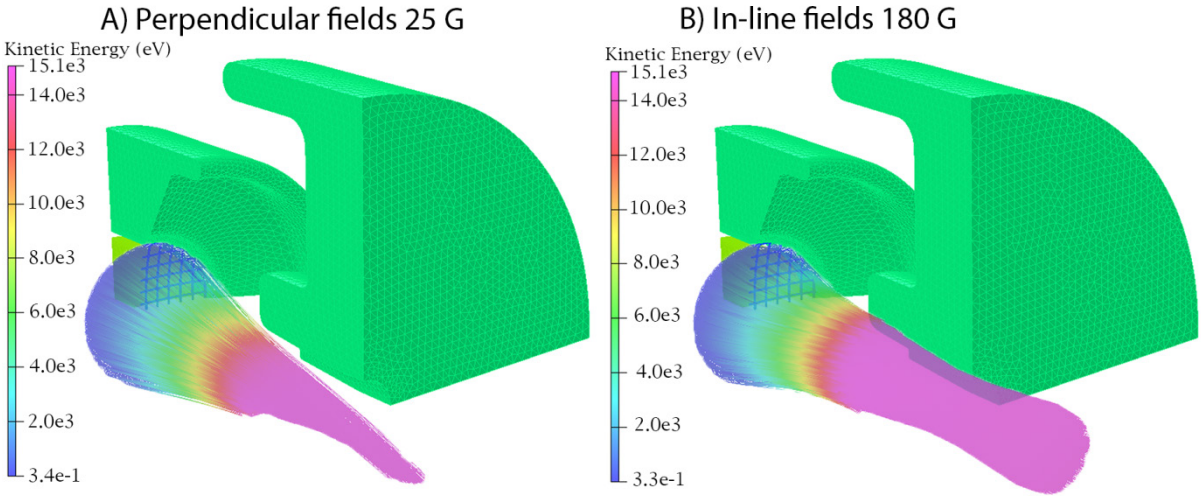

Figure 31: Examples of particle trajectories in A) Perpendicular fields of strength 25 G, and B) In-line fields of strength 180 G. Note that only one quarter of the geometry is shown so that the particle trajectories can be seen. The 6X mode is shown and the particles are color coded according to their energy. The equivalent figure for zero field trajectories is shown in Figure 3.

Figure 7 shows examples of the particle trajectories of the 6X mode in both (A) in-line and (B) perpendicular fields, whilst Figure 8 shows the injection current and emittance results of operation in magnetic fields for the 6X, 10X and 15X modes. The emittance results should be interpreted with caution; emittance only enables a particularly useful comparison between beams when the total current and energy is the same. For instance, for the in-line case, as the magnetic field increases the outer electrons are scraped off by the anode, the emittance begins to decrease – however, this comes at the cost of substantially reduced current.

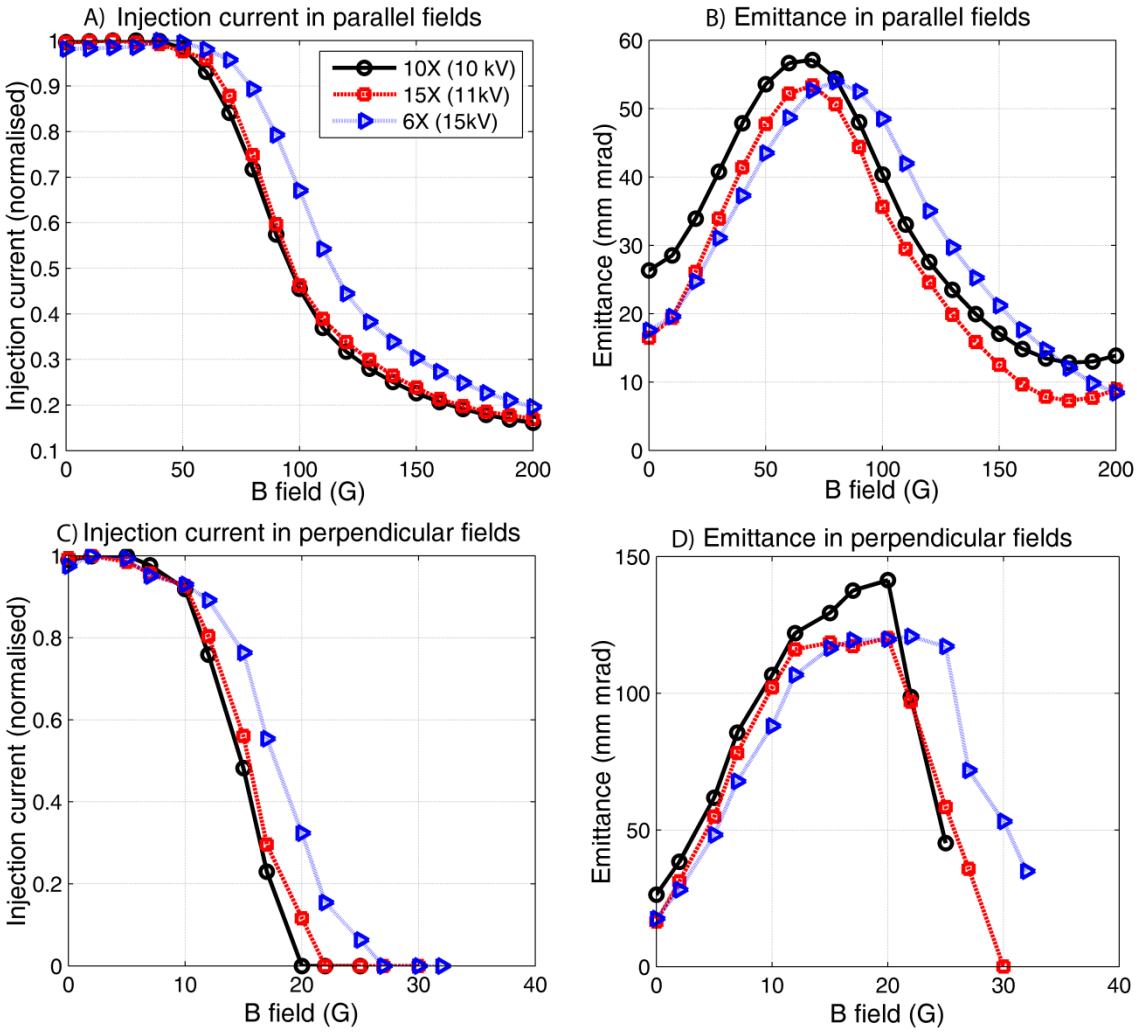

Figure 32: Injection Current for in-line (A) and perpendicular (B) fields, and emittance for in-line (B) and perpendicular (D) fields.

Table 4 shows a comparison of the three Pierce type electron guns which have previously been simulated in magnetic fields with this work. It can be seen that this model is around twice as sensitive in in-line magnetic fields as previously tested models. The sensitivity in in-line magnetic fields is primarily a function of the HV setting and the ratio of the cathode-to-anode radius. The sensitivity in perpendicular fields on the other hand is primarily a function of the HV setting; as such this model shows slightly less sensitivity to perpendicular fields than previously published models which used a lower HV setting.

Table 6: A summary of previously published results of electron gun models in magnetic fields, and this work. The 'sensitivity' metric in columns 3 and 4 is simply the field where the first beam loss occurs. Not all authors reported on this exactly, in which case it was graphically derived from plots and the '~' symbol is used.

Model	Ref	In-line field sensitivity (G)	Perpendicular field sensitivity (G)	Gun type	Boundary conditions	Geometry
Litton	[3]	~163 G	2 G	Non-gridded	From literature	From literature
Varian	[3]	~141 G	5 G	Non-gridded	From literature	From literature
Alberta	[1, 5]	100 G	NA	Non-gridded	Measured	Inferred
This work		**40 G**	**7.5 G**	**Gridded**	**Measured**	**Measured**

3. 6. Comparison of commercial electron gun solvers

A fairly large disagreement in emission current was found between the two commercial FEM solvers, even when every effort was made to match the mesh, emission parameters, and geometry. As outlined in section 3.3, the Opera solver returns a value of 322 mA for the 6X mode parameters. The CST solver on the other hand returns a value of 267 mA. This represents a 17% error compared to Opera. It is worth reiterating that no notable changes occur in these values when the mesh density is increased. Despite the differences in the calculated emission current, the two solvers give quite good agreement regarding the gun performance in magnetic fields (Figure 9A and B). This is an important observation for in-silico MRI-Linac electron gun work, as accurate calculation of the emission current is not the primary goal in this field. In Figure 9, the largest error in relative injection current is 10% for a perpendicular field of 15G.

The emission current is highly dependent on the electric field close to the cathode, and vice versa. Figure 9C shows the axial electric field up to 1mm in front of the cathode. It can be seen that in the absence of any electrons being included in the simulation, the two solvers give extremely high agreement – however, once the electron source is added, this disagreement drops substantially. Coupled with fact that the solutions do not change with increased mesh density, this strongly implies that the difference in calculated emission current is due to different implementations of the space charge limited emission models.

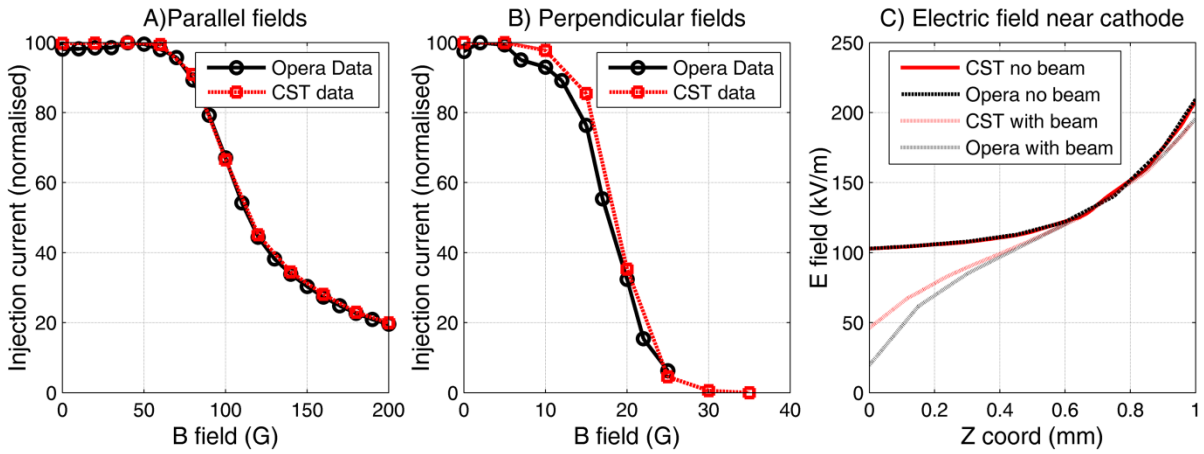

Figure 33: Comparison of the Opera solver used elsewhere in this study to CST, a separate commercial FEM based solver. The 6X mode is used in all figures. A) The normalised injection current in in-line fields, B) The normalised injection current in perpendicular fields, C) The axial electric field from the cathode surface to 1mm.

4. Discussion

We have developed a finite element model of a clinical triode (i.e. gridded) electron gun based on accurate geometric and electrical measurements. The model matches emission current measured on six TrueBeam linacs to within 6% across the 6X, 10X and 15X operating modes (section 3.3). To the best of our knowledge, this is the first time in the literature that a model of a medical triode electron gun has been published. This work also contrasts with previous publications which examined electron gun sensitivity in external magnetic fields, in that measured (rather than inferred) geometry was used (Table 4). Previous studies have either tested models which are no longer used clinically [3], or have inferred the gun geometry based on measurements of emission current and High Voltage [2]. Regarding the latter approach, it is worth noting that this is a degenerate problem, meaning that multiple geometries can produce the same emission current at a given high voltage. This is important as gun geometry directly affects sensitivity to magnetic fields - particularly for the in-line case. The model simulated in this work is approximately twice as sensitive to operation in in-line magnetic fields as previous models.

Utilising a gridded gun results in low beam on/off latency, which is an important aspect of a real time system architecture, and a current area of research for MRI-Linac systems [16, 17]. Gridded guns are also important in applications where continuously variable dose rate is desired, such as arc therapy [6]. In addition, the grid enables much better control over beam parameters such that desirable phase space distributions can be injected into the waveguide over a wide range of HV and current settings. This control is important for multi-mode linacs, because the current which arrives at the target must be appropriate when considering both dose rate and target heating. Multi-mode electron accelerators are a current area of research for MRI-Linacs [14, 15].

Of the major manufactures, Tomotherapy, Cyber-Knife, and Siemens all utilise gridded guns. Varian use gridded guns on multi-mode systems, and diode guns on low energy guides. On the other hand, Elekta utilise diode guns even on their multi-mode guides [21, 22]. Very long gating times (on the order of seconds) have been observed on these linacs, making them

unsuitable for a real time system [9, 10]. Interestingly, there is at least one system designed for real time treatments which does use a diode gun; the Brainlab Vero. This system only uses one beam mode, and is powered by a klystron [7]. Although we could not find any data on the beam latency of the Vero system, a klystron based system would likely have much lower latency than a magnetron based system, since the RF drive power could be very quickly switched. MRI-Linac systems utilising only one beam mode could feasibly be constructed with good beam latency using a klystron based power source and diode gun – however, klystrons are more expensive and complex than magnetrons, and so it may still prove simpler to utilise a triode gun.

A number of comments can be made about the results in this paper. Most strikingly, this model was at least twice as sensitive to in-line magnetic fields as previously published models (section 3.5). Including this work, four Pierce type electron gun models have now been tested in-silico in magnetic fields (Table 4), and based on the amalgamated results some general comments on the most important parameters impacting gun sensitivity can be made. For perpendicular fields, the most important characteristic is the high voltage setting (HV). Larger HV means that the beam becomes more magnetically stiff, or difficult to bend. The model tested here is more robust in perpendicular fields than some previous studies [3], primarily due to the larger HV setting used. The HV is also an important factor for in-line fields, however here the geometry of the gun plays a larger role; specifically, the ratio of the cathode radius to the (minimum) anode radius. The larger this ratio, the more the electron beam must be focused to fit through anode, and this focusing is strongly affected by in-line magnetic fields. The two previous studies which tested electron guns in in-line magnetic fields both found that under 100G, no beam loss occurred. The model investigated here shows substantially greater sensitivity, as the beam begins to drop off at 50G. Also, due to the fact that different operating modes use different HV settings, different operating modes show different sensitivity.

The design of the anode drift tube is a particularly interesting aspect of gun performance in in-line magnetic fields. When one simulates an electron gun in-silico (and without magnetic fields), this geometry has negligible impact on the results, as long as it is large enough for the beam to pass through unimpeded (it is essentially a field free region). In a real system which is coupled to a linac however, the design of this geometry is crucial, as the anode drift tube plays a key role in mitigating cathode damage from back accelerated electrons [20]. In order to minimize the amount of back accelerated electrons reaching the cathode, the drift tube should have a minimum radius corresponding to the minimum beam radius. By choosing the injection point to be close to the beam waist, the minimum anode radius can be made as small as possible, minimizing the amount of back accelerated electrons which can return to hit the cathode. On the other hand, a drift tube design like this would substantially increase sensitivity to in-line magnetic fields. As indicated in section 3.1, there was additional uncertainty in the dimensions of this geometry in this work due to limitations of the 3D scanning technology. As well as this, typically when these guns are replaced, the cathode section (Figure 1A, left) is simply bolted to the end of the accelerating waveguide, which also serves as the anode. The anode that ships with the gun is for the purpose of high voltage and

cathode processing[25]; this strongly implies that the field forming (i.e. internal) parts of the anode are identical to the true anode. However, the internal dimensions of the drift tube may differ from this, as they will have negligible impact on high voltage and cathode processing. Given these sources of uncertainty, the sensitivity of this gun to in-line magnetic fields may be even greater than simulated in this work – which was already twice as sensitive as previous work.

We observed quite large uncertainty in simulated emission current, as this was extremely sensitive to the spacing between the grid and cathode (section 3.4). We measured this as accurately as possible, but still observed variations in emission current of over 30% over the range of grid heights simulated. Moreover, different commercial solvers showed a 17% difference in the emitted current for the same geometry and similar mesh settings. In the real world this uncertainty would not be a major issue; if the current were more or less than predicted, the grid voltage could be adjusted accordingly (Figure 4). Indeed, when examining the variability in operating parameters of the six different guns measured in Table 1, this appears to be exactly what is done practically. This uncertainty is also not of particular concern when modelling the sensitivity of these guns in external fields, since these results are not affected by the emitted current (section 3.6). Accordingly, the two different FEM solvers give very similar results when simulations are carried out in magnetic fields (Figure 9).

This model operates at much higher emittance than previously published work [2-4]. Low emittance indicates a more ordered beam, and in general lower emittance results in decreased loss through a beam transport line. However, there is no published data on how important emittance actually is in medical accelerators, and there is no published data regarding the measured emittance of medical electron guns. Part of the reason for the increased emittance is the fact that the beam is simply larger than previous simulations (cathode radius of 5.53 mm compared to 2.31mm). Given the accuracy with which we have measured the geometry, and the fact that we have closely matched the emission current across multiple modes, we believe that the phase space we have calculated is likely to be highly representative of reality.

An aspect of this work which surprised us was that the 15X mode used a substantially higher injection current than the 10X mode (section 3.3). As electron energy increases, bremsstrahlung conversion efficiency increases, as does beam power (for the same current). As such, we expected to see injection current drop as the photon energy increased. A possible explanation for this result is the fact that different operating modes will have different capture efficiency, and it may simply be that the 15X mode has lower capture efficiency than the 10X mode. Another factor will is the different tungsten targets and flattening filters utilised for different operating modes.

Taken together, these results have some very clear implications for the design of electron guns which are to be used in MRI-Linac systems. Firstly, operating at a higher HV setting results in substantially increased robustness in magnetic fields (section 3.5). However, whilst certainly achievable, redesigning an electron gun to operate at higher HV is not trivial. The

[25] Altair Website, Electron sources and cathodes. Available from: http://www.altairusa.com/electron-gun.php.

HV setting directly affects the velocity of electrons injected into the waveguide, and for a given waveguide (and operating mode) there will be an injection velocity which is optimal. Additionally, higher HV increases the electrical hazard and risk of breakdown and arcing in the gun.

Although there have been proposals to redesign electron accelerators to function within magnetic fields [4, 20], at this stage all systems appear to be using conventional equipment in conjunction with magnetic shielding. It is therefore interesting to note that the external dimensions of this triode gun are much larger than the model which was previously investigated for shielding [23]. This means that as well as being more sensitive to magnetic fields, this gun would be more difficult to shield, as shielding could not be placed as close to the beam (unless internal shielding was incorporated during the manufacture stage). Whilst a triode gun design will require more complex housing to some extent, it is not clear how easily the gun could be made smaller, or whether other triode guns would also be this large.

We used homogenous parallel fields to approximate the true heterogeneous fields of the MRI scanner in this work. For the in-line case, a key difference is that a changing magnetic field strength causes magnetic lensing effects. However, because the gun is very small (~30mm) with respect to the rate at which the magnetic field changes, we would expect that homogenous fields provide a very close approximation to the heterogeneous fields. We have verified that this is a valid approximation case by comparing the gun solution in the heterogeneous field of a 1 Tesla superconducting magnet to a homogenous parallel field of the same average strength at SIDs between 0.9 and 2.1 meters (data not shown). For all SIDs greater than 1 meter, the difference in current loss was within 5% (absolute mean difference of 1.5%). For SIDs larger than this, the field gradients become much larger and larger differences are seen; however using homogenous parallel fields still predicts the current loss to within 15%.

The effects of electron beam collisions with the control grid were not modelled in this work. These include heating of the grid, electron scatter, and secondary electron emission. The former phenomenon leads to grid damage and heating. If the grid becomes hot enough, it can become a thermionic emitter and source of ‘dark current’ (i.e. current which occurs even when the gun should be gated). Moreover, secondary electron emission will occur from the grid, which degrades the overall beam quality [24]. Although the electrons hitting the grid have energy of less than 100 eV (determined by the grid voltage) the grid is thermally light and in a vacuum, so it is plausible that substantial heating occurs. A back-of-the envelope calculation would suggest that intersection of the emitted electrons with the grid may not be a major issue in terms of heating; the forward beam power at the grid is of order 30W, of which only ~20% or 6 W will actually hit the grid (Section 3.2, Figure 4C). Meanwhile, the power from backward accelerated electrons can be on the order of kilowatts [20]. The fact that the duty cycle of the gun is of the order 0.001 will further mitigate heating of the grid [6]. Nevertheless, the impact of collisions between the emitted electrons and control grid are clearly significant enough in real world applications to have been the subject of a large body of work aimed at reducing them - mainly by utilising some form of ‘shadow grid’ which reduces the number of electrons intercepting the grid [25-28]. Both mechanical and chemical

shadow grids have been proposed in the literature. The gun modelled in this work has no visible form of shadow grid; to check if some kind of chemical shadow grid may be present, we scanned the cathode surface using Raman spectroscopy. Raman spectroscopy provides information on chemical bonds via laser probing. We were not able to detect any presence of any shadow grid using this technique - however we cannot say with absolute certainty that one is not present.

In this work, we have developed and presented an FEM model of a clinical triode electron gun which matches available experimental data to within 6%. This is the first time we are aware of that a model of a clinical triode gun has been published in the literature. We have characterised the performance of this gun in both perpendicular and in-line magnetic fields, and from this data and previously published data, made a number of important inferences regarding the design of electron guns for integrated MRI-Linac systems.

Acknowledgements

Brendan Whelan would like to acknowledge the Centre for Oncology Education and Research Translation (CONCERT) and Cancer Institute NSW for scholarship support. We would like to thank Ben cooper, Bill Morrison, Alex Quinn, Curtis Allan (Altair), Stephen Gierman, Yves de Deene & Barry Napthali for interesting discussions and help with measurements. This study received funding support from the NHMRC (program grant APP1036075) and the NIH (R21 EB015957-02).

Conflicts of interest

Dragos Constantin is currently employed by Varian Medical Systems. Rebecca Fahrig is currently employed by Siemens Healthcare.

References

1. Aubin, J.S., D. Santos, S. Steciw, and B. Fallone, *Effect of longitudinal magnetic fields on a simulated in-line 6 MV linac.* Medical physics, 2010. **37**(9): p. 4916-4923.
2. Aubin, J.S., S. Steciw, C. Kirkby, and B. Fallone, *An integrated 6 MV linear accelerator model from electron gun to dose in a water tank.* Medical physics, 2010. **37**(5): p. 2279-2288.
3. Constantin, D.E., R. Fahrig, and P.J. Keall, *A study of the effect of in-line and perpendicular magnetic fields on beam characteristics of electron guns in medical linear accelerators.* Medical physics, 2011. **38**(7): p. 4174-4185.
4. Constantin, D.E., L. Holloway, P.J. Keall, and R. Fahrig, *A novel electron gun for inline MRI-linac configurations.* Medical physics, 2014. **41**(2): p. 022301.
5. St Aubin, J., S. Steciw, and B. Fallone, *Effect of transverse magnetic fields on a simulated in-line 6 MV linac.* Physics in medicine and biology, 2010. **55**(16): p. 4861.
6. Karzmark, C., C.S. Nunan, and E. Tanabe, *Medical electron accelerators*. 1993: McGraw-Hill.
7. Kamino, Y., S. Miura, M. Kokubo, I. Yamashita, E. Hirai, M. Hiraoka, and J. Ishikawa, *Development of an ultrasmall C-band linear accelerator guide for a four-dimensional*

image-guided radiotherapy system with a gimbaled x-ray head. Medical physics, 2007. **34**(5): p. 1797-1808.
8. Bigolas, J., S. Getka, A. Kucharczyk, S. Kulinski, W. Maciszewski, M. Pachan, E. Plawski, and A. Soltan. *The SW accelerating structure of variable energy electron linac for medical application*. in *Particle Accelerator Conference, 2001. PAC 2001. Proceedings of the 2001*. 2001. IEEE.
9. Evans, P.M., J.R.N. Symonds-Tayler, R. Colgan, G.D. Hugo, N. Letts, and C. Sandin, *Gating characteristics of an Elekta radiotherapy treatment unit measured with three types of detector.* Physics in medicine and biology, 2010. **55**(8): p. N201.
10. Freislederer, P., M. Reiner, W. Hoischen, A. Quanz, C. Heinz, F. Walter, C. Belka, and M. Soehn, *Characteristics of gated treatment using an optical surface imaging and gating system on an Elekta linac.* Radiation Oncology, 2015. **10**(1): p. 68.
11. Fallone, B.G. *The Rotating Biplanar Linac–Magnetic Resonance Imaging System*. in *Seminars in radiation oncology*. 2014. Elsevier.
12. Keall, P.J., M. Barton, and S. Crozier. *The Australian Magnetic Resonance Imaging–Linac Program*. in *Seminars in radiation oncology*. 2014. Elsevier.
13. Lagendijk, J.J., B.W. Raaymakers, and M. van Vulpen. *The Magnetic Resonance Imaging–Linac System*. in *Seminars in radiation oncology*. 2014. Elsevier.
14. Baillie, D., J.S. Aubin, B. Fallone, and S. Steciw, *FEM design and simulation of a short, 10 MV, S-band Linac with Monte Carlo dose simulations.* Medical physics, 2015. **42**(4): p. 2044-2053.
15. Baillie, D., J.S. Aubin, B. Fallone, and S. Steciw, *Feasibility of producing a short, high energy s-band linear accelerator using a klystron power source.* Medical physics, 2013. **40**(4): p. 041713.
16. Crijns, S., J. Kok, J. Lagendijk, and B. Raaymakers, *Towards MRI-guided linear accelerator control: gating on an MRI accelerator.* Physics in medicine and biology, 2011. **56**(15): p. 4815.
17. Heerkens, H.D., M. van Vulpen, C.A. van den Berg, R.H. Tijssen, S.P. Crijns, I.Q. Molenaar, H.C. van Santvoort, O. Reerink, and G.J. Meijer, *MRI-based tumor motion characterization and gating schemes for radiation therapy of pancreatic cancer.* Radiotherapy and Oncology, 2014. **111**(2): p. 252-257.
18. Pierce, J.R., *Theory and design of electron beams*. 1954: van Nostrand.
19. Lawson, J.D., *The physics of charged-particle beams.* 1977.
20. Whelan, B., S. Gierman, L. Holloway, J. Schmerge, P. Keall, and R. Fahrig, *A novel electron accelerator for MRI-Linac radiotherapy.* Medical Physics, 2016. **43**(3): p. 1285-1294.
21. Lee, B.-N., P. Buaphad, S.H. Lee, S. Lee, B.C. Lee, K. Song, Y. Kim, H.D. Park, and S.-S. Cha, *Status of KAERI 6 MeV 9.3 GHz X-Band Electron Linac for Cancer Treatment System.* 2014.
22. Waldron, T. *Function Requirements for IMRT*. 2003; Available from: https://www.aapm.org/meetings/03SS/Presentations/Waldron.pdf.
23. Santos, D., J.S. Aubin, B. Fallone, and S. Steciw, *Magnetic shielding investigation for a 6 MV in-line linac within the parallel configuration of a linac-MR system.* Medical physics, 2012. **39**(2): p. 788-797.
24. Shih, A., J. Yater, C. Hor, and R. Abrams, *Secondary electron emission studies.* Applied surface science, 1997. **111**: p. 251-258.

25. Lien, E.L., G.V. Miram, and R.B. Nelson, *Grid controlled electron source and method of making same*. 1976, US patent 3967150.
26. Miram, G., *Grid coating for thermionic electron emission suppression*. 1981, US patent 4263528.
27. Elfe, T., O. Koppius, and R. Willis, *Electron gun with masked cathode and non-intercepting control grid*. 1974, US patent 3818260.
28. Tornoe, R.N., *Cylindrical cathode with segmented electron emissive surface and method of manufacture*. 1982, US patent number 4359666.

Chapter 4
Passive Magnetic Shielding

1. Introduction

MRI-Linac therapy holds the promise of greatly improved cancer treatment outcomes by coupling the exquisite soft tissue contrast, high temporal resolution and functional imaging capabilities of MRI with the established therapeutic gains of radiotherapy [1]. However, realisation of these benefits requires that a medical linear accelerator (linac) function within the fringe field of a superconducting MRI magnet. Many linac components are not compatible with an external magnetic field. Prominent examples of accelerator components which may not function correctly in an external magnetic field include the linac electron gun [2-4], accelerating waveguide [4-6], and MLC motors [7]. To minimise the deleterious effects of external magnetic fields on linac performance, there are two options: redesign the components in question, or ensure that they operate in a sufficiently low magnetic field that acceptable performance is achieved. Means by which one may ensure that components operate in sufficiently low magnetic field are to incorporate a low field region into the very early stages of magnet design, and to utilise magnetic shielding to create a low field environment around the component in question – the focus of this chapter.

Two approaches exist to magnetic shielding; active and passive [8]. Active magnetic shielding utilises current loops to cancel the local magnetic field within the shielded area, whilst passive magnetic shielding utilises materials with a high magnetic permeability to redistribute the magnetic field distribution such that the field in the shielded area is lowered. For magnetic shielding in MRI-Linac systems, the challenge is to design the magnetic shield such that the very precise magnetic field needed for MRI imaging is not unduly perturbed. Both passive and magnetic shielding have been proposed and implemented in MRI-Linac systems [9, 10], although passive shielding has a number of features which make it attractive. Firstly, passive shielding takes up very little space and can be made (indeed in general should be made) conformal with the volume which is to be shielded, which is highly desirable in the space constrained design of medical linacs. Secondly, passive shielding does not generate heat. Heat management is crucial for stable accelerator operation, as deviations in temperature result in thermal expansion and hence affect resonant frequency. Also, the current loops themselves must be sufficiently cooled to avoid reaching dangerous temperatures. Finally, passive shielding does not require an external power source. On the other hand, one attractive aspect of active shielding is the fact it can be rapidly switched on and off. This could be beneficial in an MRI-Linac approach whereby the imaging and treatment functions were carried out separately. However, overall it seems likely that passive shielding will be the preferred approach in MRI-Linac systems.

The utility of passive magnetic shielding has been known since at least the 16th century [11]. As such, there is a large amount of work in the literature discussing the mechanisms and design of passive magnetic shielding. However, much of this work was before the advent of modern computational electrodynamics, which resulted in some limiting assumptions being made (very simple geometries and single valued permeability). In addition there has been very little published work looking at magnetic shielding for MRI-Linac systems – results of a passive magnetic shielding study in MRI-Linac was previously studied using modern computational techniques in [8]. However, the scope of this work was quite specific, and a

number of aspects remain unaddressed, such as the impact of: different shield geometries, distance from the scanner, different shielding materials, shield layering, and different MRI-Linac orientations. In this work, we aimed for a more general examination of passive magnetic shielding, starting from first principles and using Finite Element Modelling to examine the impact of different shield geometries, shield materials, MRI-Linac configurations. In so doing, we hope to provide a useful data set which may be used for the optimal integration of linacs with MRI. This work may also prove useful in other research areas in which magnetically sensitive components must be used in close proximity to MRI, such as interventional imaging suites [12] or linacs in adjoining rooms to MRI systems [13].

2. Theory and Background

The most widely adopted conceptual model of passive magnet shielding is the magnetic shunting or magnetic circuit description, which describes the shield as providing a low reluctance (and hence favourable) path for magnetic field lines to travel [14]. By providing a more favourable pathway for magnetic field lines to travel through, the magnetic field inside the shield can be reduced. Whilst this description is often useful, there are circumstances when the magnetic shielding effect is not well described by such an approach. An alternative conceptual model which can also be useful is to consider that often the field arising from the magnetised shield tends to oppose the applied field (the so called 'demagnetising field'). The physical interpretation of magnetic shielding effects is explored in detail in ref [15]; to quote directly from this paper, "[No single] listed mechanism taken separately can account for all aspects of magnetic shielding. This effect can be described adequately only if a complex approach is endorsed."

In order to describe the effectiveness of a magnetic shield, we introduce the shielding factor;

$$\mathbf{K = B_{NoShield}/B_{WithShield}} \qquad \textbf{Equation 15}$$

B can refer to either a point measurement or an average value within the sensitive volume. Analytic solutions for the shielding factor provided by simple geometries can be derived analytically; equations 2 & 3 show the approximate shielding factor for spherical [15] and cylindrical [16] shells in a constant field. In these formulas R is the external radius and t is the shell thickness. Equation 3 is for an infinite cylinder in a field applied perpendicular to the axis.

$$\mathbf{K}_{sphere} \simeq \frac{2t\mu_r}{3R} + 1 \qquad \textbf{Equation 16}$$

$$K_{cyl,\perp} \simeq \frac{\mu_r t}{2R} + 1 \qquad \textbf{Equation 17}$$

The analysis for a cylinder in axial magnetic fields is more difficult; however in ref [17] a formula for a finite cylinder with closed ends is presented which agrees quite well with numerical simulation; this formula is outlined in equation 4.

$$K_{cyl,\parallel} \simeq 1 + \left[\frac{4N^{elipse}K_{cyl,\perp}}{f}\right] \qquad \text{Equation 18, where}$$

$$N^{elipse} \simeq \left[\frac{1}{G^2-1}\right] * \left[\frac{G}{\sqrt{G^2-1}} * \ln\left(G + \sqrt{G^2-1}\right) - 1\right] \qquad G = \frac{L}{2R} \quad f = 1 + .5/G$$

N^{elipse} is the demagnetisation factor of an ellipse having the same aspect ratio as the shield; the demagnetisation field is the field which results from the magnetisation of the object.

In certain situations, the shielding factor can be substantially improved by utilising multiple shells of high permeability material separated by low permeability material (e.g. air, copper, foam, etc.) [18]. The mechanism for this effect can be explained in terms of the magnetic shunting model as follows (paraphrased from ref. [18]); In order to shield a given volume, it would be effective to surround it with material of either higher or lower permeability than the surrounding volumes (in the latter case, invoking the magnetic shunting model, field lines would tend to travel around the shield instead of through it). Although no materials with permeability substantially lower than air are known (except superconductors), a similar effect can be produced by utilising gaps of low permeability material spaced between high permeability materials. In layered models, the thickness of the gap as well as the thickness of each layer becomes important. In general, analytical work suggests that optimal or close to optimal shielding is obtained when each the gaps and layers are around the same size. An analytic expression for a two shell sphere has been derived in ref. [19], as well as for cylinders in perpendicular [20] and parallel fields [21]. The resultant formulas are quite complicated, and will not be reproduced here. Solutions for an arbitrary number of shells have also been derived and can be found in the provided references.

Equations 2-4 provide an excellent starting point for the design of magnetic shields, illustrating the general principles that the shielding factor increases as the magnetic permeability and thickness (or layers) of a shield is increased, and decreases as the volume being shielded increases. However, a number of factors limit the utility of these formulas for real life applications. Materials exhibiting high enough permeability to be useful as magnetic shield are ferromagnetic. Ferromagnetic materials are nonlinear; which means that μ_r in equations 2-4 is a function of the applied field. The relationship between applied field and permeability is conventionally described by BH curves, where $\mu = B/H$ and $\mu_r = \mu/\mu_0$ [22]. Examples of BH curves for two commonly used magnetic shielding materials are shown in Figure 1. In addition to the nonlinear behaviour of these materials, it can be seen that at a certain point they become magnetically saturated; this occurs when a material cannot become any more magnetised. Ferromagnetic materials are quite hard to account for in the analytic approach, and even more so in the multilayer case where the applied field, and hence permeability will substantially change from layer to layer.

Another limitation of the analytic approach is that it can only be used for very simple geometries. In real world situations (such MRI-Linac therapy) it is often necessary to utilise a shield which has complex geometry and various openings. Calculating the impact of such perturbations using the analytic approach becomes very difficult. Given these limitations, a more robust approach to shield design is to utilise computational electrodynamics such as

Finite Element Modelling (FEM). Using FEM, non-linear materials and complex geometries can easily be solved.

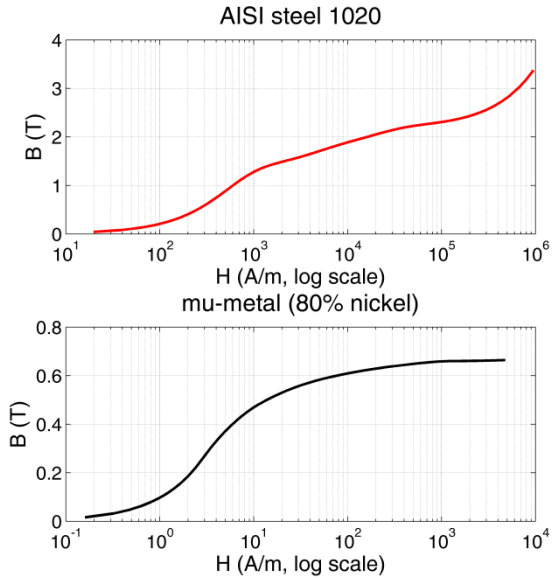

Figure 34: BH curves for Carbon Steel 1020 and mu-metal.

In general, the design goal for a passive shield is simply to optimise the shielding factor in the sensitive volume. However, shielding for MRI Linac systems introduces an additional constraint: the shield must not cause large perturbations to magnetic field in the imaging volume of the scanner. This is because MRI scanners utilise an extremely homogenous magnetic field to create images, and the introduction of magnetic material in the vicinity of the MRI scanner can cause perturbations in this field which can render the scanner useless. The homogeneity of MRI magnets is conventionally assessed in parts-per-million (PPM) within a set Diameter of Spherical Volume (DSV), even though the impact of inhomogeneity is also dependent on field strength, such that "…a 1.0 T magnet with a homogeneity of 10 ppm and a 10 T magnet with 10 ppm of homogeneity would have image distortion and signal loss problems due to static field inhomogeneities that differ in magnitude by approximately a factor of ten" [23]. To produce high quality images, a value of less than ~10 is required [23]. In order to achieve these very low PPM values, a process known as shimming is utilised. Even in the absence of a magnetic shield, shimming is typically required after the installation of a new magnet to account for factors such as magnetic impurities, machine tolerance errors, thermal contractions, mechanical stress during transport etc. [23].

An upper value of around 300 ppm from a magnetic shield has typically been quoted as being recoverable by modern shimming techniques [8, 9, 24], although we could not actually find any published literature where this value was tested. However, this value is useful to interpret the results for induced distortion which are presented in section 4.6.

3. Methods

3. 1. Comparison of analytical formulas to numerical simulation

In order to verify both the numeric simulation process and formulas 2-4, simple 2D models were implemented in Comsol version 4.4 using the reduced field formulation of the 'magnetic fields, no currents' module. For cylindrical shields in perpendicular fields, an infinite cylinder in 2D was tested. For cylindrical shields in axial fields, and for spherical shields, 2D axisymmetric models were implemented. The relative permeability and thickness of each shield was 2000 and 2 mm, and they were tested in a background field of 200 G. The shield radii were varied between 25 and 200 mm. For the cylindrical shield in axial fields, the shield length was 350 mm.

3. 2. Geometry and sensitivity of medical Linacs

In order to design a magnetic shield, two details are essential; the magnetic sensitivity and the dimensions of the object being shielded (the latter constrain the shield geometry). The dimensions of a 600C Varian Linac are shown in Figure 2. The outer dimensions were measured directly from a 600C linac, whilst the inner dimensions (which in this work are used for visualisation only) are based on ref. [5]. The linac can be considered to comprise two sensitive sub-volumes; the electron gun, and the accelerating waveguide. As each of these components show different sensitivity to magnetic fields, the shielding requirements will be considered separately. The radius of the beam line along the linac is 2.5 mm, however to account for positioning uncertainty between the Linac and MRI scanner and other errors, we will consider a 5 mm radius along both the electron gun and linac as needing to be shielded. These regions are indicated in Figure 2C.

For in-line fields, the electron gun current begins to drop off between 50-100 G [2, 4]. As explained in chapter 3, there is some uncertainty in these values in the absence of accurate geometrical measurements of the anode drift tube. Therefore, in this work we will assume that the in-line fields in the electron gun must be less than 50G. In-general, in-line magnetic fields do not strongly impact the target current of this type of accelerating waveguides [4], although as explained in ref. [6] magnetic lensing effects occur which could have major ramifications if not accounted for. Nevertheless, in this work we will assume that for the in-line direction, only the electron gun needs to be shielded. For perpendicular fields, functionality of both the electron gun and the Linac is strongly affected even at very low fields [2, 3]. Based on ref. [3], even at 2 G the target current begins to decay. In this work, we will therefore assume that transverse magnetic fields across entire electron gun and linac structure must be less than 2 G, although in both orientations, even lower fields than minimum standard would be preferred.

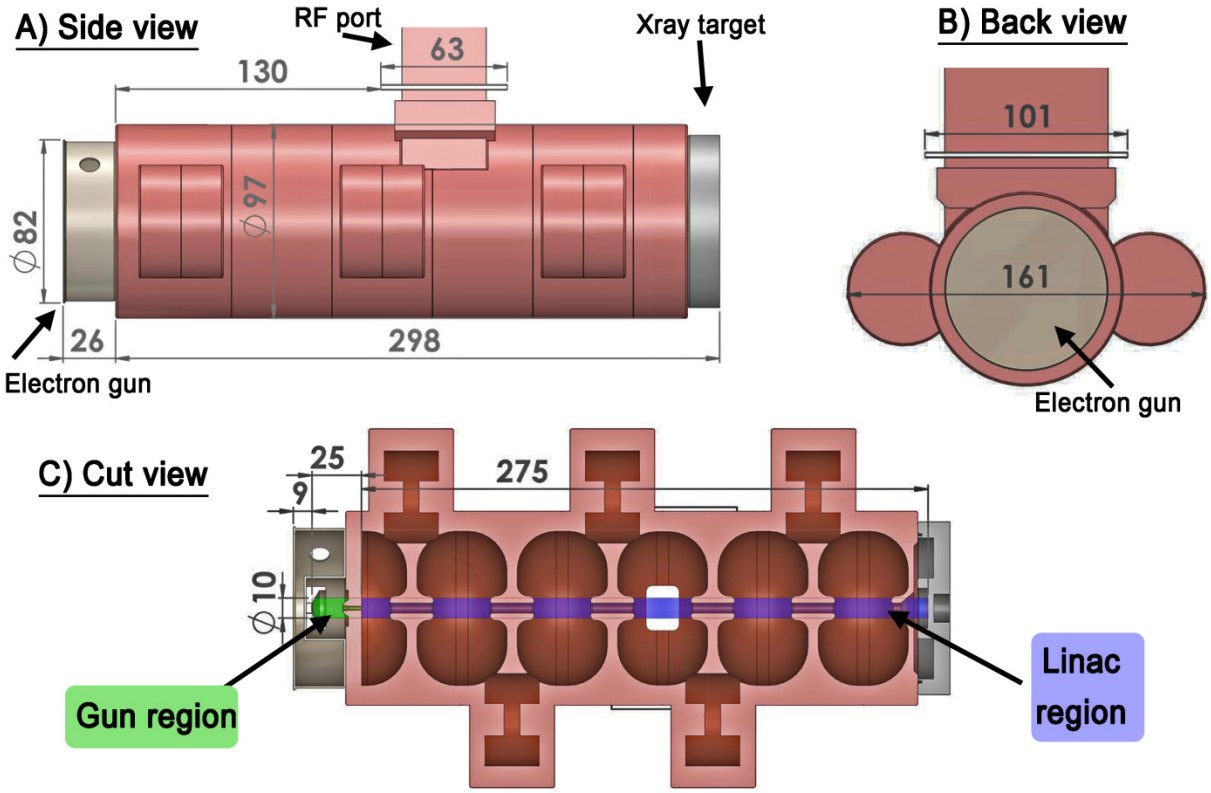

Figure 35: Dimensions of a Varian 600C Lin165/105ac. A) Side view B) Back view (from gun end) C) Cut view including the two regions which must be magnetically shielded, the gun region and the linac region. All dimensions are in mm.

3. 3. Investigation of different shield geometries

There are two basic considerations when considering appropriate shield geometry. On the one hand, it is desirable for the shield to conform as closely as possible to the volume being shielded. On the other, it is desirable for the shape of the shield to be smooth in order to maximise its efficacy as a magnetic shunt. In this section, the four different shield geometries were implemented, and the impact on the shielding factors investigated (Figure 3). All shields were designed to fit around the linac dimensions shown in Figure 2 with 2 mm of clearance. Each geometry was solved in constant perpendicular and in-line fields of 200 G for three different combinations of shield thickness and shield permeability; ST=2 and μ_r=2000, ST=4 and μ_r=10 000, and ST=2, μ_r= Carbon Steel 1020 BH curve (Figure 1). For each design, the shield length was varied between 50 and 500 mm and the shielding factor in the electron gun and linac region was recorded by dividing 200 G with the average field in the relevant volumes (Figure 2C). 200 G is quite a high field to attempt to operate an accelerator in – however, this makes it a suitably challenging environment to investigate magnetic shield design. Also, when non-linear materials are used in lower fields, the shielding performance will be better than calculated in higher fields as permeability typically increases in lower fields. For the shields longer than 150 mm, the RF input waveguide must be explicitly accounted for. In this section, the RF port was shielded with rectangular shielding to a height of 100 mm from the centre of the shields (PH in Figure 3). Perpendicular fields were always applied at right angles to the RF port to minimise the impact of the hole in the shielding. There are important cables which must somehow get through the shield; in particular the high

voltage line for the electron gun, and water cooling for the linac. We did not explicitly account for these cables in this work; however with minor modification they could travel either through the open end of the shield, or through the RF port hole.

As explained in ref. [8], the electron gun on a 600C linac is surrounded by a ferromagnetic case which is designed to shield the gun from earth's magnetic field. In this work, this casing was not included in the simulations as the primary focus was the design of standalone shielding, and including this casing would make these results more difficult to interpret. Also, the material type and magnetic properties are not known. However, it should be noted that for shielding of a 600C linac the gun shielding factor would in fact be higher than the results in this work due to this additional shielding.

As in section 3. 1, the 'magnetic fields, no currents' module with reduced field formulation was used to specify a constant magnetic field. The models were implemented in 3D. Each model was solved in a background cylinder of radius 2 m and height 3 m. Further increasing the background size had negligible impact on the results. In order to reduce the size of the simulation domain, symmetry planes were exploited using the 'magnetic insulation' and 'zero scalar potential' boundary conditions wherever possible. In practice this usually meant that half the model needed to be solved. Second order quadratic mesh elements were used to mesh the geometry. In the shield, the maximum mesh element size was set to the shield thickness and the minimum element size to the shield thickness divided by 10. We verified that these mesh settings adequately resolved the physics by running a simulation with doubled mesh density; the shielding factors changed by less than 0.05%. For simulations where ferromagnetic materials are used, an iterative solving technique is required to account for the non-linear nature of the materials. In this work, we used the 'double dogleg' method. The model was considered solved when the relative residual was less than 1E-2; the models in this section typically resolved within 30 iterations.

3. 4. Investigation of RF port shielding and partially closed shields

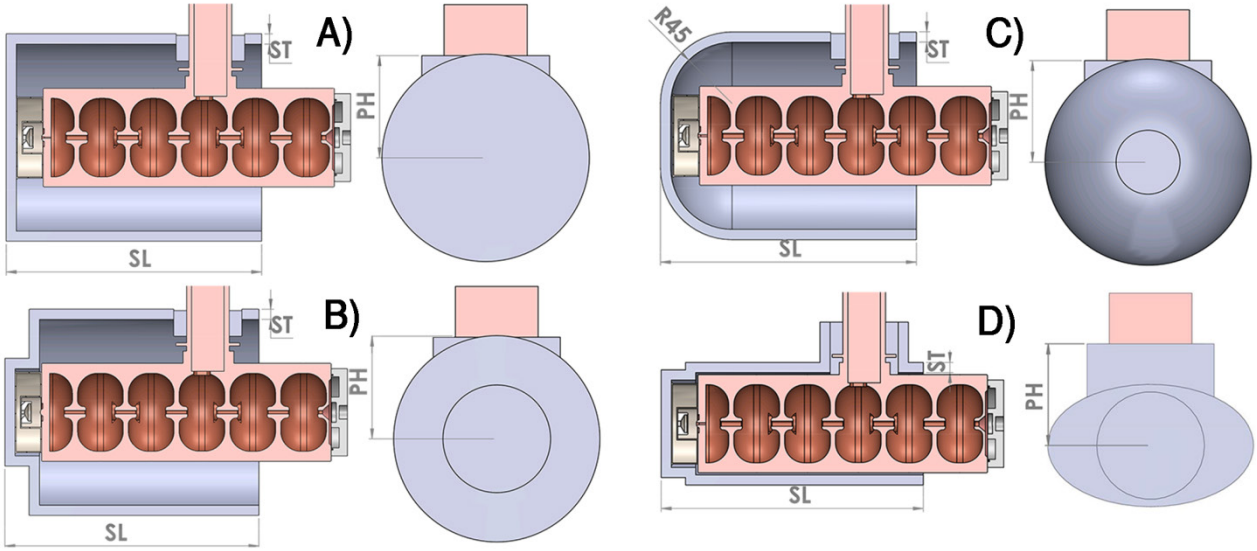

Figure 36: The four different shield geometries which were tested in this work. PH is the height of the port shielding, SL is the shield length, and ST is the shield thickness. All other dimensions are such that there is 2mm clearance around the dimensions shown in Figure 2. In this figure, ST is set to 10 mm, SL 250 mm, and PH to 100 mm.

As explained previously, it is necessary to have (at least) two openings in the shields; one for the RF port, and one for the beam to exit from (although the RF port only needs an opening for shields longer than 150 mm). In this section, we quantify the impact of these openings, and the impact of altering the height of the RF port shielding (PH in Figure 3) and partially closing the open end of the shield. Based on the results from section 3. 3 (section 4.2), Shield D was the best performed for all scenarios. Therefore, the remainder of this study focuses on this design.

To test the impact of the port opening, the shield was solved with and without the port opening, and the height of the port shield in the latter case was varied from 40 mm to 100 mm. Note that for shield D, having PH set to 40 mm is equivalent to simply having an open hole in the shield, whilst PH=100 mm is what is shown in Figure 3D. To test how much improvement could be derived from closing the end cap, the end cap was partially closed, with a circular opening in the middle of radius 10 to 50 mm. Although the shield cannot be fully closed as the beam needs to exit, it could be partially closed if this proves beneficial. As shown in 4.2, for perpendicular fields the shielding factor increases as a function of shield length, whilst for axial fields the optimal length of the shield depends on the thickness and permeability of the shielding material. However, in order to close the front of the shield the shield must be longer than the Linac; therefore, in this section we held the shield length constant at 400mm. Steel 1020 (Figure 1A) was used as the shielding material.

3. 5. Investigation of multiple shield layering

Based on the theoretical studies of shielding, it has been shown that it is often beneficial to split a given thickness of ferromagnetic material in multiple layers [25-27]. In this section, we investigated the impact of layering the shield for axial and perpendicular fields.

Shield D was solved for 1 to 5 layers, with the total thickness of ferromagnetic material held constant (i.e. for the 1 layer shield the thickness of the layer was 5 mm, and for the 5 layer shield each layer was 1 mm thick). The gap between each layer was initially the same as the thickness of each layer; as it was indicated in the literature that this should be the optimal or close to optimal configuration [27]. In this section, the axial shield length was 200 mm and the perpendicular shield length was 400 mm. This decision was based on the results from sections 3. 3 & 3. 4, which indicated that for axial fields better shielding factors of the gun were achieved with shorter shield length. However, it is important to note that this decision is based on the assumption that there is no requirement to shield the linac in the in-line configuration. It is important to take into account the non-linear nature of shielding materials for layered simulations, as the field (and hence permeability) at each layer will change. Steel 1020 was used for these simulations. Next, to test the optimal spacing of each layers, the spacing for the 5 layer shield was varied between 1 and 5 mm. Finally, we tested the impact of using mu-metal (Figure 1) for the inner layers of the shields. Mu-metal is an extremely affective magnetic shielding material, but only when the applied field is less than around ~10 G; beyond this it magnetically saturates (note that mu metal can still be used for higher fields, but higher thickness or multiple layers must be used). One advantage of the layering approach is that the outer layers can be used to reduce the field to the point that the mu metal becomes useful. For this simulation, the layer spacing was set to 10 mm based on the results of changing the layer spacing.

3. 6. Impact and efficacy of passive magnetic shielding for a 1.0 T MRI magnet

In order to put the magnetic shielding principles which are developed in the preceding sections into action, we utilised a previously developed Comsol model of the 1.0 Australian MRI-Linac program magnet [24]. To test the impact and efficacy of passive magnetic shields, various shield designs were incorporated into this model, and the distortion, average magnetic fields in the volumes of interest, and shielding factors were calculated. The elliptical shield shape (Figure 3D) was used in all cases, and all simulations were carried out in both the in-line (axial) and perpendicular orientations. Carbon steel 1020 (Figure 1A) was used unless otherwise stated. The simulations which were run were as follows: firstly, a single layer shield of length 400 mm, thickness 20 mm, and PH=100mm was solved at SIDs from 900 mm to 2500 mm every 100mm. The SID was defined as the end of the linac region (Figure 2). Based on these results, it was decided to design passive shielding for an SID of 1.5 m in both orientations. This is quite a challenging location to operate in for this magnet as the fringe fields are still quite high; in the absence of any shielding, the average fields in the gun and linac are respectively 715 G and 985 G (perpendicular) and 413 G and 504 (axial). The following simulations were carried out to optimise the shield design in each orientation. First,

the length of the same shield described above was varied between 100 mm and 400 mm in 50 mm steps. From here, the design of the shields in each orientation began to diverge based on the results of each simulation (the parameters used in each case are described in the results section 4.6). Next, shield thickness was varied between 1 mm to 20 mm every 1 mm), Shield layering was investigated for up to five layers, and layer spacing was varied between 3 mm to 12 mm every 3 mm. In each case the optimal parameters were kept for the simulations that followed. Finally, we investigated the impact of replacing the inner shield layer with mu-metal.

4. Results

4.1. Comparison of analytical formulas to numerical simulations

Figure 4 shows the comparison between the analytic formulas 2-4 and Comsol simulations. Very good agreement is obtained for a sphere and cylinder in perpendicular fields, whilst somewhat lower agreement is obtained for a cylinder in longitudinal fields. This is not particularly surprising, given that this is the most difficult case to analyse analytically. It is also notable that whilst the analytic model appears to slightly overestimate the shielding factor, it does quite accurately predict the shape of the curve. An interesting feature of both the analytical and numerical results for a cylinder in axial fields is the peak in shielding factor for a radius of ~75 mm. Considering equation 4, this maximum can be understood as by considering the interplay between the transverse shielding factor and the demagnetisation factor. The former decreases as the radius increases, whilst the latter increases. Also shown in Figure 4 are examples of the magnetic field distributions in each case, with field lines shown in black. It can be seen that the field lines tend to enter the material at around 90 degrees, as expected from magnetostatic theory [28].

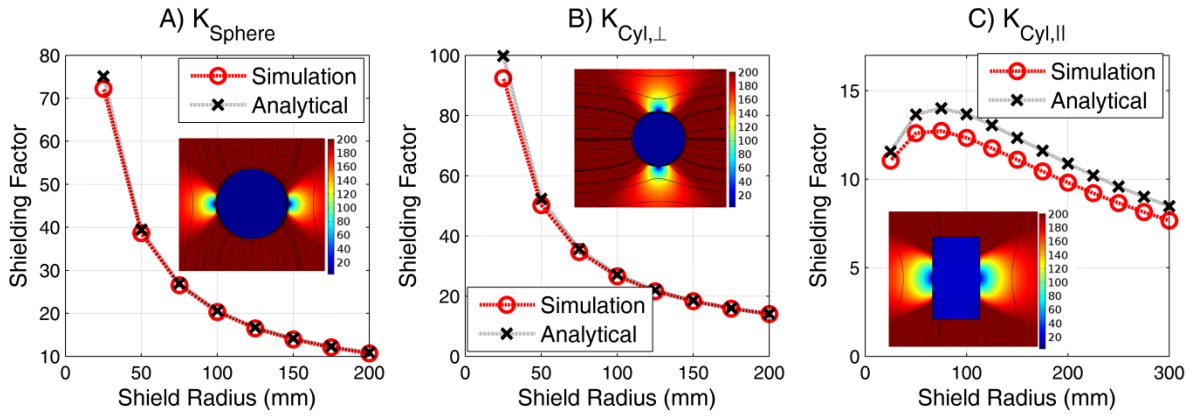

Figure 37: Theoretical versus numerical shielding factors as the shield radius is vared for for A) a sphere B) a cylinder in perpindicular fields, and C) a cylinder in axial fields. In each case, an example of the field distribution and field lines around the shield is shown. The units on the colorbar in these images are in Gauss.

4.2. Shield design in constant magnetic fields

Figure 5 shows the shielding factor for axial fields in both the gun and linac for three different factors of shield thickness and permeability. Figure 5A and B are for linear materials, whilst C is for (non-linear) Steel 1020. It can be seen that, similarly to Figure 4C in

section 4.1, there is a peak in the shielding factor as a function of length, and that this peak is dependent on both the permeability and the thickness of the shield. This maximum appears particularly pronounced for ferromagnetic materials. Of the four shields tested, the simple cylinder (shield A) performs the worst. Shields B and C offer a slight improvement in shielding around the gun, whilst shield D performs the best overall for both the gun and linac section. It also important to note that similar results regarding shielding efficacy are drawn for all three permeability and thickness combinations tested here, including non-linear materials – i.e. shield D performs the best for all situations.

Figure 6 shows the same results for perpendicular fields. Again, shields B and C offer some minor improvements in the shielding factor of the gun, whilst leaving the shielding factor of the linac unchanged. Shield D is again the best performed across all situations. It can be seen that the shielding factor for the linac rises steadily as function of length. However, increasing length also results in increasing the size of the system and amount of ferromagnetic material. Based on these results, it would seem that to achieve good shielding of the linac in the perpendicular configuration, a good compromise for shield length is around 400 mm. If the length increases beyond this, the shielding factor still increases, but at a slower rate. As was the case for the axial results, similar conclusions about the different shielding geometries are drawn regardless of the particular thickness or permeability of the shield tested.

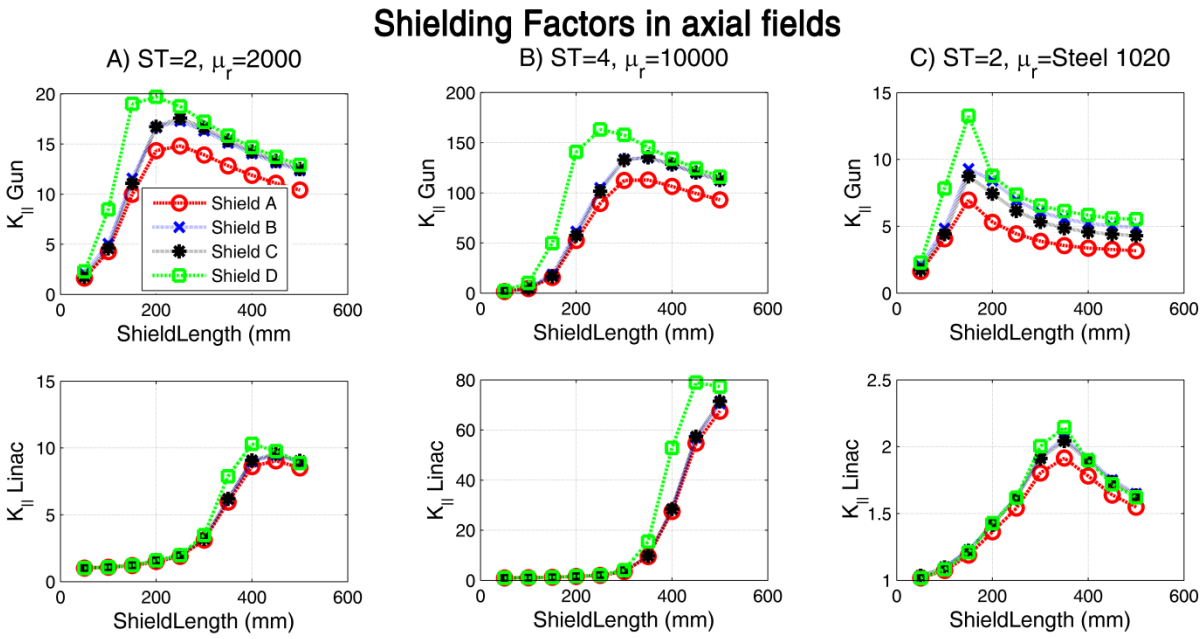

Figure 38: Shielding factors in axial fields for each of the four shield geometries shown in Figure 3. Different combinations of permeability and thickness are plotted; A) thickness =2 mm, μ_r=2000 B) thickness=4 mm, μ_r=10 000, C) thickness=2 mm, μ_r=BH curve of carbon steel (i.e. the non-linear behavior is modelled here).

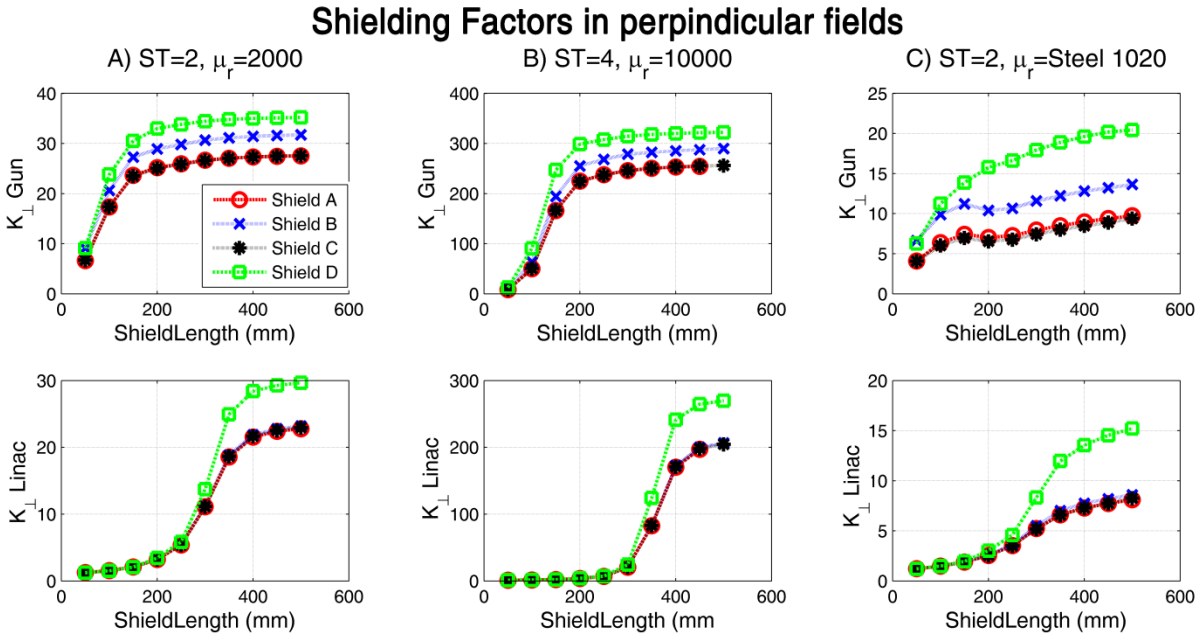

Figure 39: Shielding factors in perpendicular fields for each of the four shield geometries shown in Figure 3. Different combinations of permeability and thickness are plotted; A) thickness =2 mm, μ_r=2000 B) thickness=4 mm, μ_r=10 000, C) thickness=2 mm, μ_r=BH curve of carbon steel (i.e. the non-linear behavior is modelled here).

4.3. Investigation of RF port shielding and partially closed shields

Figure 7 shows the impact of magnetic shielding around the RF port. For both axial and perpendicular fields, the impact of shielding or not shielding the RF port is not particularly pronounced in terms of the overall effect on the shielding factors. However, a surprising result is obtained for axial fields: increasing the height of shielding around the port actually results in a decrease in shielding factor, and the peak shielding factor is obtained when the RF port hole is simply left open. Noting that the length of the shields was held constant at 400 mm in this section, the mechanism for this effect is again the demagnetising factor. When the port hole is simply left open, the same mechanism which results in a peak in the shielding factor as a function of length (Figure 5) results in increased shielding factor compared to the case with no hole left for the RF port. For perpendicular fields, the results are closer to what would be intuitively expected; as the height of shielding around the RF port increases, the shielding factor also increases. For PH=100 mm, the shielding factors are very close to what would be achieved with no hole left for the port at all.

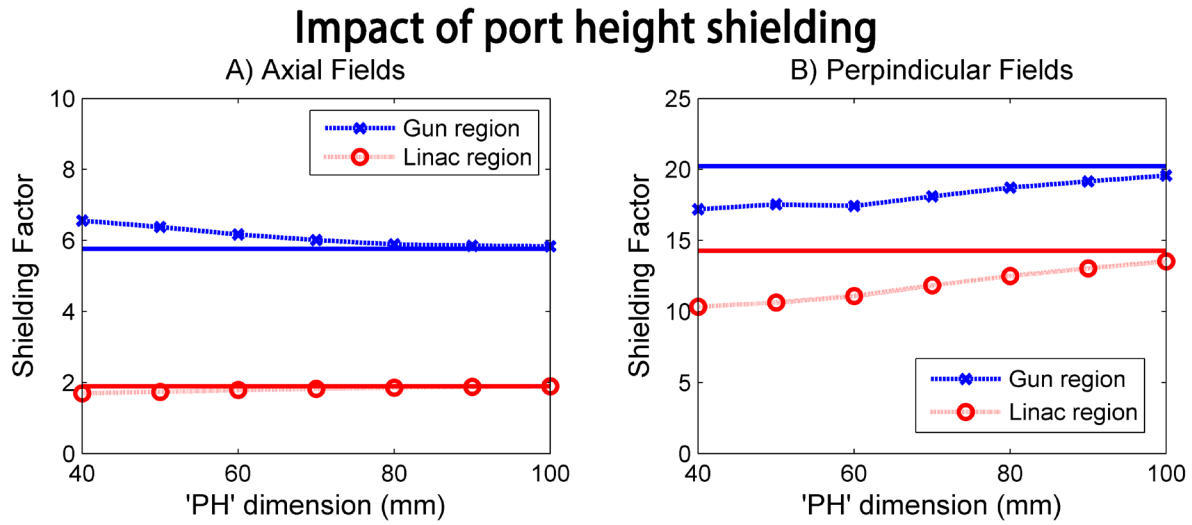

Figure 40: The impact of increasing the height of the port shielding (dimension PH in Figure 3). A) Axial fields, and B) Perpendicular fields. The solid lines show the shielding factor when the port hole is completely closed. Note that PH=40 mm is equivalent to simply having an open hole for the shield. In this data, the shield length is 400 mm, the shield thickness is 2 mm, and the material is steel 1020.

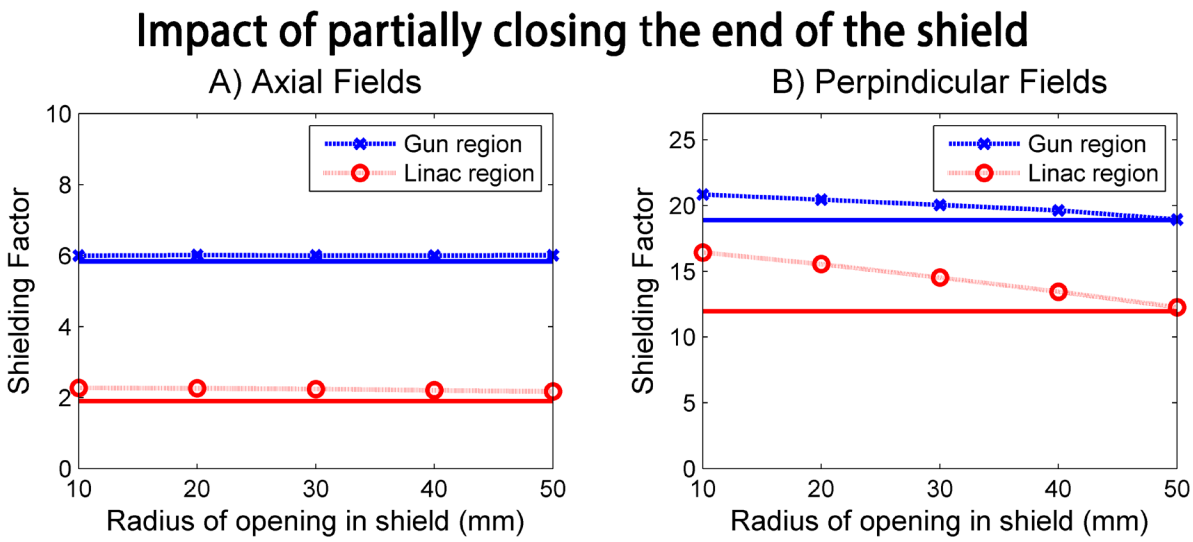

Figure 41: The impact of partially closing the open end of the shield around the linac. The solid lines show the value when the shield is completely open, whilst the dotted lines show the impact of a closed shield with a circular opening cut into the centre. In this data, the shield length is 400 mm, the shield thickness is 2 mm, and the material is steel 1020.

4.4. *Investigation of shield layering*

Figure 9 shows the impact of splitting a shield into multiple layers whilst holding the total thickness of shielding material constant. It can be seen that this approach can have a dramatic impact on the shielding factors. For axial fields, the shielding factor for the gun increases from 36.0 to 56.0 at 4 layers, an improvement of 56%. (Note the shield length for the axial case is 200 mm; in other words, no effort has been made to shield the linac region in Figure 9A). For perpendicular fields, the increase is even more pronounced; in the gun an increase of 165% is observed (62.6 to 166.1), and in the linac a 104% increase (50.4 to 102.4). Not that although the thickness of shielding is held constant, there will be a small change in the total volume of shielding material, e.g. the volume of the 5 layer shield of length 400 mm is 4% more than the single layer shield.

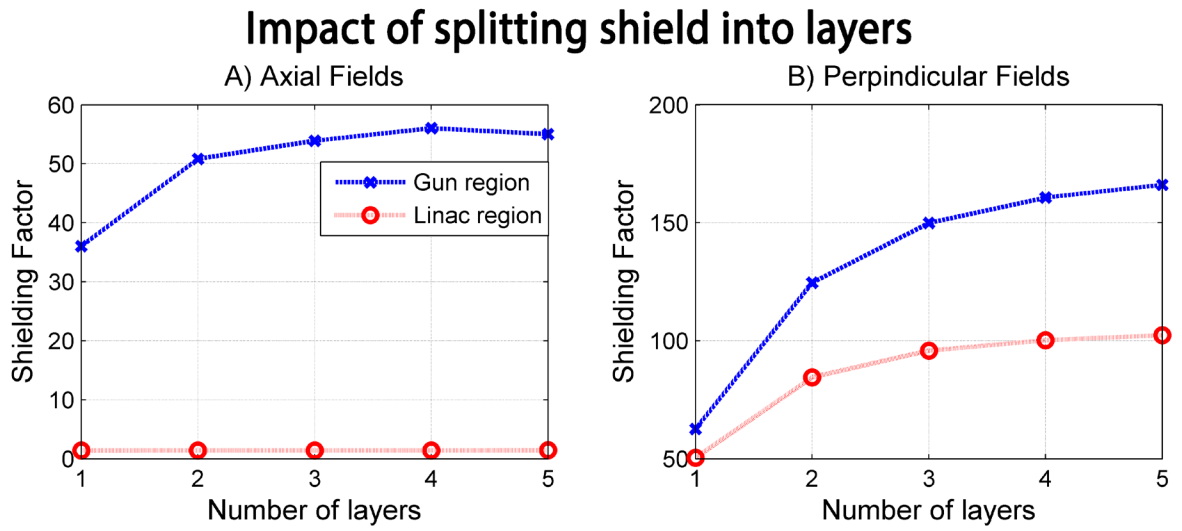

Figure 42: Impact of splitting the shield into multiple layers whilst holding the total thickness constant at 5mm. A) Axial fields. Note the length of the shield here is 200 mm, and as such the linac region is not shielded. B) Perpendicular fields. The length of the shield in this simulation is 400mm.

Figure 10 shows the impact of increasing the spacing between layers. The number of layers is held constant at 4 mm and the thickness of each layer is 1.25 mm. For axial fields (Figure 10A), the gun shielding factor has a maximum of 73.7 at a layer spacing of 10.25 mm. This represents a 105% increase on the shielding factor of the single layer shield, and a 32% increase compared to the initial spacing of 1.25 mm. For perpendicular fields, maximum shielding factor factors in the gun and linac are 404.2 and 173.0 respectively, again for a layer spacing of 10.25 mm. These values represent increases of 548% and 243% compared to the single layer case, and 252% and 72% compared to the initial layer spacing. It can also be seen that when the spacing is increased beyond 10.25 mm, the shielding factors start to decrease. This also occurs for the axial case although it is not as obvious. When the material innermost layer was replaced with mu-metal and the spacing kept set to 10 mm, even more substantial improvements in shielding were seen; for axial fields, the gun shielding factor increased to 252.7 – a further increase of 240% to the layered steel shield. For perpendicular fields, the gun shielding factor increased to 16047, and the linac shielding factor to 2612. This represents percentage increases of 3900% and 1400% on the steel shields! Figure 11shows the field distribution in the layered shields with an inner mu-metal layer.

These results show that care must be taken in choosing the number and configuration of shield layers. The obvious disadvantage in increasing either the number of layers or (particularly) the spacing of layers is the total volume taken up by the magnetic shield will increase. The significance of this depends on the extent to which one is volume constrained. In the case of MRI-Linacs, it is unlikely to be a major issue. However, what has yet to be assessed is the impact on magnet homogeneity of these techniques. The fact that utilising shield gaps substantially larger than the thickness of each layer yielded these improvements came as something of a surprise, given the literature had indicated that maximal shielding should be obtained when the thickness and the gap were equal or approximately equal [27].

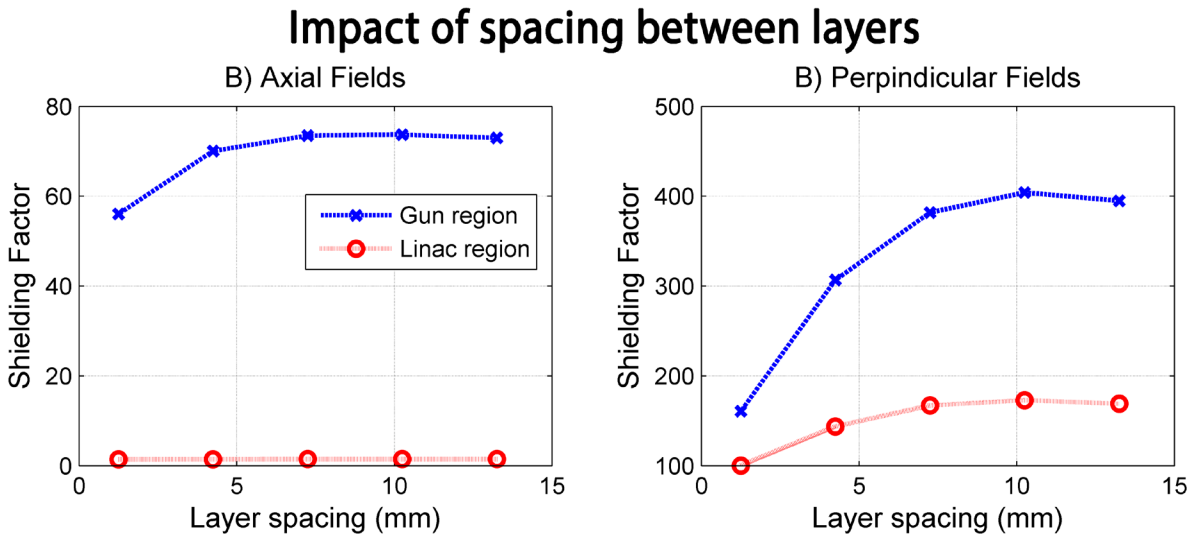

Figure 43: Impact of increasing the spacing between layers for A) Axial fields, and B) Perpendicular fields. In this data, there are 4 layers, each of thickness 1.25 mm. The shield length for the axial simulations is 200 mm, whilst for the perpendicular case it is 400 mm.

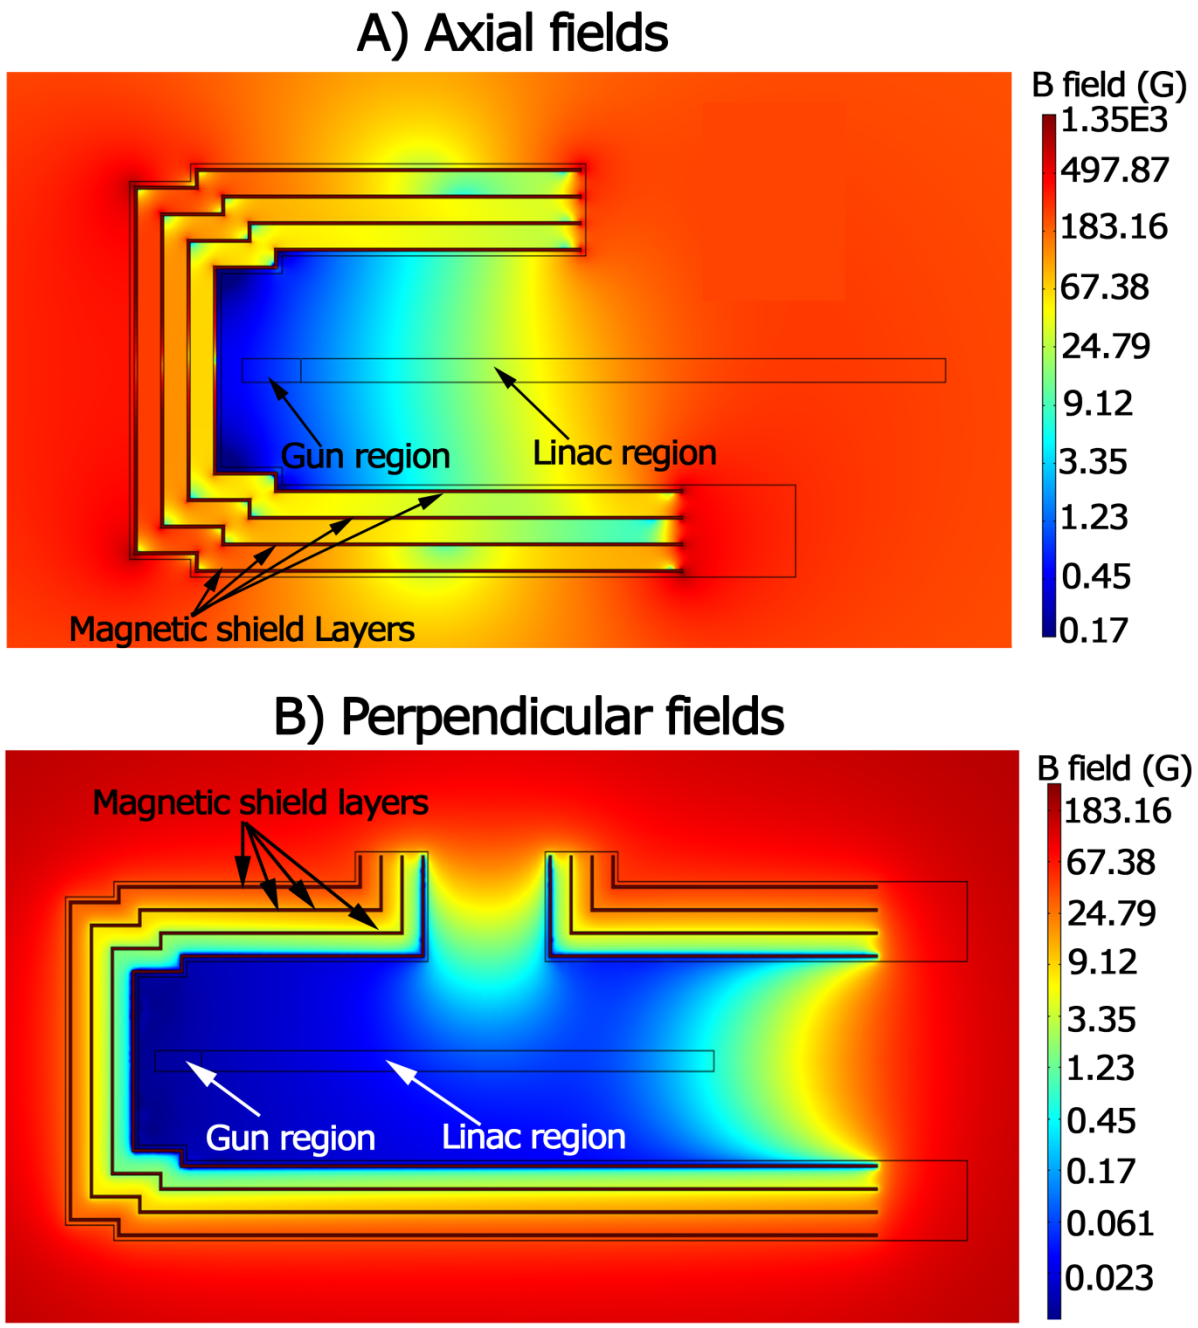

Figure 44: Examples of the magnetic field distribution in one of the shield designs. These shields had a total thickness 5 mm split across 4 layers, with the inner layer being made of mu-metal and the outer three carbon steel 1020. The shield length for axial fields is 200 mm, and the perpendicular case is 400 mm.

4.5. MRI magnet model

The MRI model with no magnetic shielding solved in around 5 minutes, and had homogeneity of 7ppm within a 30 cm DSV. Figure 12 shows an example field plot from a converged solution, whilst Figure 13 shows the fringe field in both the axial and perpendicular directions as a function of distance from the magnet.

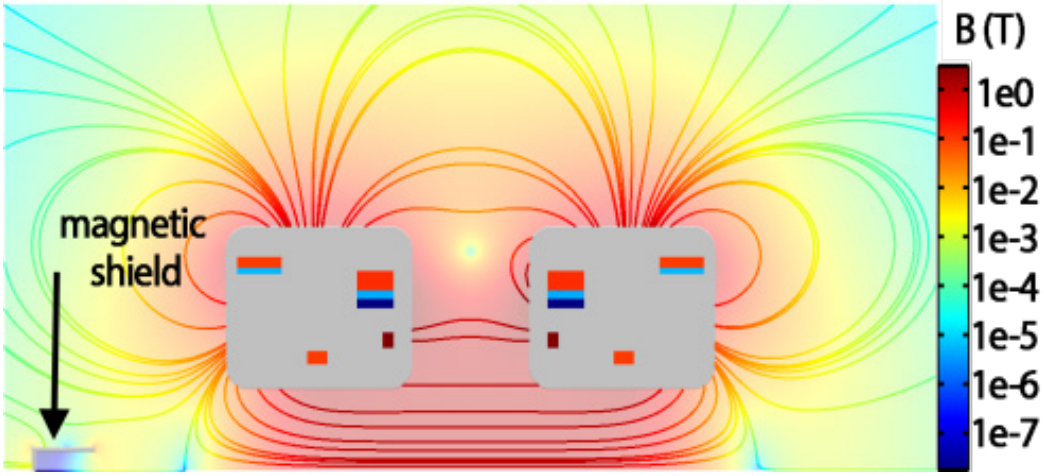

Figure 45: An example of a solved simulation of a magnetic shield. The fields and field lines are potted on a log scale as indicated by the color bar. In this instance, the shield is 400 mm long and 20 mm thick. In this case, the shield is in the 'in-line' configuration.

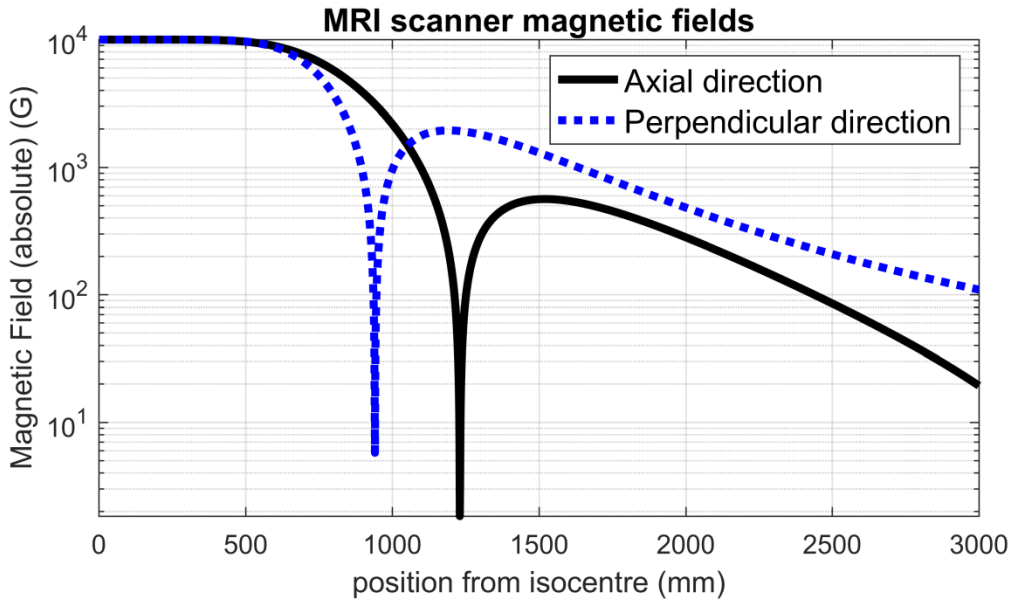

Figure 46: the absolute magnetic field of the MRI scanner in both the axial and the perpendicular direction. Note that the Y axis uses a log scale.

4.6. *Impact of magnetic shields on MRI magnet homogeneity*

Figure 14 shows the impact of incorporating a large (length=400 mm, thickness =20 mm) elliptical shield into the MRI model at various SIDs. The main consideration here is the DSV distortion (Figure 14A), as the goal of this simulation was not to design an optimal shield, but rather to get a sense of the distortion which could be induced. It can be seen that overall, the distortion falls well below the 300 ppm criteria. Only two data points exceed this, (they are not shown on the graph as the distort the range too much) For the axial orientation at SIDs of 900 mm and 1000 mm, the DSV distortion was 1090 and 589 respectively. Also, in general the distortion is less in the perpendicular direction than the axial direction (the exceptions to this statement are around the zero field position in the axial direction at 1243 mm (Figure 13). This is quite an interesting result, as the perpendicular field is substantially higher than the axial field (Figure 13), and as such one might intuitively expect the shield would become more magnetised and hence induce greater distortion. The following simulations all utilised an SID of 1500 mm.

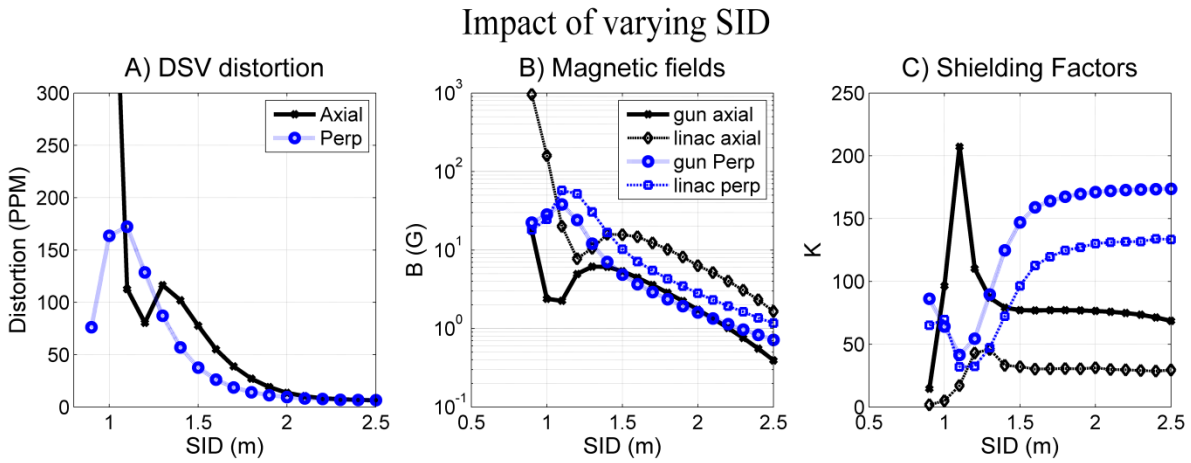

Figure 47: The impact of varing the SID on A) The distortion in the DSV B) The magnetic fields within the gun and linac volumes and C) on the shielding factors of these volumes. The shield simulated here was a single layer shield of thickness 20 mm, length 400 mm and with the port height shielding set to 100 mm from the centre of the shield.

Figure 15 shows the impact of varying the shield length. Again, it can be seen that for the same shield geometry the axial configuration tends to induce more distortion than the perpendicular orientation. Also, there is a peak in the axial shielding factor of the gun as a function of length, similarly to what was observed in the preceding sections. Based on this data, the shield length of the perpendicular shield was set to 400 mm, and the axial shield to 200 mm. For the axial case, we again assume that shielding of the linac is not an important consideration. (Also, slightly better shielding could be obtained for a shield length of 250 mm, but by keeping it at 200 mm we better avoid mechanical interference with the RF port).

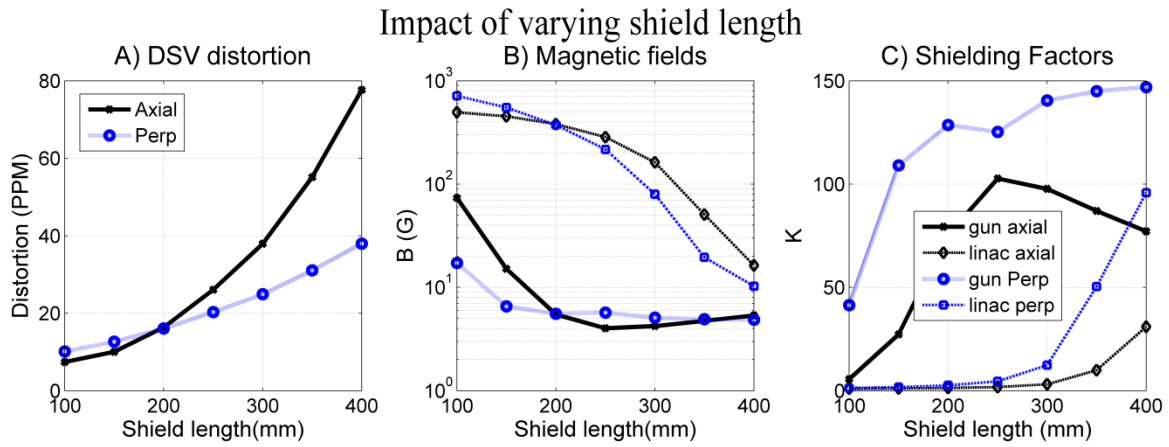

Figure 48: The impact of varying the shield length on A) The distortion in the DSV B) The magnetic fields within the gun and linac volumes and C) on the shielding factors of these volumes. The SID for this simulation was 1500 mm, and the shield thickness was 20 mm.

Figure 16 shows the impact of varying the shield thickness. Again, it is important to stress that the axial shield had a length of 200 mm at this point compared to 400 mm for the perpendicular case. This explains the fact that the perpendicular shields in this case induce more distortion in the DSV. It can be seen that increasing the thickness results in a steady increase in both distortion and shielding efficacy. However, at some point it becomes inefficient to continue to increase the shielding thickness as the increase in shielding efficacy begins to slow. Based on this data, the thickness of the axial shield was set to 10 mm, and the perpendicular fields 20 mm.

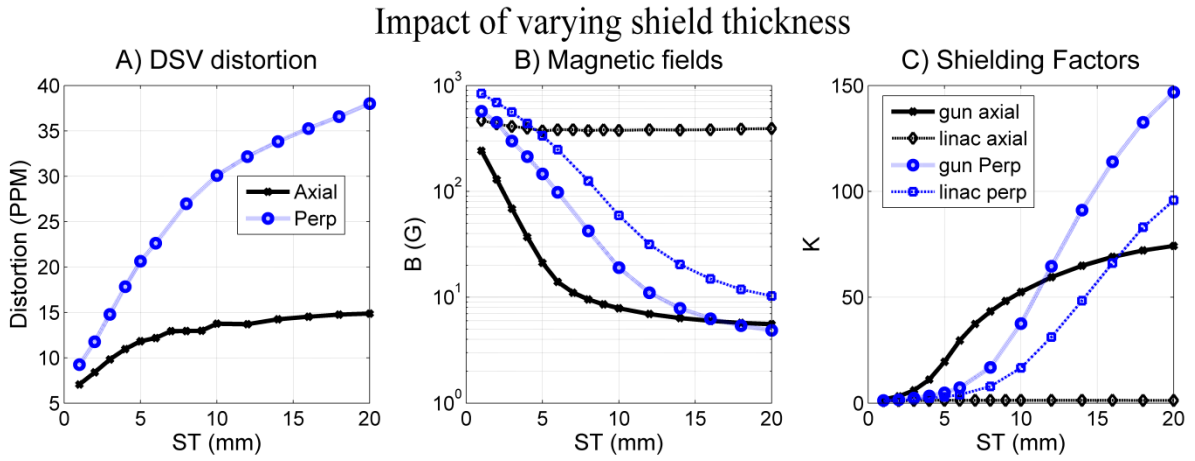

Figure 49: The impact of varying the shield thickness on A) The distortion in the DSV B) The magnetic fields within the gun and linac volumes and C) on the shielding factors of these volumes. The SID for this simulation was 1.5 m. A different shield length was used for the two different MRI-Lianc orientations on the basis of Figure 15; the axial set up had a shield length of 200 mm and the perpendicular congiuration had a shild length of 400 mm. This is the reason that the perpindicular shield induced higher distortion than the axial set up.

Figure 17 shows the impact of splitting the shield into different layers. For the axial fields, small increases in both the shielding factor and induced distortion were observed. For perpendicular fields, a large increase in the shielding factors is observed – for instance the gun shielding factor goes from 147 to 1428, with an associated magnetic field in the gun of 4.7 G to 0.5 G, and causes only a small increase in distortion (38 ppm to 58 ppm). On the

basis of this data, the axial shield was split into three layers, and the perpendicular shield into five layers.

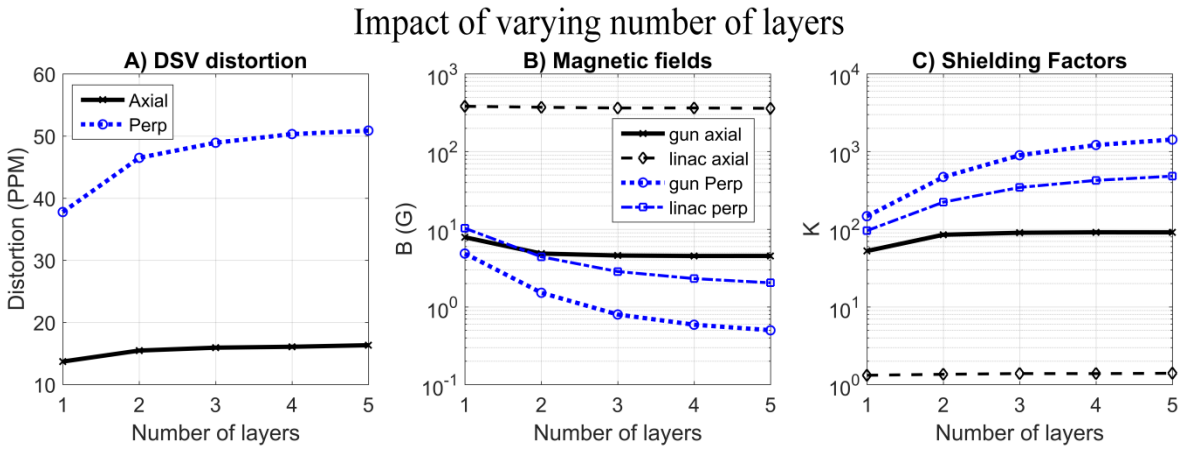

Figure 50: The impact of varying the number of layers on A) The distortion in the DSV B) The magnetic fields within the gun and linac volumes and C) on the shielding factors of these volumes. The SID for this simulation was 1.5 m. A different shield length and total thickness was used for the two different MRI-Lianc orientations on the basis of the data in Figure 15 & Figure 16; the axial set up had a shield length of 200 mm and total thickness of 10 mm and the perpendicular congiuration had a shild length of 400 mm and total thickness of 20 mm.

Figure 18 shows the impact of increasing the spacing between each layer. In constant fields, this yielded quite large increases in the shielding factors (Figure 10) however this effect is not reproduced to the same extent here. When the spacing is increased from 3 mm to 6 mm, a moderate increase in the shielding factors is observed (for perpendicular fields, the gun shielding factor increases from 1279 to 1494, for axial fields from 88.8 to 96.21). Beyond this, the shielding factors begin to drop. Again, the impact on DSV distortion is quite small. On this basis, the spacing was set to 6 mm for both cases. The shielding factors, magnetic fields in the volumes of interest, and induced distortion in the MRI scanner for the final shield geometry are outlined in Table 1.

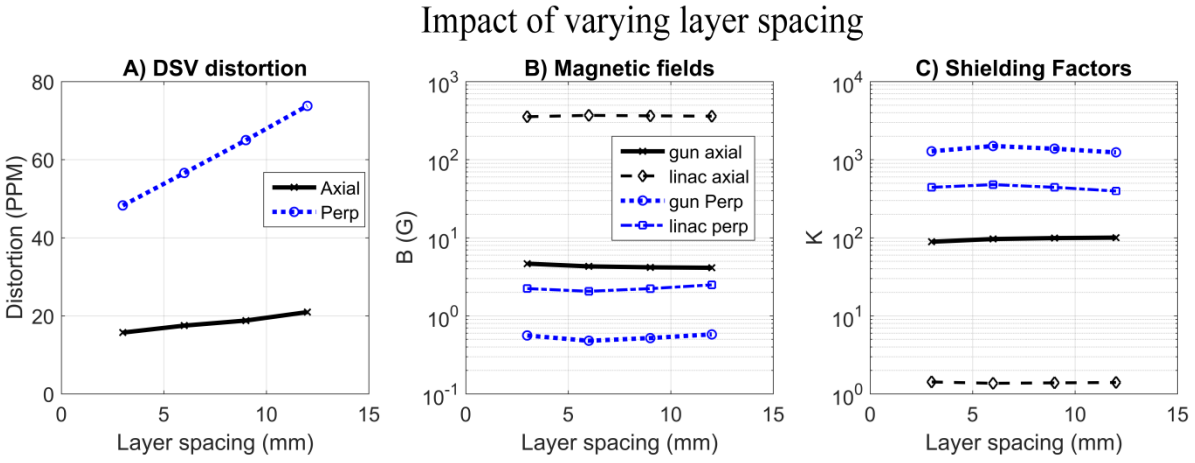

Figure 51: The impact of varying the layer spacing on A) The distortion in the DSV B) The magnetic fields within the gun and linac volumes and C) on the shielding factors of these volumes. The SID for this simulation was 1.5 m. A different shield length and total thickness was used for the two different MRI-Lianc orientations on the basis of the data in Figure 15, Figure 16 & Figure 17; the axial set up had a shield length of 200 mm, total thickness of 10 mm and three layers. The perpendicular congiuration had a shild length of 400 mm, total thickness of 20 mm, and 5 layers.

Table 7: Parameterisation, impact, and efficacy of the shield designs which were arrived for shielding at an isocentre of 1.5 m in the 1.0 T superconducting MRI magnet. Note that the thickness of each layer is the total thickness divided by the number of layers. The shape of each shield is very similar to those shown in Figure 11, although the parameterisation is different.

Parameter	In-line shield	Perpendicular shield
SID	1500 mm	1500 mm
Total shield thickness	10 mm	20 mm
Shield length	200 mm	400 mm
Number of layers	3	5
Port height shielding	0	100 mm
Gun shielding factor	96.2	1494
Average field in gun	4.3 G	0.5 G
Linac shielding factor	1.4	478
Average field in linac	354.8 G	2.1 G
Distortion induced in DSV	18 ppm	57 ppm

5. Discussion

In this work, passive magnetic shielding for MRI-Linac systems was explored. This was undertaken in two phases. In the first phase, basic principles of magnetic shielding were explored based references available in the literature and simulations within a constant magnetic field environment. In the second phase, these principles were utilised to design passive magnetic shields in both the perpendicular and in-line MRI-linac orientations of a superconducting 1.0 Tesla MRI magnet.

The advantage of simulating shields in constant magnetic fields (Section 4.1 to 4.4) is advantageous as the models can be more quickly solved; also the results tend to be cleaner and simpler to interpret. Using this approach, we demonstrated the impact that different shielding geometry can have on shielding efficacy (Figure 6). It was shown that an elliptical shield shape was superior to the previously presented cylindrical shield shape [8] across a wide range of scenarios, despite the fact that it actually contains less ferromagnetic material than a corresponding cylindrical shield (e.g. for length 400 mm and thickness 20 mm, the elliptical shield volume is 17% less than the cylindrical shield) and as such is expected to induce less distortion into the magnet. We also showed that for the in-line case, there is a peak in the shielding efficacy as a function of length. A similar effect was observed in [8], which was ascribed to saturation effects. Although saturation may play a role, the fact that this effect was also observed in linear shield models (Figure 5) shows that this is not the main cause of this effect; instead this effect is a result of the demagnetising fields from the shields. In fact, it makes sense that there must be peak at some length when one considers the fact that for both zero length infinite cylinders in axial fields, the shielding factor must be zero [15].

Also in this section, we observed that the same shield shape tends to have much better shielding efficacy in perpendicular fields than axial fields (Figure 5 and Figure 6). We also investigated the effect of the two shield openings (one for the RF port, and one at the end of the linac) in the constant field environment and found that the impact tended to be quite small (Figure 7 & Figure 8). For this reason, the shields were left open ended throughout the remainder of the study as this also facilitates ease of installation and access. For axial fields, shielding around the RF port was actually a disadvantage. Again the mechanism for this counterintuitive result is that the demagnetising fields are decreased by shielding the port.

In phase two of the study, we aimed to apply the principles learnt in the first phase to develop passive magnetic shields for a 1.0 Tesla MRI magnet (Section 3. 6). A number of comments can be made about the findings of these simulations. Firstly, the distortion induced in the DSV of the magnet was in general quite low compared to the 300 ppm limit that is conventionally adopted. An interesting finding of this work is that the distortion tended to be lower in the perpendicular configuration when the same shield geometry was solved (Figure 14 and Figure 15). Given that the fringe field of the magnet is actually substantially higher in the perpendicular direction (Figure 13), this leads us to hypothesise that this is a result which could be extended to other magnets as well. Coupled with the fact that shielding factors tend to be intrinsically higher for perpendicular fields (Figure 5 & Figure 6), this goes some way to compensating for the fact a much larger volume needs to be shielded and substantially increased sensitivity in the perpendicular orientation (Section 3. 2).

We aimed to implement passive shields at an isocentre distance of 1.5 m. The fringe fields were still quite high here, so this was considered a challenging case (Section 4.6). Nevertheless, we were able to reduce the fields to well below those which we initially aimed for (section 3. 2). In the axial direction, the field in the gun was reduced from 413 G to 4.3 G at a cost of 18 PPM induced distortion. In the perpendicular direction, the fields in the gun have been reduced from 715 G to 0.5 G, and in the Linac from 985 G to 2.1 G, at a cost of 57 PPM distortion. In both cases, the induced distortion is well below the 300 PPM limit. (It is worth noting however that the magnetic shields are far from the only parameters which influence linac homogeneity). In terms of designing a magnetic shield in magnetic fields, we suggest the approach we took her (sequentially optimise length, thickness, and layering) is the best way to approach the first run of a magnetic shield. The basic shield which results from this process could then be further optimised if necessary (this is discussed below).

A few limitations and suggestions for future work should be mentioned. Firstly, we have not addressed in this study is the practical fabrication and implementation of magnetic shields. Fabrication is important is the processing of magnetic steel can affect the magnetic properties. Also, real magnetic materials exhibit hysteric effects, which are not included in this study; nor are the effect of magnetic domains. In addition, the exact BH curve of the material from which the shield is constructed may not be precisely known. For all these reasons, there is expected to be some variation between simulated results and experimental results; we hope to quantify this in future work. Magnetic shielding material is normally sold in thin layers; therefore in order to fabricate a shield of s given thickness, the simplest approach is probably to 3D print a cast of the desired shape and wrap an appropriate number

of layers around it. Closely spaced layers of thin material should behave quite similarly to a single thicker layer. Shield layering could then be achieved by adding spacers of any nonmagnetic material (foam, rubber, cloth, etc.). A more finessed approach would be to have the desired shield formed from scratch; of course this would also involve additional expense.

Secondly, throughout this study, we assumed that the linac did not have to be shielded in the in-line direction. A number of studies have demonstrated that beam current loss is minimal in the linac for axial fields .On the other hands, magnetic lensing effects can occur which cause the spot size to either focus or defocus [4, 6]. Over focusing of the electron beam could lead to target life time issues. For this reason, whilst it is possible to operate the linac unshielded, the magnetic field environment it is operated in should still be chosen with care.

Because a non-linear solver is used in this study, the starting point of the solver can result in slightly different final solutions. This did not cause any measurable difference in the shielding factors, but we did observe that the PPM values could vary slightly (on the order of 1 PPM) when different runs of the same geometry were solved. This effect can be seen can be seen for instance in the axial series of Figure 18, which was ruin over a few different times instead of one smooth parametric sweep. The PPM is by definition a very sensitive quantity (parts per million) and this is why a small effect can be observed. This could be reduced by a reduction in the stopping criterion of the non-linear solver; however this would substantially increase the solving time and the small variations seen in the ppm values are not considered important.

Finally, although we have demonstrated the appropriate principles to be taken into account in shield design, we have not designed or attempted to design truly optimal shields. This could be achieved in a proper optimisation framework where the distortion induced in the DSV was minimised while requiring a certain maximum or average field within the volume of interest. In particular, the thickness and spacing of successive layers could be optimised further than we have done here.

In this work, the fundamental principles of passive magnetic shielding have been explored in an FEM framework, and then applied to the design of passive shielding of linear accelerators for MRI-Linac systems. Passive magnetic shields were designed for an Isocentre of 1.5 meters in both the axial and in-line directions which reduced the magnetic fields in the linac and gun region to levels such that the linac should perform acceptably. The distortion induced in the MRI magnet was 18 ppm for the in-line orientation and 57 PPM for the perpendicular orientation, both falling well below the upper limit of 300 ppm.

References

1. Ménard, C. and U. van der Heide. *Introduction: Systems for magnetic resonance image guided radiation therapy*. in *Seminars in radiation oncology*. 2014. Elsevier.
2. Constantin, D.E., R. Fahrig, and P.J. Keall, *A study of the effect of in-line and perpendicular magnetic fields on beam characteristics of electron guns in medical linear accelerators.* Medical physics, 2011. **38**(7): p. 4174-4185.
3. St Aubin, J., S. Steciw, and B. Fallone, *Effect of transverse magnetic fields on a simulated in-line 6 MV linac.* Physics in medicine and biology, 2010. **55**(16): p. 4861.
4. St. Aubin, J., D. Santos, S. Steciw, and B. Fallone, *Effect of longitudinal magnetic fields on a simulated in-line 6 MV linac.* Medical physics, 2010. **37**(9): p. 4916-4923.
5. St. Aubin, J., S. Steciw, and B. Fallone, *The design of a simulated in-line side-coupled 6 MV linear accelerator waveguide.* Medical physics, 2010. **37**(2): p. 466-476.
6. Whelan, B., S. Gierman, L. Holloway, J. Schmerge, P. Keall, and R. Fahrig, *A novel electron accelerator for MRI-Linac radiotherapy.* Medical physics, 2016. **43**(3): p. 1285-1294.
7. Yun, J., J.S. Aubin, S. Rathee, and B. Fallone, *Brushed permanent magnet DC MLC motor operation in an external magnetic field.* Medical physics, 2010. **37**(5): p. 2131-2134.
8. Santos, D., J.S. Aubin, B. Fallone, and S. Steciw, *Magnetic shielding investigation for a 6 MV in-line linac within the parallel configuration of a linac-MR system.* Medical physics, 2012. **39**(2): p. 788-797.
9. St. Aubin, J., S. Steciw, and B. Fallone, *Magnetic decoupling of the linac in a low field biplanar linac-MR system.* Medical physics, 2010. **37**(9): p. 4755-4761.
10. Lagendijk, J.J., B.W. Raaymakers, A.J. Raaijmakers, J. Overweg, K.J. Brown, E.M. Kerkhof, R.W. van der Put, B. Hårdemark, M. van Vulpen, and U.A. van der Heide, *MRI/linac integration.* Radiotherapy and Oncology, 2008. **86**(1): p. 25-29.
11. Wadey, W., *Magnetic shielding with multiple cylindrical shells.* Review of Scientific Instruments, 1956. **27**(11): p. 910-916.
12. Fahrig, R., K. Butts, J.A. Rowlands, R. Saunders, J. Stanton, G.M. Stevens, B.L. Daniel, Z. Wen, D.L. Ergun, and N.J. Pelc, *A truly hybrid interventional MR/X-ray system: Feasibility demonstration.* Journal of Magnetic Resonance Imaging, 2001. **13**(2): p. 294-300.
13. Kok, J., B. Raaymakers, J. Lagendijk, J. Overweg, C. De Graaff, and K. Brown, *Installation of the 1.5 T MRI accelerator next to clinical accelerators: impact of the fringe field.* Physics in medicine and biology, 2009. **54**(18): p. N409.
14. Paperno, E. and I. Sasada, *Magnetic Circuit Approach to Magnetic Shielding.* 日本応用磁気学会誌, 2000. **24**(1): p. 40-44.
15. Reutov, Y.Y., *Physical interpretation of magnetostatic shielding.* Russian journal of nondestructive testing, 2000. **36**(2): p. 117-126.
16. Chung, Y.-C., B.-C. Woo, and D.-C. Park, *Shielding factors of grain oriented 3% SiFe cylinders.* Magnetics, IEEE Transactions on, 1999. **35**(5): p. 3454-3456.
17. Mager, A., *Magnetic shielding efficiencies of cylindrical shells with axis parallel to the field.* Journal of Applied Physics, 1968. **39**(3): p. 1914-1914.
18. Rücker, A., *VII. On the magnetic shielding of concentric spherical shells.* The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science, 1894. **37**(224): p. 95-130.
19. Schweizer, F., *Magnetic shielding factors of a system of concentric spherical shells.* Journal of Applied Physics, 1962. **33**(3): p. 1001-1003.
20. Sumner, T., J. Pendlebury, and K. Smith, *Convectional magnetic shielding.* Journal of Physics D: Applied Physics, 1987. **20**(9): p. 1095.
21. Paperno, E., H. Koide, and I. Sasada, *A new estimation of the axial shielding factors for multishell cylindrical shields.* Journal of Applied Physics, 2000. **87**(9): p. 5959.
22. Griffiths, D.J. and R. College, *Introduction to electrodynamics*. Vol. 3. 1999: prentice Hall Upper Saddle River, NJ.
23. Brown, R.W., Y.-C.N. Cheng, E.M. Haacke, M.R. Thompson, and R. Venkatesan, *Magnetic resonance imaging: physical principles and sequence design*. 2014: John Wiley & Sons.

24. Kolling, S., B. Oborn, and P. Keall, *Impact of the MLC on the MRI field distortion of a prototype MRI-linac.* Medical physics, 2013. **40**(12): p. 121705.
25. Burt, E. and C. Ekstrom, *Optimal three-layer cylindrical magnetic shield sets for scientific applications.* Review of scientific instruments, 2002. **73**(7): p. 2699-2704.
26. Mager, A., *Magnetic shields.* IEEE Transactions on Magnetics, 1970. **6**(1): p. 67-75.
27. Reutov, Y.Y., *Choice of the number of shells for a spherical magnetostatic shield.* Russian journal of nondestructive testing, 2001. **37**(12): p. 872-878.
28. Jackson, J.D., *Classical electrodynamics*. 1999: Wiley.

Chapter 5
A novel electron accelerator for MRI-Linac radiotherapy

A Novel electron accelerator for MRI-Linac radiotherapy

Brendan Whelan[1,2,3] Stephen Gierman[4], Lois Holloway[2,3], John Schmerge[4], Paul Keall[1,2] Rebecca Fahrig[5]

8. Radiation Physics Laboratory, University of Sydney, Sydney (NSW), 2006, Australia

9. Ingham Institute for Applied Medical Research, Liverpool (NSW), 2170, Australia

10. Liverpool Cancer Therapy Centre, Liverpool Hospital, Liverpool (NSW), 2170, Australia

11. SLAC National Laboratory, Menlo Park (CA), 94025, USA

12. Stanford University, Palo Alto (CA), 94305, USA

Abstract

Purpose

MRI guided radiotherapy is a rapidly growing field; however current electron accelerators are not designed to operate in the magnetic fringe fields of MRI scanners. As such, current MRI-Linac systems require magnetic shielding, which can degrade MR image quality and limit system flexibility. The purpose of this work was to develop and test a novel medical electron accelerator concept which is inherently robust to operation within magnetic fields for in-line MRI-Linac systems.

Methods

Computational simulations were utilised to model the accelerator, including the thermionic emission process, the electromagnetic fields within the accelerating structure, and resulting particle trajectories through these fields. The spatial and energy characteristics of the electron beam were quantified at the accelerator target and compared to published data for conventional accelerators. The model was then coupled to the fields from a simulated 1 T superconducting magnet and solved for cathode to isocenter distances between 1.0 and 2.4 meters; the impact on the electron beam was quantified.

Results

For the zero field solution, the average current at the target was 146.3 mA, with a median energy of 5.8 MeV (interquartile spread of 0.1MeV), and a spot size diameter of 1.5 mm Full-Width-Tenth-Maximum (FWTM). Such an electron beam is suitable for therapy, comparing favourably to published data for conventional systems. The simulated accelerator showed increased robustness to operation in in-line magnetic fields, with a maximum current loss of 3% compared to 85% for a conventional system in the same magnetic fields.

Conclusions

Computational simulations suggest that replacing conventional DC electron sources with an RF based source could be used to develop medical electron accelerators which are robust to operation in in-line magnetic fields. This would enable the development of MRI-Linac systems with no magnetic shielding around the linac and reduce the requirements for optimization of magnetic fringe field, simplify design of the high-field magnet, and increase system flexibility.

1. Introduction

Several research groups are developing coupled medical linear accelerator and Magnet Resonance Imaging devices (MRI-Linac). The goal of these efforts is to enable in-room MRI for anatomic and physiological treatment adaptation and response monitoring [1]. Integration of these two devices is challenging, as both produce external electromagnetic fields, and the combined device must function within the net electromagnetic field environment. The resultant electromagnetic coupling of the two devices results in many engineering and design challenges, one of which is the production of an acceptable treatment beam - the subject of this paper.

Typical treatment beam requirements for photon radiation therapy are beam energy of 4-20 MV and dose rate on the order of 500 cGy per minute at the treatment isocenter [2]. A simple medical linear accelerator (linac) can be considered to comprise of two subcomponents: a thermionic electrostatic electron gun which serves as the source of the beam, and a series of coupled resonant radiofrequency (RF) cavities which are used to accelerate the electron beam to MeV energies. The accelerated electron beam is then typically collided with a tungsten target to produce a bremsstrahlung photon beam. It is challenging to operate a linac in MRI-Linac systems, as moving electrons are subjected to Lorentz forces from the magnetic fields of the MRI scanner. If not compensated for, this can cause severe aberrations in the linac behaviour, up to and including complete beam loss [3, 4]. The exact behaviour of a linear accelerator when subjected to external magnetic fields depends on the magnitude and orientation of those fields. As such, the orientation of the accelerator with respect to the MRI scanner becomes important. Two orientations are feasible; the in-line setup, in which electrons are accelerated in the same direction as the magnetic field of the MRI-scanner, and the perpendicular setup, in which the electrons are accelerated perpendicular to the magnetic field. Each of these configurations has unique advantages and disadvantages associated with it which have been discussed elsewhere [5] – however, if one considers particle acceleration in isolation, then the in-line configuration is indisputably the superior option. This is because magnetic force on a charged particle is minimised when the particle is travelling in the same direction as the magnetic field lines (to be precise, the magnitude of magnetic force is zero when a particle travels parallel to a magnetic field, and maximal when it travels perpendicular).

The effects of both in-line and perpendicular magnetic fields on linear accelerator operation have previously been studied via computational simulations. For the perpendicular case, total beam loss occurred at 14 G, and 45 % at 6 G [3, 4]. This means that in order to produce a treatment beam for the perpendicular orientation, the linac must be operated in a near zero field environment. This can be achieved by modifying the magnet and magnetically shielding the linac [6] – however, MRI magnet design (and redesign) is not a trivial task, and magnetic shielding perturbs the field homogeneity in the imaging volume of the MRI scanner. For the inline case, the maximum beam loss was 79 %, which occurred at a field of 600 G [3, 7]. It was also shown that the effect of magnetic fields on the electron accelerator is nearly entirely on the electron gun – that is, operation of the accelerating cavities is largely unaffected [8, 9].

As such, to produce a treatment beam in the in-line orientation, the only component which

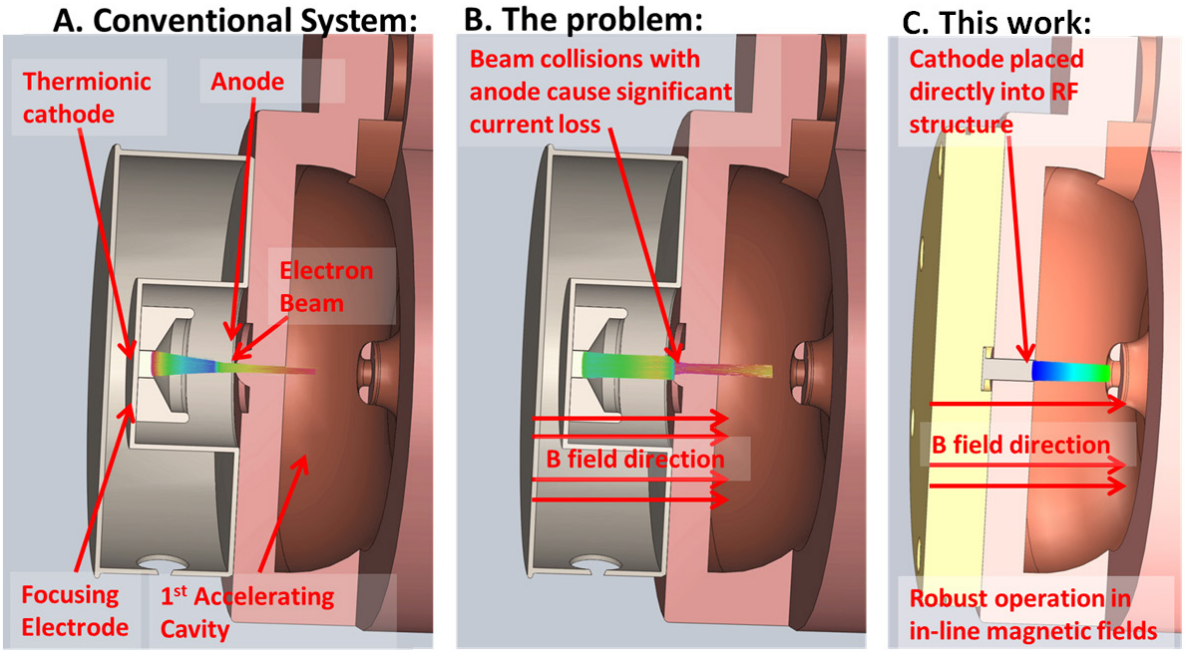

Figure 52: A) conventional medical electron accelerator utilizing a steady state stream of electrons, however B) performance is compromised in in-line MRI-Linac systems. C) In this work we are proposing a novel electron accelerator utilizing an RF electron source which is robust to operation in magnetic fields.

needs to be customised is the electron gun. Two solutions for operating the electron gun in in-line magnetic fields have been proposed. The first is to redesign the optics of the electron gun taking the presence of in-line magnetic fields into account such that the modified gun functions within these fields [9]. The second is to place magnetic shielding around the gun such that the field is reduced enough that acceptable gun performance is obtained [10]. Both these approaches were shown to be very effective, however both have drawbacks. Redesign of the gun optics requires a bespoke gun design for each different field it is to be used in. The published solution operates optimally only in the relatively high field of 1000 G or higher, and it is not clear if acceptable solutions of this nature exist at lower field strengths. The alternative, ferromagnetic shielding, causes distortion in the imaging field of the MRI scanner. This distortion can be corrected up to a point, as evidenced by current first generation MRI-Linac systems which successfully utilize magnetic shielding for either the electron gun or the entire linac [1]. However for higher field strengths and smaller SIDs, this becomes more difficult. It also limits the flexibility one has with which to compensate for other components which can cause magnetic distortion or require shielding, such as multi leaf collimators [11, 12].

The ideal accelerator for in-line MRI-Linac systems would be robust to operation in a range of field strengths without magnetic shielding. A solution which could meet these criteria and that has not previously been explored is a radiofrequency (RF) electron gun based system. As the name implies, instead of the steady state fields used to produce an electron beam in conventional systems, RF electron guns utilise RF fields. RF guns are widely used in other

particle accelerator fields, for instance as injectors to synchrotron beams [13]. An RF gun based accelerator has two theoretical advantages over a DC based gun setup for MRI-Linac systems. (1) The beam is accelerated to relativistic velocities over a much shorter distance, resulting in a 'stiffer' electron beam which is more difficult to bend, and (2) the cathode is subjected to much higher electric fields, which could reduce the need for transverse beam focusing – the main problem with conventional electron gun operation in in-line MRI-Linac systems [9]. The purpose of this paper was to investigate these hypothetical advantages, and determine if an RF-gun based accelerator could produce a beam suitable for radiotherapy treatments.

2. Methods and Materials

Computational simulations were utilised to investigate the behaviour of an RF gun based accelerator in in-line magnetic fields.

2.1. Model of thermionic emission

In conventional medical DC electron guns, a relatively low (kilovoltage) electrostatic field is applied to a thermionic cathode, resulting in space charge limited thermionic emission (the space charge of the beam limits the emitted current).

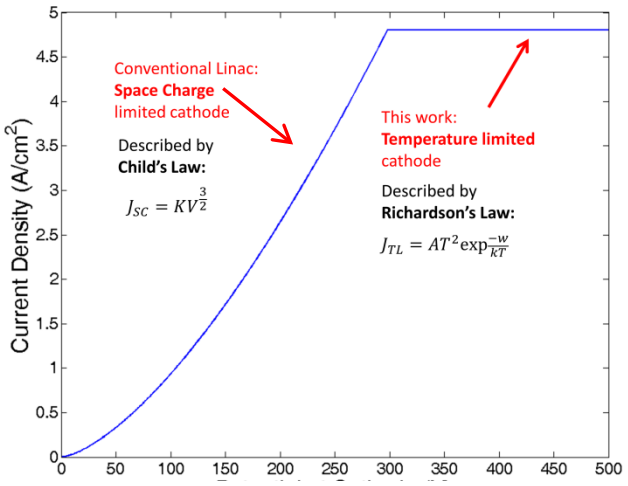

Figure 53 Theoretical current extracted from a thermionic cathode versus potential difference. Note that the potential here includes the field of the extracted beam itself, and is assessed at a distance 0.5 mm from the cathode

Table 8: Parameter values used in Richardson's equation (Equation 1)

Parameter	Value
A	60 A/(cm^2.°K)
T	1245 K
w	1.8 eV

$$J_{TL} = AT^2 \exp \frac{-w}{k_B T}$$

Equation 19: Richardson's law for temperature limited thermionic emission. A is Richardson's constant, T is temperature, w is work function, and k_B is Boltzmann's constant.

$$J_{SC} = KV^{3/2}$$

Equation 20: Child's law for space charge limited emission. K is a geometry dependent constant often termed the perveance, and V is the electric potential assessed close to the cathode (a typical sampling distance would be 0.5 mm)

$$\varepsilon_N = \frac{r_c}{2}\sqrt{\frac{k_b T}{m_0 c}}$$

Equation 21: Thermal emittance on the surface of a thermionic cathode (in m rad), where r_c is the cathode radius, k_b is Boltzmann's constant, T is the cathode temperature, m_0 is the electron rest mass, and c is the speed of light.

In the proposed configuration, the megavoltage electric field at the cathode is constantly changing as the electric field oscillates back and forth. Half the time no current will be emitted at all as the electric field is pointing in the wrong direction. For the other half, current emission is modelled using Richardson's law for temperature limited thermionic emission (Equation 1), which describes the current extracted from a thermionic cathode assuming no space charge effects. This emission model is appropriate as the electric field in the RF cavity is hundreds of times higher than that in a DC electron gun, and temperature limited emission will dominate over space charge limited emission. An estimate of the total space charge limited contribution to the current can be obtained as follows: an accelerating RF pulse is defined by an input value for the peak RF field at the cathode. The point at which the cathode becomes temperature limited is then calculated; i.e. where J_{SC} (Equation 2) > J_{TL} (Equation 1). The remainder of the pulse, which is considered space charge limited, is spilt into 10 bins, and in each bin the space charge limited current is calculated according to Equation 2. The electric field is converted to potential difference for Child's law using V=E.d, where d was chosen as 0.5 mm – a typical value for this kind of calculation. Using a value of 4 MV/m for peak field at the cathode we estimate space charge limited emission will account for less than 5% of the emitted current. As will be seen in the results section 3.2), 4 MV/m is a very conservative value. Higher fields further limit the impact of space charge. The values used in Richardson's equation are show in Table 1, and match commercially available cathodes and experimental observations [14-16].

The transverse RMS emittance of the beam is often used as a figure of merit to quantify electron beam quality [17]. For a thermionic cathode, the intrinsic, or thermal RMS emittance is given by Equation 3 and is generally considered a lower bound on what can be practically achieved [17, 18].

2.2. Radiofrequency field calculation

The next step was to calculate the RF fields to which the electrons will be subject while in the accelerating structure. We have utilised the basic S band cavity design presented by St. Aubin [19]. This basic structure was staggered to develop an RF structure with five full accelerating cavities and one half cavity (Figure 4). The fields within the cavities are solved using an

eigenmode solver (CST, Darmstadt, Germany). The solver is based on the finite element method and utilises a tetrahedral mesh with quadratic shape functions. An adaptive meshing strategy was utilised such that the discretisation error in the frequency of the returned solution was less than 0.5 MHz. In order to minimise computational cost, a lossless eigenmode solver was utilised, which assumes perfect conductivity at all boundaries – a reasonable approximation for copper. The losses in the waveguide were calculated as a post processing step, which uses perturbation theory based on the magnetic field distribution at the walls. A conductivity of 5.8×10^7 S/m (copper) was used. Further details on this method can be found in [20]. A small ring like structure was added around the cathode in order to increase the radial focusing fields at the point of emission.

2.3.Particle trajectories

In order to calculate the particle trajectories, the RF fields from section 2. 2) were imported into a Particle in Cell (PIC) solver (CST, Darmstadt, Germany) as ASCII files. The electric fields in the central region where the radial coordinate is less than 5 mm were sampled on a 0.1 mm Cartesian grid, whilst all other fields were sampled on a 1 mm grid. The resulting ASCII data took up around 6 Gb of disk space. The PIC solver is a fully integrated solution which incorporates space charge and wake field effects of the electron beam based on the finite integration technique (FIT), a formulation of the finite difference time domain (FDTD) method. Exactly the same geometry as in section 2. 2) is used. The waveguide structure was discretised into 4.97e6 hexahedral mesh cells. The electron source was defined based on the data from section 2. 1). Particle tracking was carried out over a time period of 1850 ps, representing around 6 RF cycles and 2 full electron bunches at the target plane. Rather than explicitly simulate each individual electron trajectory, electrons are grouped into macro particles. Each macro particle in this work represented around 200 electrons, and around 3 million macro particles reached the target in each simulation.

2.4. Beam Assessment at the Target

In order to assess the performance of the RF gun based accelerator, a beam monitor was placed at the exit of the simulation; particle information was scored as it crossed this monitor and exported to Matlab for further analysis. The mean current, spatial and energy distributions were evaluated to assess suitability of the beam for radiotherapy treatments. The normalised emittance at the target was calculated using Equation 4 [17]. Note that we calculate normalised emittance instead of the geometric emittance which has been used in previous recent publications in this journal [3, 9, 10, 19, 21, 22]. This is because the former quantity is a more appropriate and robust metric in instances where the beam energy is changing [17, 23]. If it is assumed that if the relativistic γ and β distributions are single valued, normalized emittance can be also described as $\varepsilon_G = {}^{\varepsilon_N}/_{\beta\gamma}$ [17], where ε_G is the geometric emittance which has been utilised in other recent publications [9, 19, 21].

$\varepsilon_N = \sqrt{\langle x^2\rangle\langle p_x^2\rangle - \langle x.p_x\rangle^2}$	**Equation 22: Normalized emittance of a particle beam, where x is the x coordinate, and p_x is the scaled momentum, defined using the special relativity parameters such that $p_x=\gamma\beta_x$.**

2.5. Back-bombardment power

An issue for all microwave accelerators which utilize a thermionic cathode is back bombardment. This refers to electrons which are accelerated back towards the cathode, where they deposit unwanted power. This has two effects – firstly, it can damage the cathode and reduce its lifetime. Secondly, whilst the RF is on (~ 5 us in medical systems), the temperature of the cathode can be expected to rise steadily at a rate dependent on the power being delivered by back accelerated electrons [2, 18, 24]. As such, back bombardment is one of the principle effects which limit the achievable pulse length in thermionic guns used in both accelerator physics and in therapeutic systems. Conventional DC medical electron guns can be expected to have two advantages compared to the RF type cathode described here when considering back bombardment. Firstly, because they are operated in a space charge limited mode of emission (as opposed to temperature limited) the impact of additional heating during the pulse should be smaller. Secondly, any electrons striking a DC cathode have to first navigate the anode drift tube and overcome the DC electric potential of the electron gun – meaning a cathode in a DC system has inherently greater protection from back bombardment.

In order to quantify the extent of back bombardment occurring in the proposed design, the back accelerated electrons striking the cathode plane in the simulation described in section 2.3) were exported to matlab for further analysis. In order to compare this to a conventional system, the RF source from section 2.1) was replaced with a DC source exported from an Opera electron gun simulation which has been previously described [9, 22]. The particles striking the wall of the first cavity of this simulation were imported into a separate PIC simulation which incorporated DC electron gun geometry along with the associated electrostatic field, and the electrons reaching the DC cathode were exported to matlab. As will be seen in section 3.4) it takes some time for back bombardment power to reach a steady state, so the simulation time was extended to 4000 ps for these simulations.

2.6. Performance in magnetic fields

In order to assess the performance of the RF gun accelerator in the presence of in-line magnetic fields, the particle in cell model from section 2. 3) was coupled to a previously published model of a 1 Tesla MRI magnet [11]. This magnet is being constructed for the Australian MRI-Linac program [5]. Since this field is axially symmetric in the in-line orientation, it can be characterised by the central axial field as outlined in [9]. This field expansion is accurate to within a few gauss within the 2.5 mm radius beam line of the present accelerator structure. CST contains a built-in interface for adding a magnetic field in this manner which was utilised. The simulation was repeated for cathode to isocenter distances from 1 to 2.3 meters in 0.1 meter steps; the magnetic field at the cathode ranged from 141 to 2186 G over this range. For each step the beam assessment was repeated. Note that the source to isocenter distance (SID) is approximately 300 mm smaller than the cathode to isocenter distance.

3. Results

3.1. Model of thermionic emission

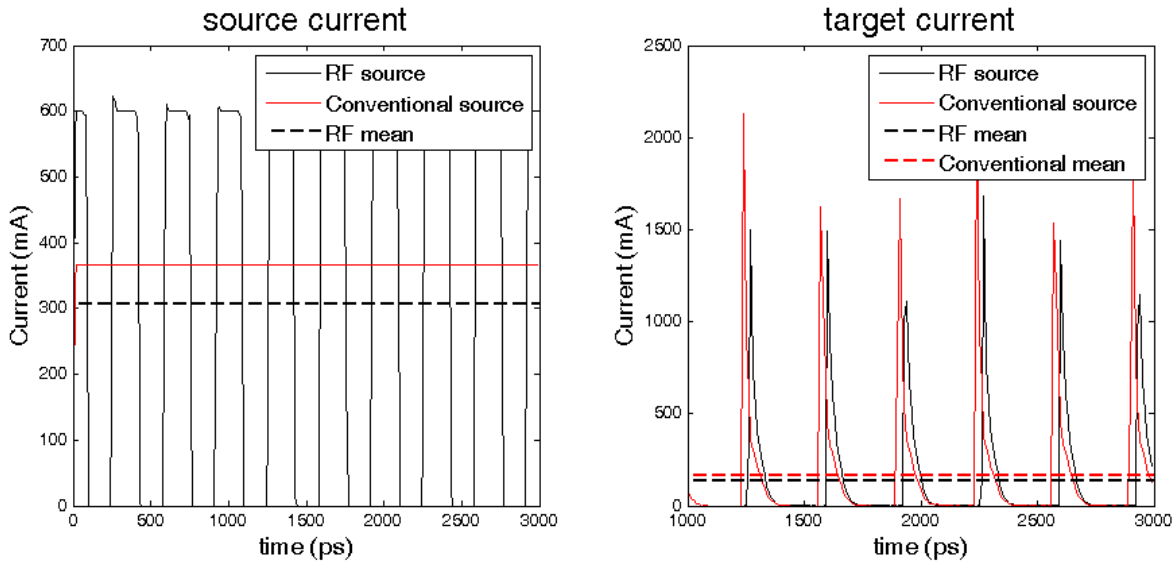

Figure 54: Comparison of the electron current from a conventional and RF electron source at the beginning (source) and end (target) of the accelerator. Note that current is not a particularly well defined concept when dealing with electron bunches over short time scales, and this is the cause of the variability in the target current peaks.

Based on Equation 1, a cathode of radius 2 mm and work function 1.8 eV operated at a temperature of 1245 K will emit a current of 600 mA in the temperature limited regime. These values are all well within the capabilities of modern tungsten dispenser cathodes [14]. Figure 3 shows a plot of the net current passing through a plane placed 1 mm in front of the start of the accelerating waveguide and at a plane at the end of the accelerating waveguide for the RF based electron source and a DC based electron source. The conventional DC source modelled here is the diode gun published by St. Aubin [22] and frequently utilised in publications in this area [9, 10, 21]. It can be seen that whilst the source currents of the two electron sources show considerable difference, the target currents for the RF based model and the DC based model are very similar. The mean target current of the RF based source is 146.3 mA; well within the current requirements outlined by Karzmark for a low energy medical linac and experimental values for similar systems [2, 22]. Also, if needed the current can be further increased by increasing either the size or temperature of the cathode. As such, we conclude that a temperature limited RF cathode can easily generate the target currents required for radiotherapy. The expected thermal emittance of the electron beam on the cathode surface is 0.46 mm mrad (Equation 3).

3.2. Radiofrequency field calculation

The electromagnetic field solution is shown in Figure 4 (Multimedia view). Table 2 shows the frequency, shunt impedance and quality factor for a single uncoupled accelerating cavity and for the final coupled structure. For comparison, previously published values for the same accelerating cavity are also included. Although further optimisation of the RF structure could be undertaken, we did not do this as our goal was to provide a proof of principle, and this was achieved with minimal modifications to the original geometry.

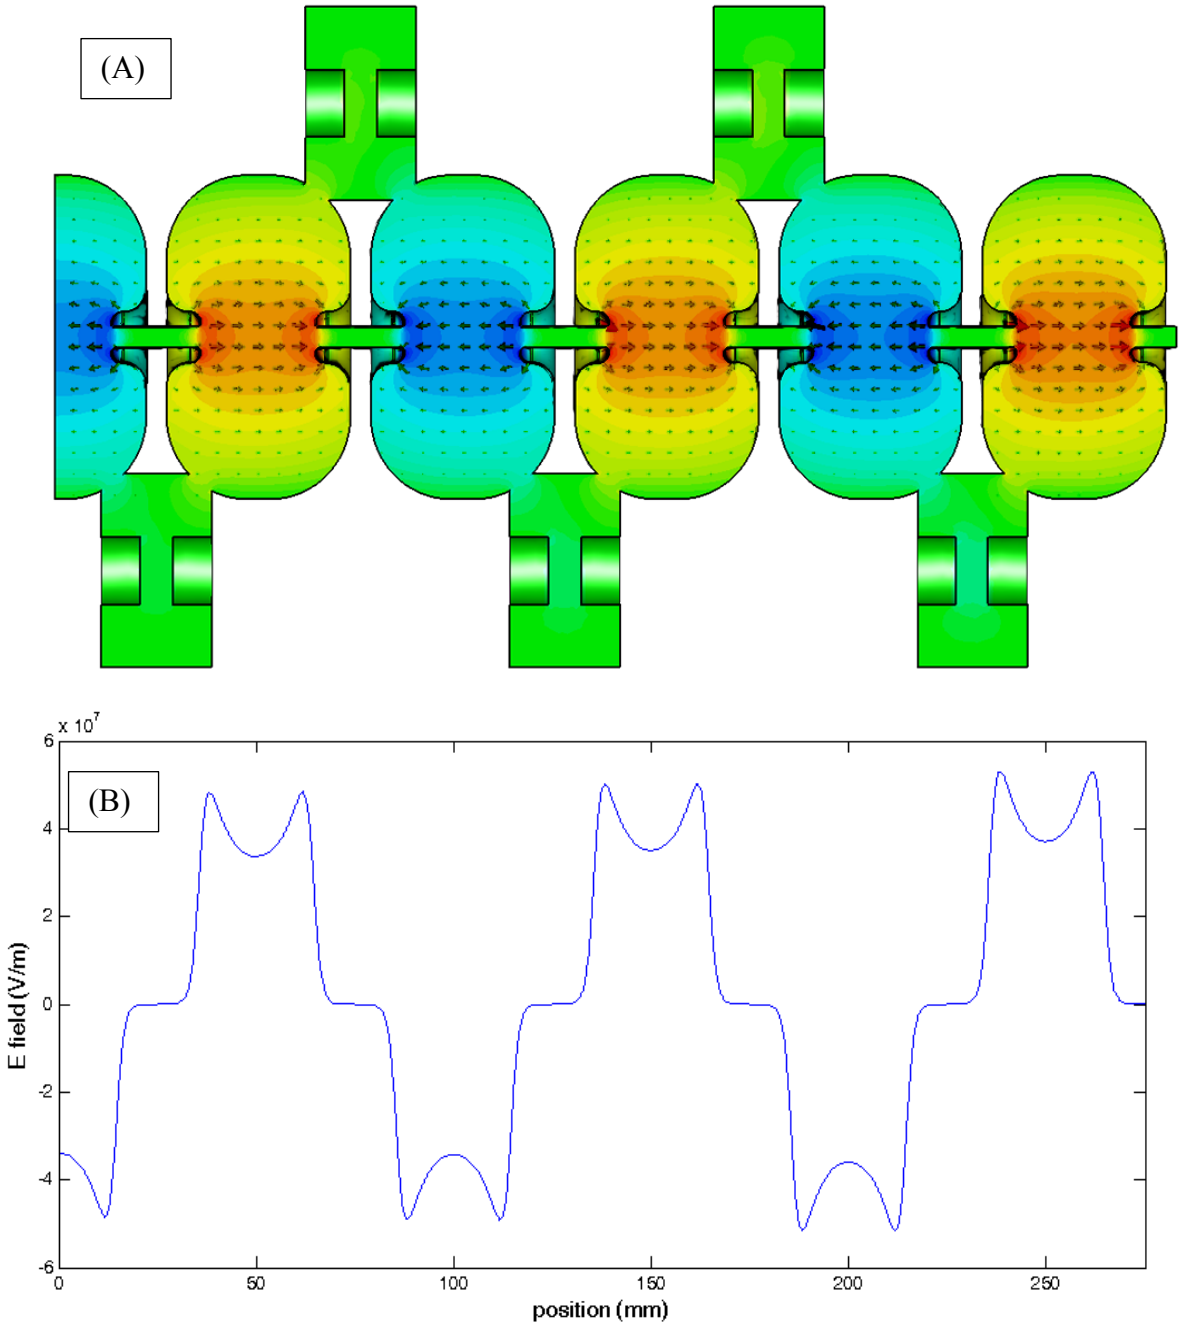

Figure 55: (A) The accelerator structure used and the electric field eigenmode solution (B) The axial electric field along the length of the accelerator (Multimedia view)

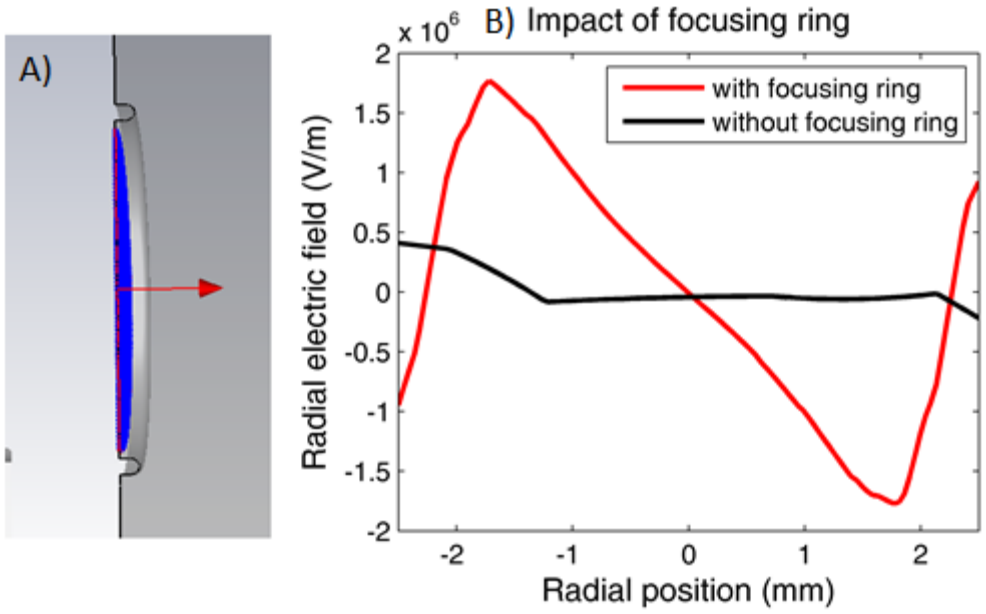

Figure 56: A) shows a cutaway view of the focusing ring around the cathode, shown in blue. B) shows the impact of this structure of the radial fields, plotted 1 mm in front of the cathode. Note that the asymmetry evident in the plot without the focusing ring is due to the presence of the coupling cavity.

As such, there are only two changes to the initial geometry published by St. Aubin. Firstly in the last cavity, the nose cone length was increased by 0.08 mm to compensate for the fact that there is only one coupling slot present.

Secondly, we added a focusing ring around the cathode to increase the radial focusing fields at the point of emission. This ring can be described as a half torus with major radius of 2.1 mm and inner radius of 0.2 mm. Further optimisation of this structure would be needed for a physical system, but this simple implementation was sufficient for the present work. Although the addition of the focusing ring did slightly affect the RF solution, the frequency, shunt impedance and quality factor all changed by less than 1%, so no further corrections were made for this.

In general, the field amplitudes obtained from an eigenmode solver are normalised in a manner specific to a given solver implementation, and must then be scaled to levels appropriate for the need at hand. However in the present instance, the original amplitudes were adequate and no further scaling of the fields was required.

Table 9: Various figures of merit extracted from the simulations. "Single Cell" refers to a single uncoupled accelerating cell, whilst the final structure is that shown in Figure 4. For comparison purpose, the values published by St. Aubin are also included. Note that the uncertainties here refer only to the numerical noise introduced by the mesh size – no other uncertainties are included.

Metric	Single cell values	Single cell values published in [19]	Complete Waveguide
Frequency (MHz)	3007.5 ± .5	3007.23 ± .01	2998.5±.5
Quality Factor Q0	17509	17521.3	16542
Shunt Impedance (MΩ/m)	165.5	165.24	109.5 (effective)
Transit Time Factor	.8371	.8381	N/A

3.3. Particle Trajectories & Beam assessment

Figure 6 shows a representative snapshot of the particle trajectories (Multimedia view). The particles were scored at the accelerator exit for assessment. The following values were calculated; mean target current 146.3 mA, Median Energy 5.8 MeV, Energy spread (Interquartile range, IQR) 0.1 MeV, spot size (Full width at tenth maximum, FWTM) 1.5 mm, and normalised emittance 6.5 mm mrad. The geometric emittance is 0.6 mm mrad. The energy and spatial distributions of the beam are shown in Figure 7. The spot size is quantified by fitting a circle to the 10^{th} percent intensity iso-line (effectively FWTM). Although we have not explicitly modelled the bremsstrahlung phase space resulting from this electron beam, the spatial and energy parameters listed above are comparable with published values [25, 26]. Also, relatively little sensitivity in radiation dose distributions has been shown to the target electron beam parameters - the most important factors are mean energy and current [26]. As such, we can conclude that an RF source based accelerator is capable of producing an electron beam suitable for radiotherapy.

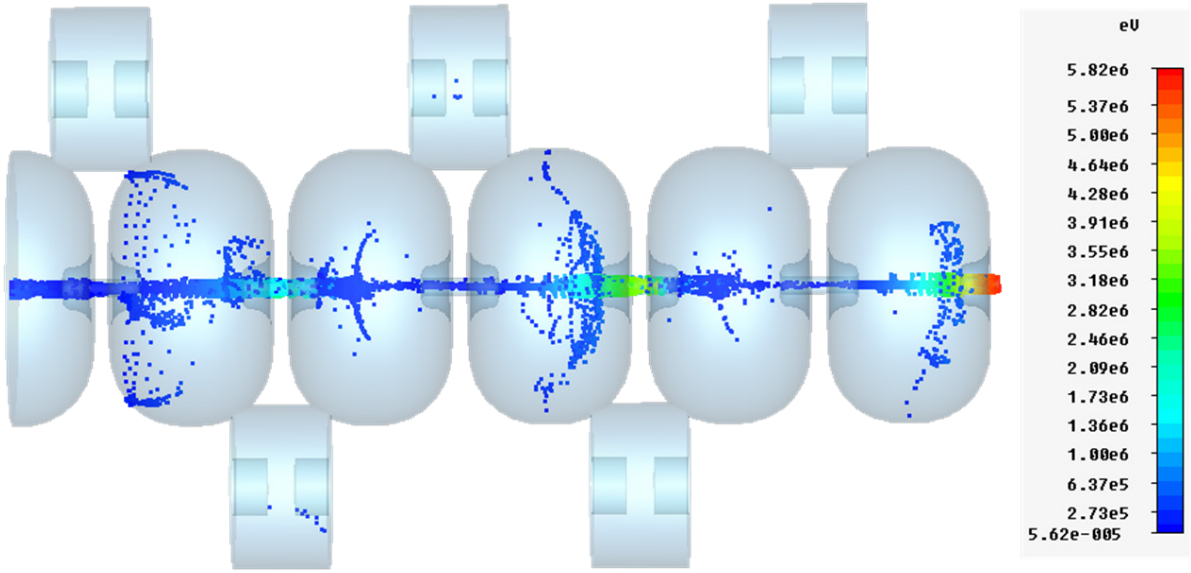

Figure 57: Electron trajectories. Note the way that low energy electrons are deflected when they pass through the nose cones at the wrong phase (Multimedia View)

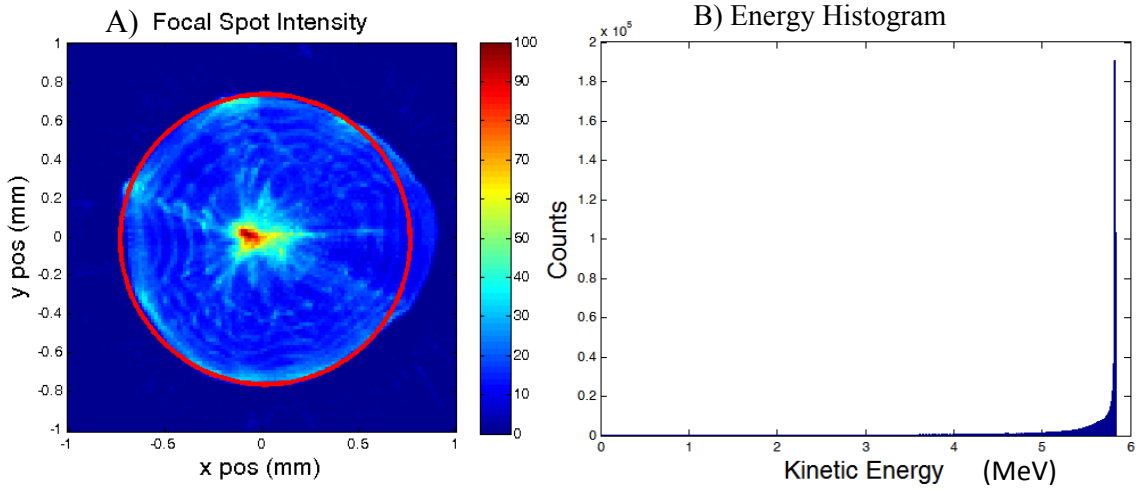

Figure 58: (A) Spatial distribution at the target. FWTM is 1.5 mm (red circle) (B) Energy histogram at the target. Median Energy is 5.8 MeV, Interquartile range is .1 MeV.

3.4. Back-bombardment power

Figure 8 shows the instantaneous power (defined as the energy delivered every 10 ps) and the bunch power, defined as the energy delivered in a bunch multiplied by the frequency of the RF fields (2998.5 MHz). It can be seen that the bunch power rises steadily before plateauing; this behaviour occurs because the mean energy of the back bombarded electrons increases as the forward directed beam propagates further down the accelerator. The steady state back bombardment power is 23.2 kW for the RF system, and 14.3 kW for the DC system. The mean electron energies are 0.16 MeV and 0.13 MeV respectively. These numbers take into

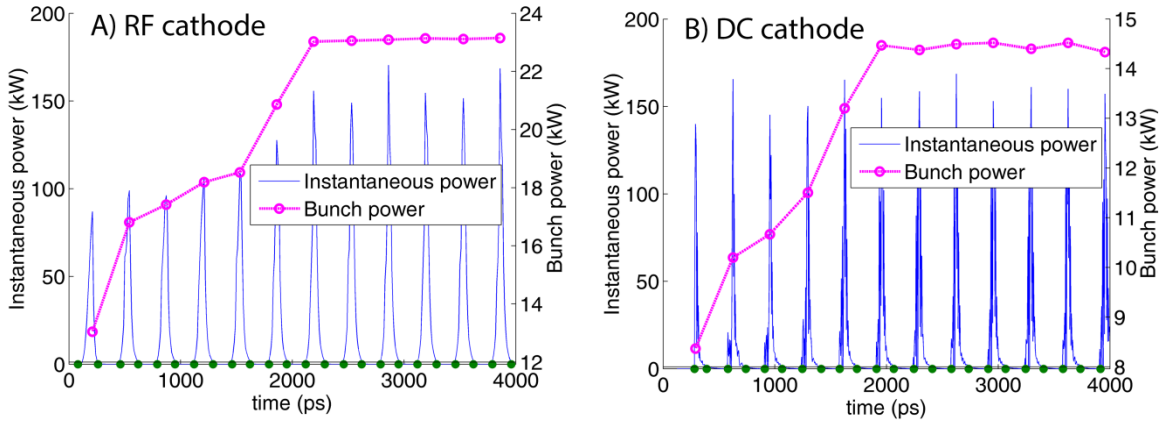

Figure 59: Back bombardment power for A) the RF system proposed in this work, and B) a DC based system which has previously been published. The bunch power indicates the mean power delivered per bunch, and steadily rises before reaching a steady state. The bunch segmentation is indicated by the dots along the horizontal axis.

account all back accelerated electrons. If electrons with a radial coordinate greater than 2 mm (the cathode radius used in this work) are filtered out then the RF back bombardment power decreases to 19.6 kW, whilst the DC back bombardment power remains unchanged. It is

important to note that none of the models tested here were actually designed to mitigate back bombardment. It is highly likely that the simulated back bombardment power could be greatly reduced for both systems modelled here – this is discussed in more detail in section 4. The stated results are during the ~5 us beam pulse – the overall mean will be around an order of magnitude lower than this given medical systems are typically operated with a duty cycle of 0.1% [2].

3.5. Performance in magnetic fields

The last step of the study was to compare the zero magnetic field behaviour of the novel accelerator with its performance in magnetic fields. Figure 9(A) shows the target current as a function of SID within the fringe field of the 1 T superconducting magnet – for comparison purpose, the current of the conventional diode gun is also plotted. Note that since the accelerator itself is not strongly affected by in-line fields, the losses in target current are proportional to the losses in the electron gun plotted here.

It can be seen that the target current of the proposed system is far more robust to operation in in-line magnetic fields, experiencing a maximum of 3 % current loss versus 85 % for the conventional system. Again, it can be seen that these fields do not affect the behaviour of the accelerator. As such, we conclude that an RF gun based accelerator is capable of robust performance without magnetic shielding in a wide range of in-line magnetic fields.

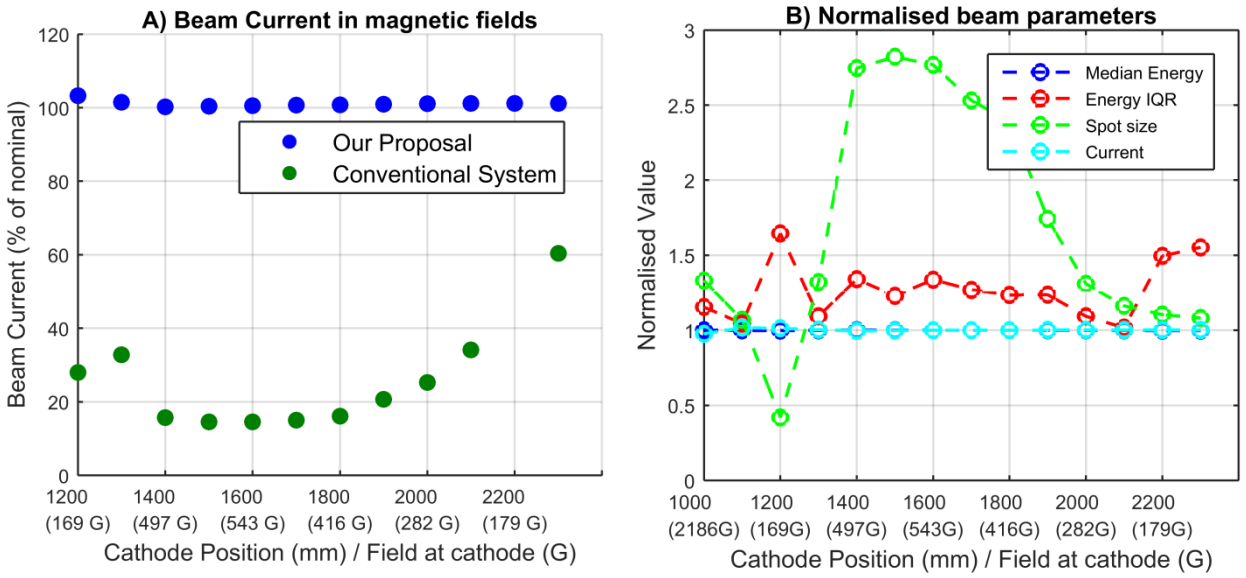

Figure 60: (A) Comparison of the current loss versus distance from magnet isocenter for proposed system and a conventional DC electron gun. (B) Various beam metrics at the target versus isocenter position. All metrics are normalized to the zero field values from section 3.3.

In Figure 9B, the median energy, energy interquartile range, spot size (FWTM) and current are plotted. Each metric is normalised to the zero field values from section 3.3; respectively 5.8 MeV, 0.1 MeV, 1.5 mm, 143.6 mA. It can be seen that whilst the median energy and current are barely changed by the addition of in-line fields, the spot size and interquartile range undergo large deviations compared to the zero field values. The maximum value of the

interquartile range is 0.16 MeV, which is still a very small energy spread. However the spot size ranges from 0.32 mm to 2.2 mm – this is discussed further in section 4.

4. Discussion

In this work we have proposed a novel medical electron accelerator with robust target current in a wide range of axial magnetic fields. Further, it does not require ferromagnetic shielding and so does not interfere with the MRI scanner. There are currently three MRI-Linac systems under development. Two utilise the in-line configuration tested in this work, while the third utilises the perpendicular configuration. Whilst the proposed accelerator design is likely to be slightly more robust to operation in perpendicular fields due to the stiffer beam, this has not been tested in this work. The potential improvement will not be anywhere near as stark as for the in-line case, and a combination of careful placement in low field and magnetic shielding would still be required for optimal performance.

One interesting result is the large variations which occur in spot size as a result of in-line magnetic fields. This is a result of well understood magnetic lensing effects and is explained by Busch's theorem [27]. These effects will occur regardless of the electron source used and have not previously been reported in other work looking at in-line systems. Although this will impact penumbral width, it is expected to have minimal impact on clinical dose distributions, as a number of studies have found that the radiation dose is quite insensitive to the spot size [22, 26, 28]. More significant effects may be the local heat load on the tungsten target and source occlusion in small field radiotherapy. Both these effects need be explored in more detail in future work.

We have assumed perfect alignment of the accelerator with the magnetic field in this work. In reality, perfect alignment is not possible. Misalignment between the accelerator and magnetic field increases the radial fields the electrons are subject to, which will cause the electron beam to bend. For small offsets, this will result in a shift in the focal spot, whilst for large enough offsets, partial or full beam loss could result. However, such effects were previously investigated by St. Aubin and found to be small even for exaggerated misalignments [29]. Although the fields tested here are much stronger than those previously investigated, preliminary investigations suggest that the impact of misalignments remain small as long as the alignment is within about 2 mm and the SID is greater than 1 meter. For SIDs less than 1 meter, the magnetic field becomes much stronger, and substantial beam loss could occur. In this scenario more precise alignment would be required.

There are a number of potential downsides to the use of RF gun based electron sources for medical accelerator systems which have not been simulated here, and indeed would be difficult to simulate within the present framework. One of these is beam stability. Operating an electron source in the space charge limited regime provides inherent stability, as the resultant current is dependant only on voltage – something which can be controlled with high accuracy and precision. However, in the temperature limited regime, the current is

dependent on the cathode temperature and the work function, both of which are more difficult to control. The temperature can fluctuate due to the large temperature gradients present and due to the energy deposited by back accelerated electrons. In existing systems this is compensated for by active feedback, resulting in pulse to pulse variations of less than 1% [23]. As such, we believe that with the present ability to accurately monitor integrated dose, this could also be overcome.

A second issue which would need to be quantified for the proposed system is that of beam on/ off latency. Most conventional medical Linacs 'gate' the beam using a triode electron gun. In the case of systems which do not use triode guns, rather long beam on/off latencies have been observed [30]. The current system would likely exhibit similar latency – however it is worth noting that the same criticism can be made of currently proposed MRI-Linac systems, which appear to be using diode rather than triode guns at this stage. There have also been triode based RF guns proposed – such a structure may be able improve the system latency to that achievable with DC triode guns [31].

Perhaps the most substantial barrier to clinical implementation of a system such as that described here is the increased back bombardment power deposited by back accelerated electrons. The results presented in this work suggest that the back bombardment power is only around 1.5 times higher for the RF cathode than the DC system. However, these results are probably not representative of what would be obtained in real systems for several reasons. Firstly, the field magnitude in the first half cell is typically reduced compared to the field in the rest of the accelerator. This was not done in this work, and would reduce the amount of back bombarded power. In the case of the DC system, the anode drift tube geometry could also be optimized such that back bombardment was substantially reduced. As such, the results presented in section 3.4) must be considered preliminary. We were not able to find any published literature on the typical extent of back bombarded power in medical DC electron guns or any information on how much back bombardment would be acceptable or unacceptable. However, from speaking to industry representatives this appears to be an area to which each accelerator manufacturer has devoted substantial in-house effort. A naïve interpretation of the data presented in 3.4) coupled to CSDA electron ranges [32] and data on tungsten cathodes presented in reference [18] would lead one to conclude that the temperature of the cathode proposed in this system might rise by over 100 degrees during the beam pulse. However, this ignores that both radiative and conductive heat dissipation is occurring. Based on the published data we could find (which concerns synchrotron injector guns), a better (but still extremely rough) estimate would be around 50 degrees [18]. This would still be a serious issue, resulting in a change in the emitted current by roughly a factor of two, which would also effect the beam loading and hence energy spectrum of the beam at the target. On the other hand, if the dose per pulse remains consistent (i.e. the effect is reproducible), this may not be as large a problem therapeutically as it is for high energy applications where consistent beam quality is extremely important. Further investigation into back-bombardment, its impact on therapeutic beam quality, and mitigation strategies would require full Monte Carlo simulation of electron transport through the cathode coupled to models (e.g. FEM) of thermal energy loss. However, this is beyond the scope of this work. A

large amount of strategies to mitigate the back-bombardment effect in RF guns have already been published, and given that there is already at least one system in existence operating at similar pulse lengths and energies as would be required for this system (The Kyoto University free electron laser injector) it is reasonable to believe that this effect could be managed [18, 31, 33].

A potential limitation of this work is that no power input port was incorporated into the simulation. It has been previously shown that the presence of a coupling port does introduce further asymmetry into the beam [19]. It will also lower the loaded quality factor and shunt impedance of the final structure. However, this work is intended as a proof of principle, and the explicit modelling of an input power port is not anticipated to significantly affect the results. The required input power is dependent on the power losses within the accelerator. As the sources of loss (losses to the conducting walls and to the beam) are very similar to previously published work, the required input power will also be similar to this – around 2.3 MW [19]. In a real system, further asymmetry would be introduced into the final electron beam distribution, however this is the case regardless of the electron source, and is not a significant problem for therapeutic beams in any case [25, 26]. Previous publications have used a technique whereby the first and last bunch (or half bunch) was removed from the analysis pipeline in order to remove 'end effects' from the simulation. We observed minimal bunch to bunch variation in this study (less than 3%) and as such, all macro particles reaching the target were included in the analysis.

We have made minimal effort to optimise the RF structure for a temperature-limited cathode – for instance, further optimisation of the radial focusing fields around the cathode could be undertaken, and the spatial energy distribution at the target could be optimised by lowering the fields in the first half cavity [19]. The focusing ring used in this study would not be suitable for a physical system, as both joule heating from the RF fields and thermal isolation from the hot cathode must be considered. However, both issues have been solved in many other RF guns previously and as such are not anticipated to present major challenges [13, 23]. Again, the purpose of this study was to provide a proof of principle and whilst the above are all interesting directions for future research, they do not counter the proof of principle that has been provided.

5. Conclusion

We have investigated the use of an RF electron source based linear accelerator for delivering MRI-Linac radiation therapy. Through the computational simulations we have shown that such a system is capable of generating an electron beam suitable for therapeutic applications. We have also shown that such a system is far more robust than conventional systems to the presence of in-line magnetic fields, and as such could be an ideal solution for next generation in-line MRI-Linac systems.

Acknowledgements

Brendan Whelan would like to acknowledge the Centre for Oncology Education and Research Translation (CONCERT) and Cancer Institute NSW for scholarship support. The authors would like to acknowledge funding support from the NHMRC and the NIH.

References

1. Ménard, C. and U. van der Heide. *Introduction: Systems for magnetic resonance image guided radiation therapy*. in *Seminars in radiation oncology*. 2014. Elsevier.
2. Karzmark, C., C.S. Nunan, and E. Tanabe, *Medical electron accelerators*. 1993: McGraw-Hill, Incorporated, Health Professions Division.
3. Constantin, D.E., R. Fahrig, and P.J. Keall, *A study of the effect of in-line and perpendicular magnetic fields on beam characteristics of electron guns in medical linear accelerators.* Medical physics, 2011. **38**(7): p. 4174-4185.
4. St Aubin, J., S. Steciw, and B. Fallone, *Effect of transverse magnetic fields on a simulated in-line 6 MV linac.* Physics in medicine and biology, 2010. **55**(16): p. 4861.
5. Keall, P.J., M. Barton, and S. Crozier. *The Australian magnetic resonance imaging–linac program*. in *Seminars in radiation oncology*. 2014. Elsevier.
6. Overweg, J., B. Raaymakers, J. Lagendijk, and K. Brown, *System for MRI guided radiotherapy.* Proc. Int. Society Magnetic Resonance in Medicine, 2009: p. 593.
7. St .Aubin, J., D. Santos, S. Steciw, and B. Fallone, *Effect of longitudinal magnetic fields on a simulated in-line 6 MV linac.* Medical physics, 2010. **37**(9): p. 4916-4923.
8. St. Aubin, J., D. Santos, S. Steciw, and B. Fallone, *Effect of longitudinal magnetic fields on a simulated in-line 6 MV linac.* Medical physics, 2010. **37**(9): p. 4916-4923.
9. Constantin, D.E., L. Holloway, P.J. Keall, and R. Fahrig, *A novel electron gun for inline MRI-linac configurations.* Medical physics, 2014. **41**(2): p. 022301.
10. Santos, D., J. St. Aubin, B. Fallone, and S. Steciw, *Magnetic shielding investigation for a 6 MV in-line linac within the parallel configuration of a linac-MR system.* Medical physics, 2012. **39**(2): p. 788-797.
11. Kolling, S., B. Oborn, and P. Keall, *Impact of the MLC on the MRI field distortion of a prototype MRI-linac.* Medical physics, 2013. **40**(12): p. 121705.
12. Yun, J., J. St. Aubin, S. Rathee, and B. Fallone, *Brushed permanent magnet DC MLC motor operation in an external magnetic field.* Medical physics, 2010. **37**(5): p. 2131-2134.
13. Rimjaem, S., K. Kusoljariyakul, and C. Thongbai, *RF study and 3-D simulations of a side-coupling thermionic RF-gun.* Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 2014. **736**: p. 10-21.
14. Cronin, J., *Modern dispenser cathodes.* IEE Proceedings I (Solid-State and Electron Devices), 1981. **128**(1): p. 19-32.
15. Seely, S., *Work function and temperature.* Physical Review, 1941. **59**(1): p. 75.
16. Semicon Associates. Available from: http://www.semiconassociates.com/products/microwave.aspx.
17. Reiser, M., *Theory and design of charged particle beams*. 2008: John Wiley & Sons.
18. Bakr, M., R. Kinjo, Y. Choi, M. Omer, K. Yoshida, S. Ueda, M. Takasaki, K. Ishida, N. Kimura, and T. Sonobe, *Back bombardment for dispenser and lanthanum hexaboride cathodes.* Physical Review Special Topics-Accelerators and Beams, 2011. **14**(6): p. 060708.
19. St. Aubin, J., S. Steciw, and B. Fallone, *The design of a simulated in-line side-coupled 6 MV linear accelerator waveguide.* Medical physics, 2010. **37**(2): p. 466-476.
20. Weiland, T., *High precision eigenmode computation.* Part. Accel., 1996. **56**: p. 61-82.
21. Baillie, D., J. St. Aubin, B. Fallone, and S. Steciw, *FEM design and simulation of a short, 10 MV, S-band Linac with Monte Carlo dose simulations.* Medical physics, 2015. **42**(4): p. 2044-2053.
22. St. Aubin, J., S. Steciw, C. Kirkby, and B. Fallone, *An integrated 6 MV linear accelerator model from electron gun to dose in a water tank.* Medical physics, 2010. **37**(5): p. 2279-2288.
23. Borland, M., *A high-brightness thermionic microwave electron gun*. 1991, to the Department of Applied Physics.Stanford University.

24. Kowalczyk, J.M., M.R. Hadmack, and J.M. Madey, *Measurement of back-bombardment temperature rise in microwave thermionic electron guns.* Review of Scientific Instruments, 2013. **84**(8): p. 084905.
25. Jaffray, D., J. Battista, A. Fenster, and P. Munro, *X-ray sources of medical linear accelerators: Focal and extra-focal radiation.* Medical physics, 1993. **20**(5): p. 1417-1427.
26. Chibani, O., B. Moftah, and C.-M.C. Ma, *On Monte Carlo modeling of megavoltage photon beams: a revisited study on the sensitivity of beam parameters.* Medical physics, 2011. **38**(1): p. 188-201.
27. Kumar, V., *Understanding the focusing of charged particle beams in a solenoid magnetic field.* American Journal of Physics, 2009. **77**(8): p. 737-741.
28. Sheikh-Bagheri, D. and D. Rogers, *Sensitivity of megavoltage photon beam Monte Carlo simulations to electron beam and other parameters.* Medical physics, 2002. **29**(3): p. 379-390.
29. Aubin, J.S., D. Santos, S. Steciw, and B. Fallone, *Effect of longitudinal magnetic fields on a simulated in-line 6 MV linac.* Medical physics, 2010. **37**(9): p. 4916-4923.
30. Freislederer, P., M. Reiner, W. Hoischen, A. Quanz, C. Heinz, F. Walter, C. Belka, and M. Soehn, *Characteristics of gated treatment using an optical surface imaging and gating system on an Elekta linac.* Radiation Oncology, 2015. **10**(1): p. 68.
31. Masuda, K., K. Kusukame, T. Kii, H. Ohgaki, H. Zen, T. Fukui, Y. Nakai, K. Yoshikawa, and T. Yamazaki. *Particle simulations of a thermionic RF gun with gridded triode structure for reduction of back-bombardment*. in *Proceedings of the 27th International Free Electron Laser Conference*. 2005.
32. Berger, M., *ESTAR, PSTAR, and ASTAR: Computer programs for calculating stopping-power and range tables for electrons, protons, and helium ions.* Unknown, 1992. **1**.
33. Kii, T., K. Yamane, I. Tometaka, K. Masuda, H. Ohgaki, K. Yoshikawa, and T. Yamazaki, *Improvement of electron beam properties by reducing back-bombardment effects in a thermionic RF gun.* Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 2003. **507**(1): p. 340-344.

Chapter 6
Experimental characterisation of an RF gun in magnetic fields

1. Introduction

MRI-Linac therapy holds great promise for increasing radiation therapy patient outcomes, and is under development by several academic and commercial groups [1-5]. It is a requirement of MRI-Linac therapy that that a linear accelerator (linac) functions within the magnetic fringe field of a MRI magnet. Conventional medical linacs are extremely sensitive to external magnetic fields, as has been detailed in a number of publications [6-8]. As such, integration into an MRI Linac system requires that the linac is (1) operated in a very low magnetic field (requiring intelligent magnet design and magnetic shielding) or (2) redesigned to be more robust in magnetic fields. The latter option was previously investigated in-silico for the 'in-line' MRI-Linac configuration (Chapter 5, reference [9]). In this work, a novel medical accelerator based on an RF gun configuration was proposed. The performance of the proposed device was simulated within the fringe field of a 1 Tesla MRI magnet, and found to operate with a maximum of 3% current loss when operated at Source to Isocentre Distances (SIDs) of 0.9-2 meters. This compared with a (simulated) conventional system, which had a maximum current loss of 85% over the same conditions. As well as this, it was found that magnetic lensing effects within the accelerator could cause large variations in spot size at the x-ray target [9].

The purpose of this work was to collect experimental data supporting these findings (robustness to axial fields, and magnetic lensing), as well as to benchmark the accuracy with which the experimental results could be predicted in-silico. The latter data is important not only for the medical RF-gun approach, but also in the broader field of MRI-Linac therapy. Whilst a large number of publications exist in which the impact of magnetic fields on electron accelerators has been quantified in-silico [6-8, 10-14], this is the first time that experimental verification of in-silico results has been carried out. In order to provide experimental data, a purpose built beamline was designed, constructed, and installed at SLAC national laboratory. The electron source was a 1.5 cell RF gun, and the magnetic field was generated using room temperature magnetic coils.

2. Methods and Materials

I. Experimental beam line

A purpose built beam line was designed and constructed at SLAC national laboratory. A CAD layout of this design is shown in Figure 1. The beam line was installed in a radiation shielded bunker and utilized the infrastructure from the (now defunct) SSRL gun test facility [15]. The beam line consists of a 1.5 accelerating cell RF gun (detailed in section II) located between two magnetic coils (detailed in section III). In order to monitor the beam current, a toroid current monitor was installed (Figure 1D). In order to monitor the beam size, two fluorescent screens with viewports and cameras (Allied Vision Technology, Manta G-033) were installed (Figure 1B). The screens can be moved in and out of the beam line with the use of pneumatic actuators (Figure 1C). The first of these screens is a 10 mm piece of phosphor, whilst the second is a 18 mm circular screen of Yttrium aluminum garnet doped with cerium (hereafter referred to as YAG) [16]. These choices were made largely on the

basis of what was available at the time of construction. The beam is disposed of by dumping it into water cooled and a radiation shielded slug of copper. Prior to installation, the entire beam line was vacuum baked to ~200℃ over two days. Vacuum is maintained by two ionic vacuum pumps; one in the gun and one near the beam dump RF power was supplied to the gun by a XK-5 SLAC Klystron [17].

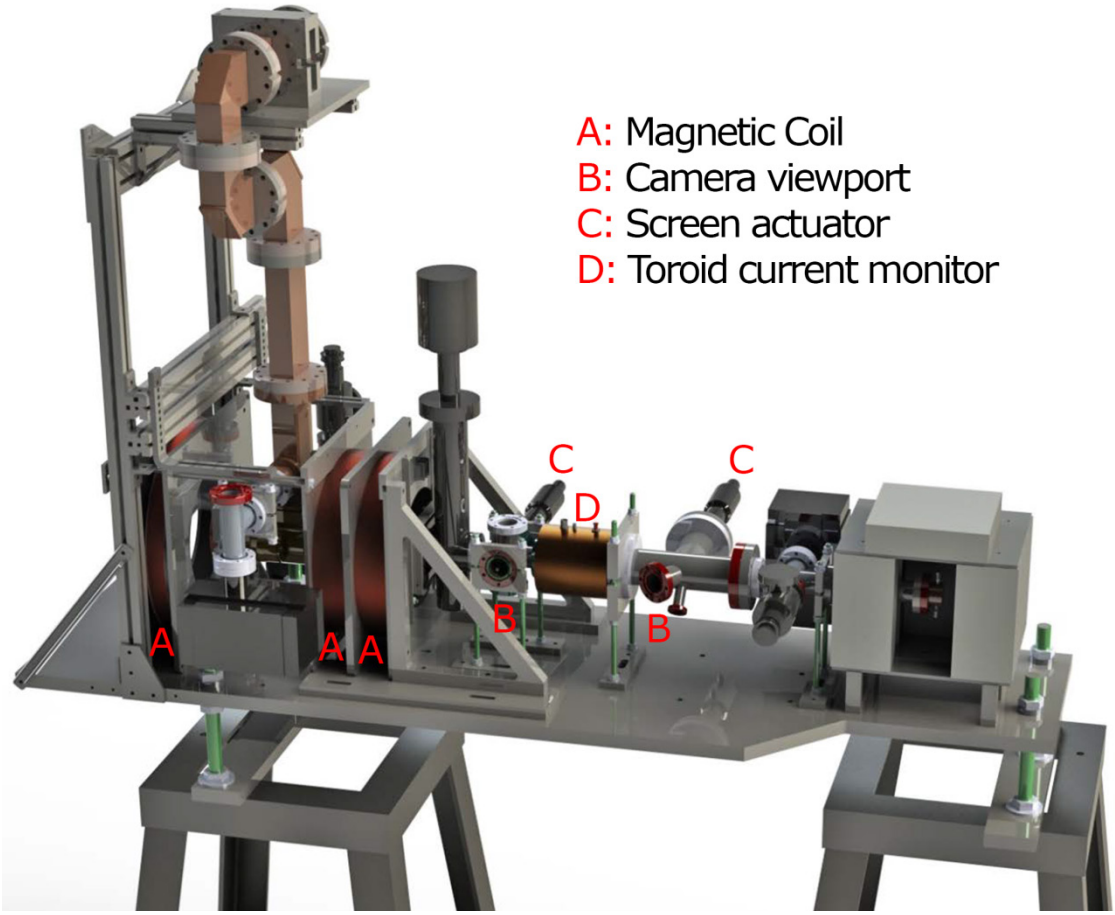

Figure 61: The experimental beam line which was developed and used for these experiments. The source of the electron beam is a 1.5 accelerating cell RF gun, which is located between the three magnetic coils. The labels indicate particular components of interest; A) Magnetic coils generate an axial magnetic field along the beam line (note that the far left coil was not used in these experiments) B) Camera viewports (cameras can be installed to capture images of the beam colliding with fluorescent screens) C) Screen actuators (these pneumatically move fluorescent screens in or out of the beam line) D) Toroid beam current monitor.

Figure 2 shows the geometry of the YAG screen and camera relative to the beam. The screen is tilted 60° relative to the beam in the x direction. Trigonometric calculations can be used to show that the width in the horizontal axis is half the width of the vertical axis from the beam perspective. A coordinate system for the images was constructed based on the known dimensions and orientations of the screen.

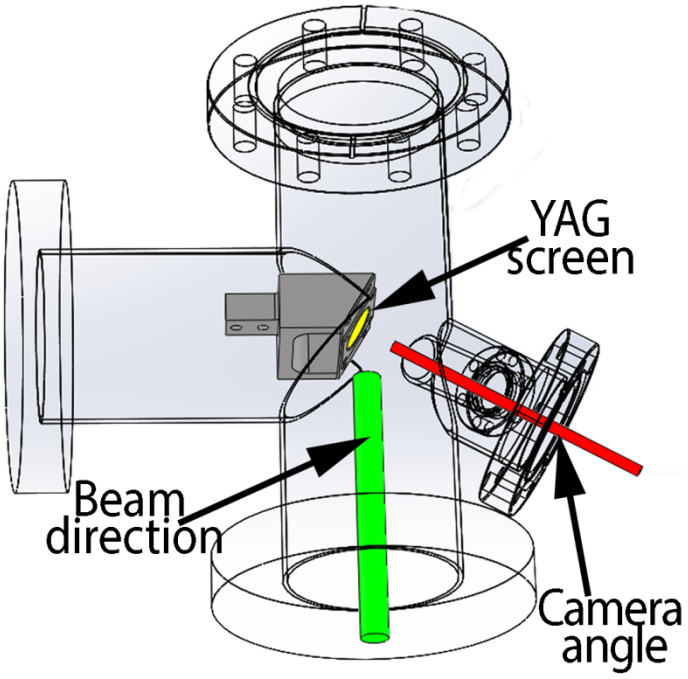

Figure 62: A close up of the second viewport from Figure 1, showing the screen geometry which was used for image collection. The screen has an angle of 60° relative to the beam in the x direction and can be retracted in and out of the beam line using pneumatic actuators.

II. The SSRL RF gun and computational simulation of beamline

We did not have the resources available to build an RF gun prototype from scratch; therefore we utilized an existing RF gun which was available at SLAC. The RF gun which was tested in this work is normally used as the injector for the Stanford Synchrotron Radiation Light Source (SSRL) [18]. The SSRL gun is a thermionic side coupled S band RF gun consisting of 1.5 accelerating cells, and can generate energies up to ~3 MeV. In contrast, a typical low energy medical linac would consist of 5.5 accelerating cells and produce energies of ~ 6 MeV[9, 19]. Despite the differences in these two systems, the SSRL gun is similar enough to the design proposed in chapter 5 in that it provides an excellent basis for experimental testing. An in-silico model of this gun was developed based on the information available in ref. [18], and using CST particle studio (Darmstadt, Germany) as detailed in Chapter 5. Whilst the geometry of the two accelerating cells is known, the exact geometry of the side coupled cell is not. Therefore, we developed a model of the gun which did not incorporate the side coupling cell. This means that the field in each cell is solved separately, before being combined with the correct scaling in the PIC solver based on ref. [18]. The magnetic fields generated by the magnetic coils (Section III) were also added into the PIC model, which was then solved for coil currents between 0 and 20 A.

III. Magnetic coils and magnetic lensing theory

To generate magnetic fields along the beam line, two magnetic coils (Stangenes Industries) were installed (Figure 1). These coils are powered in parallel by a single DC power source capable of generating currents up to 160 A, meaning a theoretical maximum of 80 A was

available for each coil. When the coils are operated, they heat up due to resistive losses. The maximum coil temp is limited by the melting temperature of the insulation around the copper wire, which based on feedback from the manufacturer, was 'at least' 130° C. FEM models of the coils (Comsol, Burlington) were developed based on gauss probe measurements of the axial field and geometry (external dimensions and wire dimensions) of the coils, and time domain thermal modeling was carried out in using a convective heat flux condition on the coil boundary. The heat transfer coefficient was set to 10 W/K/m^2[26]. These simulations suggested that it should be possible to utilize at least 40 A through each coil for five minutes at a time without the temperature exceeding 100 degrees Celsius. Thermocouples were installed on the coils which cut power to the coils if there temperature exceeded 70˚C (a large safety margin was used as the uncertainty in the FEM models is not quantified). The axial magnetic field produced by these coils at various locations along the beam line is shown for a coil current of 30 A. The magnetic field scales linearly with coil current.

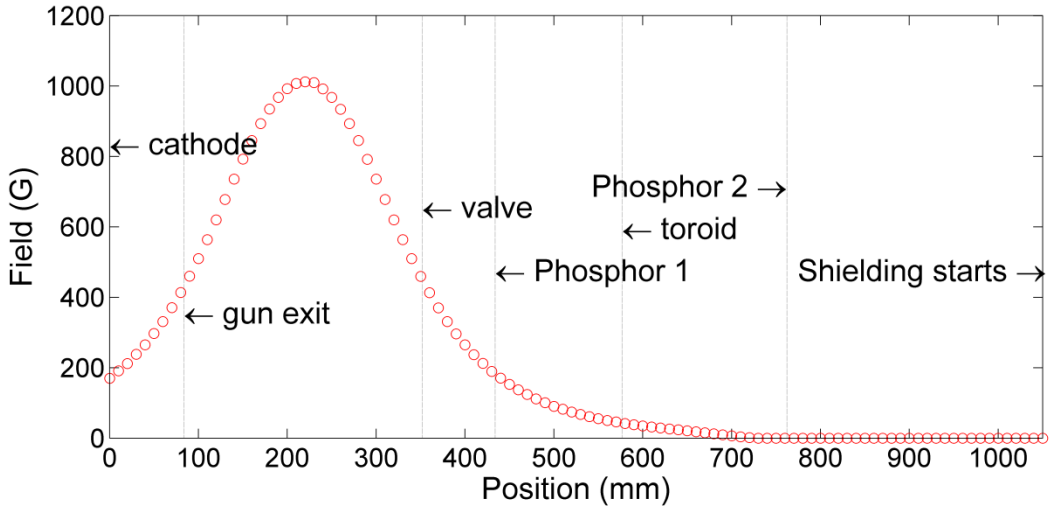

Figure 63: The axial magnetic field amplitude produce by running 30 A through all three coils. The 'solenoid' field refers to the two downstream coils (Figure 1) whilst the 'bucking' field refers to the single upstream coil. The location of various important beam line diagnostics is indicated on the figure.

The theoretical description of the impact of electron kinematics in axial magnetic fields is given by Busch's theorem, which describes the trajectory of a charged particle in axial fields as a function of conserved canonical angular momentum [20]:

$$\gamma m_0 r^2 \theta' + \frac{e_0}{2\pi}\Phi_B = \text{constant} \qquad \text{Equation 23}$$

In equation 1, γ is the relativistic gamma factor, m_0 is the electron rest mass, e_0 the electron charge, r is the radial coordinate of the electron, θ' is the angular velocity, and Φ_B is the magnetic flux passing through a cylinder of radius r. Equation 1 is difficult to interpret at first glance, but it describes the fact that charged particles travelling in axial fields will tend to circle around the field lines, and each circular trajectory will tend to cross through the axis

[26] http://www.engineersedge.com/heat_transfer/convective_heat_transfer_coefficients__13378.htm

of magnetic symmetry. When the strength of the axial field is increased, the radius of each individual electron trajectory decreases, and vice versa. As such, electron beams in axial fields tend to exhibit 'scalloping' behaviour where the width of the beam increases and decreases as a function of distance [6]. This means that under various conditions, either focusing or defocusing can be caused by axial fields, and exactly what is observed depends on both the field which is applied, and the location of the screen. An excellent overview of the effect of axial magnetic fields on electron trajectories is given in [20].

IV. Control of beam energy, beam current, and RF power

The gun utilizes a 6 mm thermionic tungsten dispenser cathode and is operated in the temperature limited mode of emission current (chapter 5). As such emission is controlled by the power supplied to the cathode heater. The beam current is simply the emission current multiplied by the proportion of this current which actually exits the gun. This proportion is determined by the magnitude of the electric field in the gun, and hence RF power [18]. Therefore, the beam current exiting the gun is a function of both the cathode heater power and the RF power levels. The RF power generated by the Klystron is determined by three factors: the high voltage (HV) setting in the klystron, the klystron beam current (which in fact is also a function of the HV), and the power of the input RF power, called the drive power. In practice, the drive power was used to adjust the RF power level. This is because the Klystron and pulse forming network is 'happy' operating at a given high voltage, and moving it far from these levels results in performance degradations, such as less homogenous power pulses, multipacting in klystron, and electric break down in the klystron or PFN. In this work, the klystron was operated with a high voltage of ~230kV and a current of ~150 A. The drive power was set to ~100 W. The ~symbol is used as all of these parameters were tweaked at various points throughout the experiments to maximize stability. These settings resulted in a power output (forward power) from the klystron of ~4 MW. The forward power will be reflected at any location there is an impedance mismatch, the most important location being the gun itself. To optimize impedance matching to the gun, minor frequency tuning can be achieved by changing the water cooling temperature to control thermal expansion of the cavities. This means that there is both forward and backward travelling power in the RF waveguides, which was monitored using diode type detectors (Keysight technology, USA) in two locations: just outside klystron, and just before the ceramic window separating the RF waveguide from the gun. In order to avoid damage to the klystron, power reflected from the gun is directed to a cooled load using a RF ferrite circulator. The power in the gun (and hence beam energy) must be inferred from the forward and backward power at the power monitor outside the gun as there is no RF power probe inside this gun.

V. Experimental data collection

Prior to turning on the beam for the first time, RF processing, interlock testing, and radiation surveys were carried out. The first beam on was on the 31st of July 2016. The first data under magnetic fields was collected on the 1st of August with the following system settings: Klystron HV: 232 kV, Klystron beam current: 147 A, Forward RF power at gun: 3.75 MW, Backward RF power at gun: 0.13 MW, Inferred power inside the gun: 3.62 MW, Inferred beam energy (median): 3 MeV, inferred beam current (during beam pulse): 25 mA. Low

beam current was used, as this minimizes scattered radiation and the risk of overheating the beam diagnostics or creating vacuum faults. At the time of running, only one camera was available; since only one of the two installed screens could be utilized, it was decided to use the downstream YAG screen. This decision was made on the basis that this screen was less likely to saturate, less likely to cause vacuum issues, was larger, and was easier to access for camera installation. Next, current to the two solenoids was increased in 5 A up to 20 A (note that as they are wired in parallel, the current in each coil is half the current output of the power source). The bucking coil was not powered at this time. Images of the beam on the YAG screen were collected at each magnetic field strength.

The beam was next run on 6th of August; unfortunately in the afternoon of the 6th of August a vacuum fault developed on the beam line which meant we were not able to run the next day. It is believed that the electron beam was being scattered onto the ungrounded viewport window, which eventually arced to the nearest grounded component causing cracking and the resultant vacuum fault. This was being resolved at time of writing, with an aim to complete experimental data collection in October.

VI. Image analysis

Due to the fact that the vacuum leak developed before the experimental procedure was complete, the experimental data which exists to date is images of the beam colliding with the YAG screen with coil currents of 5, 10, 15 and 20 A. To quantitatively analyze these images, each image was masked such that all background except the screen was deleted. Next, a coordinate system was constructed based on the known dimensions and orientation of the YAG screen (Figure 2). The masked images were median filtered to remove speckle noise (filter size 3x3). Next, the coordinates of pixels having greater than 50% relative intensity were extracted. A confidence ellipse was fitted based on the covariance matrix of these points. From this ellipse, a measure of spot size and position was extracted. In addition, to obtain some measure of the total current intersecting the screen, the sum of all pixel values was calculated for each image. However, this result should be interpreted with caution, as the exact response of the YAG screen to energy and current is not well understood at this stage. The same procedure was run on the simulated data so that the two could be compared.

3. Results

I. Experimental Data

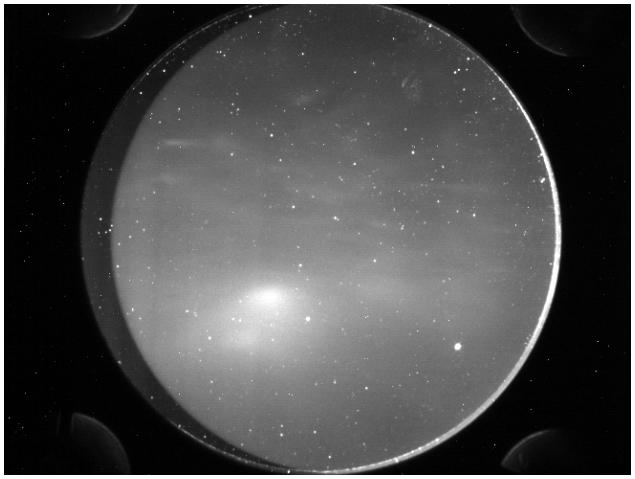

Figure 64: The first image acquired, showing the electron beam impinging onto the YAG screen. The speckle noise which can be seen is radiation damage to the camera. The bright and dark edges are simply the result of optical shadows.

Figure 4 shows the first collected image of the beam colliding with the YAG screen. The bright spot in this image is the electron beam. This image was taken at very low current and high camera gain as the beam current was carefully increased. The speckle noise seen in this image is due to radiation damage to the camera. Figure 5 shows beam images which were collected under the influence of magnetic fields. The ellipse which was fitted to the spot size is shown in red. The influence of the magnetic fields on the electron beam can be seen in the way the spot size rotates and changes size. The rotation of the electron beam around the centre of the screen indicates that imperfect alignment between at least one of the coils, beam line, YAG screen, and electron beam. This misalignment is not a huge surprise given that the beam line had to be constructed very quickly.

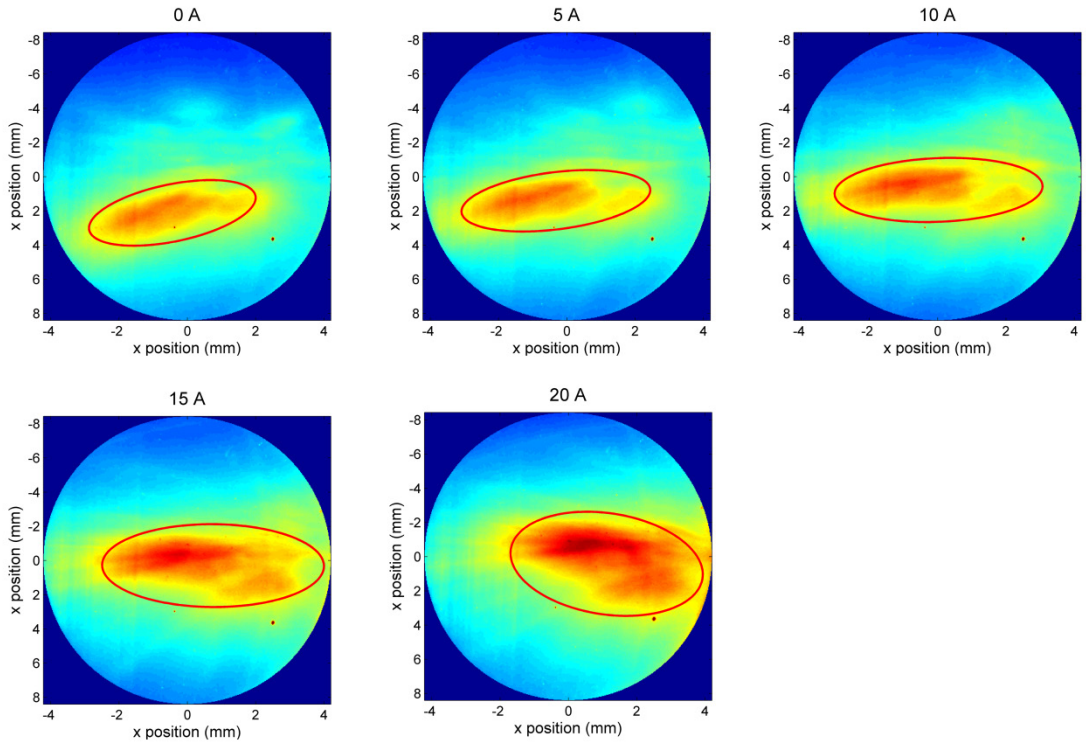

Figure 65: The images collected under magnetic fields. The raw image shown in figure 4 has been masked and median filtered to remove speckle noise. An ellipse has been fitted to the 50% intensity line of the image.

Figure 6 shows A) the spot size (as defined by the area of the ellipse in Figure 5 versus the current in the coils, and B) the sum of every pixel in the image, which is a surrogate for beam power. It can be seen that both the spot size and total YAG signal steadily increase as the coil current is turned up.

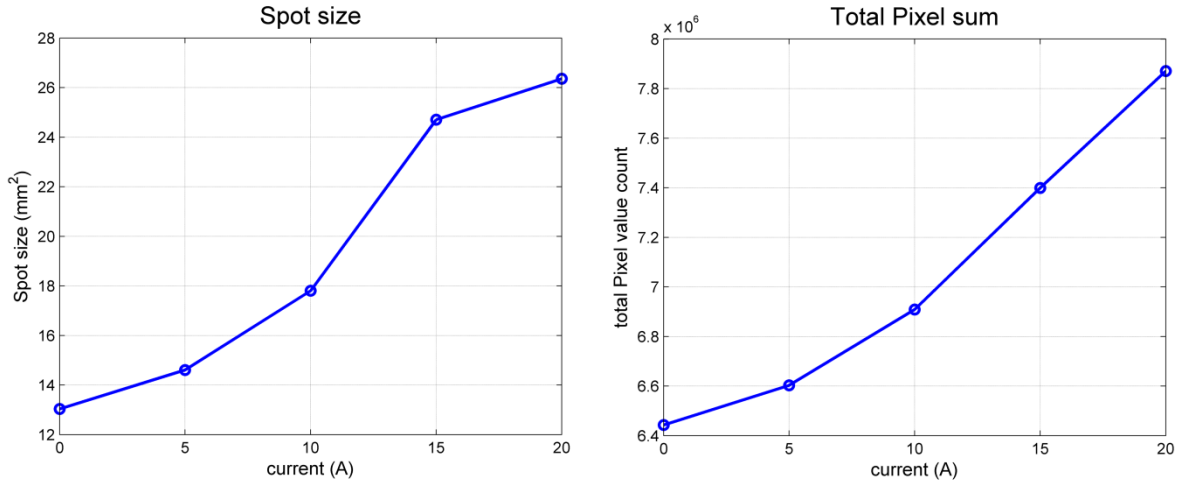

Figure 66: A) The spot size of the beam (defined by the ellipse in Figure 5) versus coil current. B) The sum of all pixel values in each image (Figure 5) versus coil current.

II. Simulated gun model

Figure 7 shows a 2D cross section of the gun model and the electric field distribution within it. Note that the actual gun has a side coupling port which couples the two cells, and an RF input port in the second cell. These have not been included in the gun model as their geometry is not known. Whilst both will perturb the electromagnetic field distribution in the gun, and hence the electron beam, simulation of the axial geometry should be adequate to obtain a first order estimate of the beam behavior. The following beam parameters were calculated at the exit of the gun model: median energy 2.7 MeV, beam current 25 mA, emittance 5.7 mm mrad, showing good agreement with ref. [18].

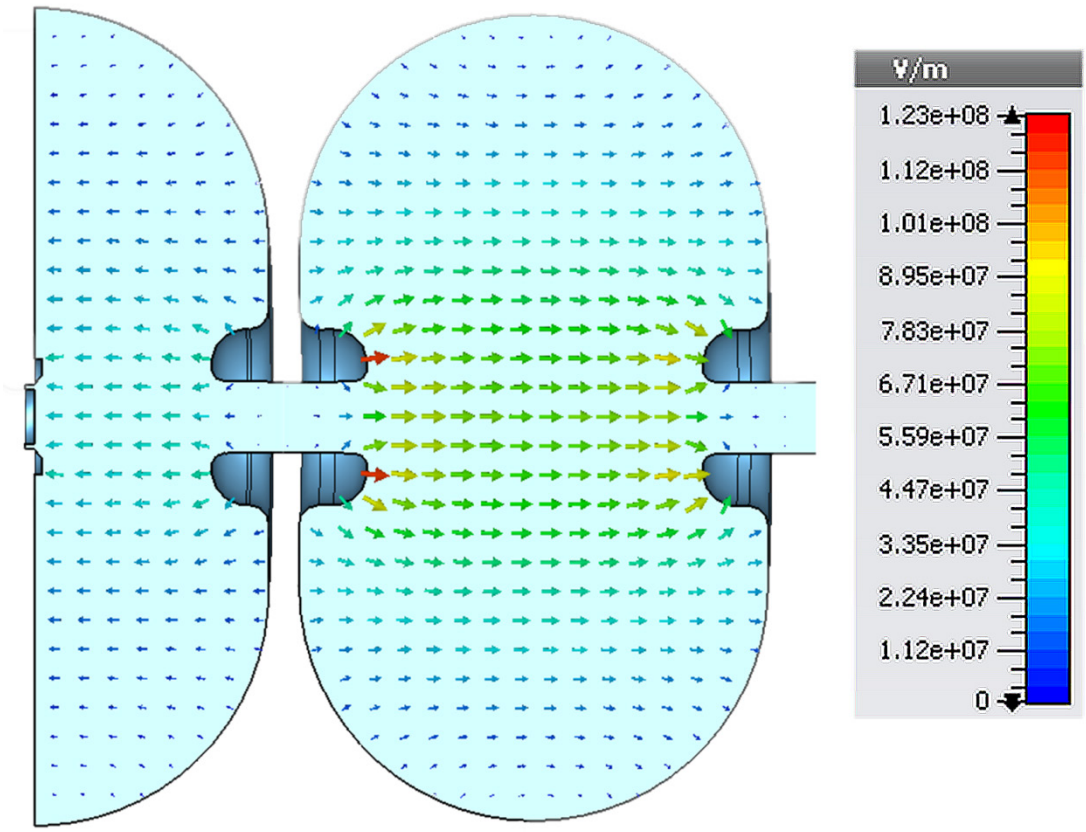

Figure 67: A 2D profile of the gun geometry, showing the electric field distribution at 0 phase.

III. Simulated data

Figure 8 and Figure 9 show the results of passing the simulated data through the same analysis pipeline as the experimental data (Section 2 VI). It can be seen that there is not particularly good agreement between the two data sources; In particular the spot size calculated by the 50% thresholding method yield is substantially smaller in the simulated data (~4 mm^2 versus ~20 mm^2). Also, the simulated data predicts a minimum in spot size at A (Figure 9) which was not seen in the experimental data (Figure 6). However, after this both data sets show that the beam size and total pixel count increases when the magnetic field is increased. A few comments should be made about the simulated data. Firstly, it can also be seen that the resolution of the simulated data is quite poor (Figure 8) which makes the images quite difficult to interpret. This low resolution is a result of the simulation settings, in particular the number of particles simulated. This can (and will) be improved, but comes at the cost of a substantial increase in simulation time. Also, it cannot be assumed that the direct comparison of Figure 8 and Figure 5 is valid; whilst Figure 8 shows the (normalised) total beam energy passing through each pixel, the signal in the YAG images will be some (potentially non-linear) function of current and energy. At this stage, this function is not well known and further review of the literature is needed to better understand this. Finally, whilst the position of the two screens in the beam line is the same, the geometry is different; whilst

the YAG screen is angled at 60 degrees with respect to the beam (Figure 2), the simulated screen is normal to it. An angled particle monitor is not supported by the CST software.

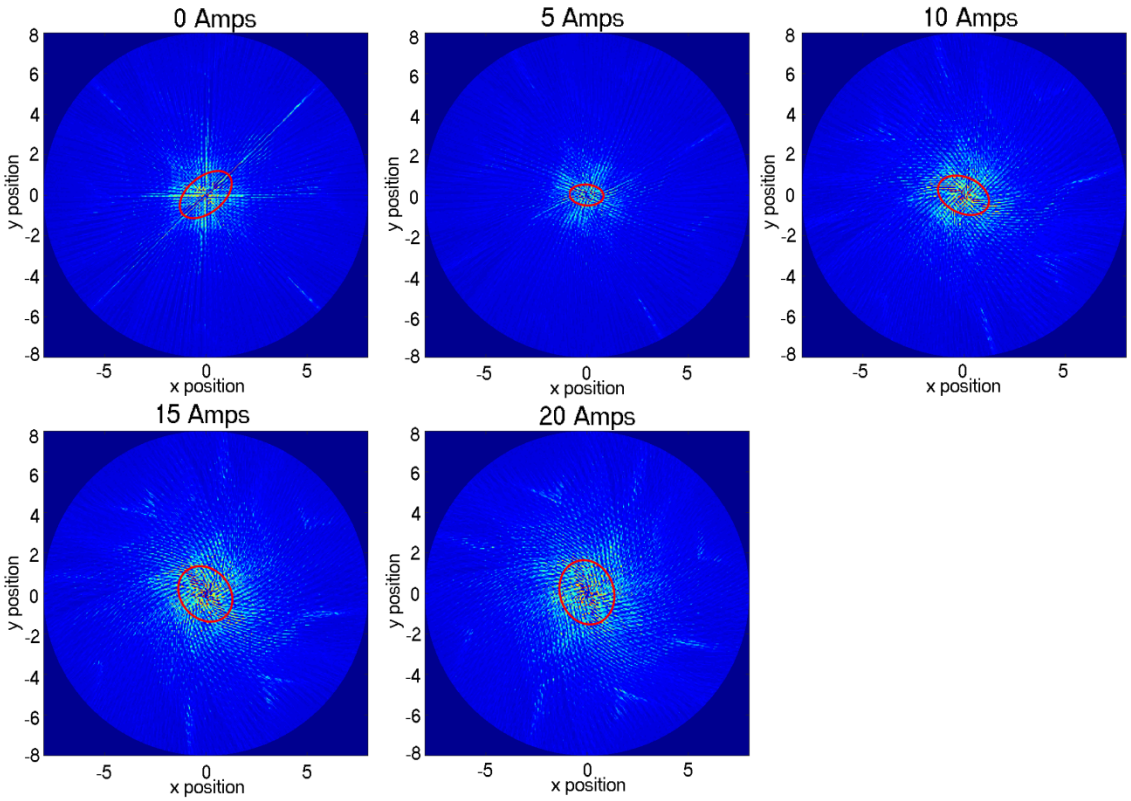

Figure 68: Simulated data passed through the same analytic pipeline as the experimental data. Note that the 'screen' in this case is normal to the beam direction, which is different from the geometry of the screen in Figure 5. This is because CST does not support non-normal beam monitors at his point.

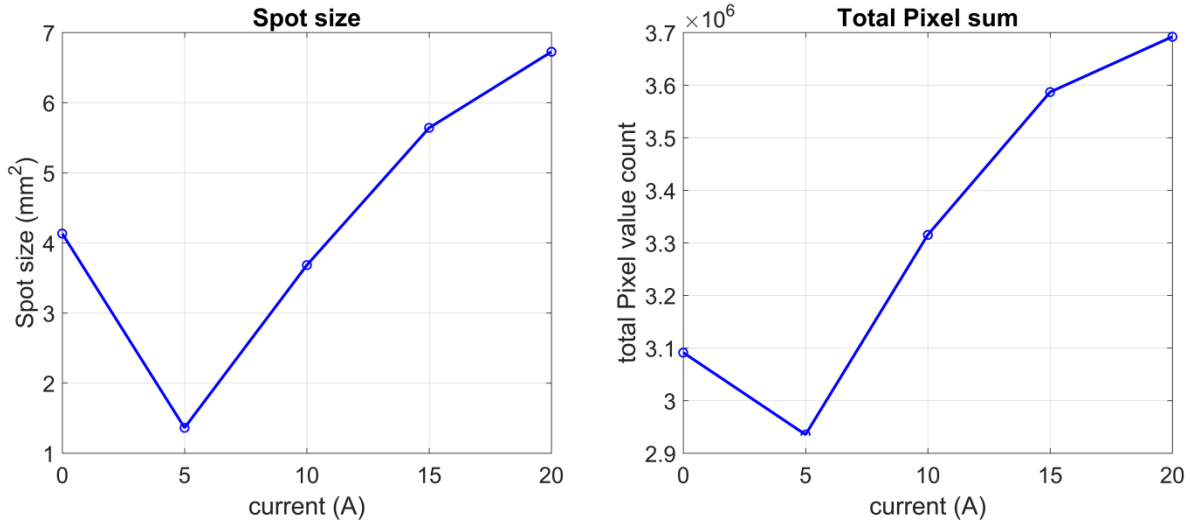

Figure 69: Simulated spot size and total pixel count versus coil current, as based on the simulated data

4. Discussion

In this work, an experimental beam line has been designed, constructed and installed at SLAC national laboratory. This project is ongoing. To date, the first beam on has been

achieved, and preliminary data under magnetic fields collected. The next beam on is planned for October 2016. The data collected and presented thus far is only preliminary, and as such so are any resultant conclusions.

The aims of this project are to (1) experimentally demonstrate that an RF gun can operate with minimal current loss in axial magnetic fields and (2) test how accurately computational electrodynamics codes predict the behavior of Megavoltage electron beams in magnetic fields. Because we do not currently have the torroid current monitor data, the first aim cannot be addressed at this stage. Regarding the second aim; only moderate agreement has thus far been obtained between the experimental and simulated data. There are a number of uncertainties which could limit the agreement between the two sets of data. For the experimental beam, a major source of uncertainty is mechanical misalignment (the effects of which can be clearly seen in the off center experimental images (Figure 5). A better understanding of the alignment issues could be obtained by taking measurement using both viewports and screens (Figure 1). As yet, it has not been possible to use the first screen as we did not have a camera available. Even if a camera is available, it is not clear whether it will ultimately be possible to use the first screen since it may cause vacuum trips (it is very close to the gun where one of the vacuum pumps is, and phosphor is known to cause vacuum issues under irradiation. Another source of uncertainty is the power levels of the gun. We have inferred the power levels from the forwards and backward power monitors just before the ceramic window which separates the gun from the SF6 filled RF waveguide. From this and from previous data [18] we inferred the energy distribution of electrons leaving the gun, but this is not equivalent to proper spectroscopic data.

For the simulated data, the major uncertainty is the resolution of the data, which will be improved in the future by running simulations with more particles and a more densely discretized cathode surface. In addition, the simulated gun model does not include the side coupling or input RF port, which will alter the modeled particle distribution exiting the gun.

Despite these shortcomings, there is also some agreement between the two sets of data. The most important is simply that axial magnetic fields have an observable impact on the radial beam intensity distribution at a given location. This was predicted in our previous work, and as discussed this could have serious implications on the x-ray target lifetime. In addition, there is general agreement between the two sets of data that the spot size increases as the magnetic fields are increased, although the simulated data predicts a minimum at 5 A which was not seen experimentally. One of the goals of this work was to quantify agreement between simulated and experimental data. However, given that the experimental work has only just begun, we feel that it would be premature to attempt to make any conclusive statement on this. Future experiments will be carried out to complete data collection, and hopefully to collect more data at higher magnetic field strengths.

5. Conclusion

An experimental beam line incorporating a thermionic RF electron gun has been designed, constructed, and installed at SLAC national laboratory, with the goal of benchmarking and

verifying the in-silico predictions made in our previous work. Thus far, some early data showing the electron beam impinging on a YAG screen in axial magnetic fields has been collected. Only moderate agreement with the simulated results has been obtained at this stage, however both data sources clearly show that the presence of axial magnetic fields can have a substantial impact on electron spot sizes. Both the in-silico modeling and experimental data collection are ongoing, and the results presented here are preliminary.

References

1. Keall, P.J., M. Barton, and S. Crozier. *The Australian Magnetic Resonance Imaging–Linac Program*. in *Seminars in radiation oncology*. 2014. Elsevier.
2. Fallone, B.G. *The rotating biplanar linac–magnetic resonance imaging system*. in *Seminars in radiation oncology*. 2014. Elsevier.
3. Lagendijk, J.J., B.W. Raaymakers, and M. van Vulpen. *The magnetic resonance imaging–linac system*. in *Seminars in radiation oncology*. 2014. Elsevier.
4. Mutic, S. and J.F. Dempsey. *The ViewRay system: magnetic resonance–guided and controlled radiotherapy*. in *Seminars in radiation oncology*. 2014. Elsevier.
5. Mutic, S., D. Low, T. Chmielewski, G. Fought, M. Hernandez, I. Kawrakow, A. Sharma, S. Shvartsman, and J. Dempsey, *TU-H-BRA-08: The Design and Characteristics of a Novel Compact Linac-Based MRI Guided Radiation Therapy (MR-IGRT) System.* Medical Physics, 2016. **43**(6): p. 3770-3770.
6. Constantin, D.E., R. Fahrig, and P.J. Keall, *A study of the effect of in-line and perpendicular magnetic fields on beam characteristics of electron guns in medical linear accelerators.* Medical physics, 2011. **38**(7): p. 4174-4185.
7. St .Aubin, J., D. Santos, S. Steciw, and B. Fallone, *Effect of longitudinal magnetic fields on a simulated in-line 6 MV linac.* Medical physics, 2010. **37**(9): p. 4916-4923.
8. St Aubin, J., S. Steciw, and B. Fallone, *Effect of transverse magnetic fields on a simulated in-line 6 MV linac.* Physics in medicine and biology, 2010. **55**(16): p. 4861.
9. Whelan, B., S. Gierman, L. Holloway, J. Schmerge, P. Keall, and R. Fahrig, *A novel electron accelerator for MRI-Linac radiotherapy.* Medical physics, 2016. **43**(3): p. 1285-1294.
10. Constantin, D.E., L. Holloway, P.J. Keall, and R. Fahrig, *A novel electron gun for inline MRI-linac configurations.* Medical physics, 2014. **41**(2): p. 022301.
11. St. Aubin, J., S. Steciw, and B. Fallone, *The design of a simulated in-line side-coupled 6 MV linear accelerator waveguide.* Medical physics, 2010. **37**(2): p. 466-476.
12. St. Aubin, J., S. Steciw, C. Kirkby, and B. Fallone, *An integrated 6 MV linear accelerator model from electron gun to dose in a water tank.* Medical physics, 2010. **37**(5): p. 2279-2288.
13. Baillie, D., J. St. Aubin, B. Fallone, and S. Steciw, *FEM design and simulation of a short, 10 MV, S-band Linac with Monte Carlo dose simulations.* Medical physics, 2015. **42**(4): p. 2044-2053.
14. Baillie, D., J.S. Aubin, B. Fallone, and S. Steciw, *Feasibility of producing a short, high energy s-band linear accelerator using a klystron power source.* Medical physics, 2013. **40**(4): p. 041713.
15. Schmerge, J.F., D.A. Reis, M. Hernandez, D.D. Meyerhofer, R.H. Miller, D.T. Palmer, J.N. Weaver, H. Winick, and A.D. Yeremian. *SLAC RF photocathode gun test facility*. in *Photonics West'97*. 1997. International Society for Optics and Photonics.
16. Moszyński, M., T. Ludziejewski, D. Wolski, W. Klamra, and L. Norlin, *Properties of the YAG: Ce scintillator.* Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 1994. **345**(3): p. 461-467.
17. Konrad, G.T., *High power RF klystrons for linear accelerators*. 1984.
18. Borland, M., *A high-brightness thermionic microwave electron gun*. 1991, to the Department of Applied Physics.Stanford University.
19. Karzmark, C., C.S. Nunan, and E. Tanabe, *Medical electron accelerators*. 1993: McGraw-Hill, Incorporated, Health Professions Division.
20. Kumar, V., *Understanding the focusing of charged particle beams in a solenoid magnetic field.* American Journal of Physics, 2009. **77**(8): p. 737-741.

Chapter 7
Patient reported outcomes of slow, single arc rotation

Patient reported outcomes of slow, single arc rotation: Do we need rotating gantries?

Brendan Whelan[1, 2, 3] brendan.whelan@sydney.edu.au

Miriam Welgampola[4]

Leigh McGarvie[4]

Kuldeep Makhija[1]

Robin M Turner[5]

Lois Holloway[2]

Ilana Feain[1]

Michael Jackson[6]

Michael Barton[2,]

Paul Keall[1,2]

1. *Radiation Physics Laboratory, University of Sydney, Sydney (NSW), 2006, Australia*
2. *Ingham Institute for Applied Medical Research UNSW, Liverpool (NSW), 2170, Australia*
3. *Central Clinical School, University of Sydney, Sydney (NSW), 2006, Australia*
4. *Institute of Clinical Neurosciences, Royal Prince Alfred Hospital, Sydney (NSW), 2006, Australia*
5. *School of Public Health and Community Medicine, University of New South Wales, Sydney (NSW), 1466, Australia*
6. *Prince of Wales Clinical School, University of New South Wales, Sydney (NSW), 2031, Australia*

Abstract

Purpose: Slow patient rotation in radiotherapy could greatly simplify radiation therapy delivery, with particularly important ramifications for fixed beam treatment with protons, heavy ion, MRI-Linacs, and low cost linear accelerators. Patient tolerance is often cited as a barrier to widespread implementation to patient rotation, however no quantitative data addressing this issue was found in the literature. In this study, patient reported experiences of slow, single arc rotation in both upright (sitting) and lying orientations were collected.

Methods: Fifteen patients previously or currently being treated for cancer were slowly (~2 rpm) rotated in both upright and lying orientations using an existing medical device normally used to assess balance disorders, the Epley Omniax. Patients were rotated a full 360 degrees in increments of 45 degrees. The rotation was paused for 30 seconds at each 45 degree increment to simulate beam delivery; in total 8 beams were simulated. Patients were rotated in both an upright and lying position in the same session. Response was monitored via validated psychometric questionnaires for claustrophobia, anxiety and motion sickness. The Wilcoxon signed rank test was used to test for significant differences in anxiety and motion sickness before, during and after the study.

Results: No significant differences in anxiety or motion sickness were found between before and after the study, or upright and lying rotation ($p > 0.05$). The median percentage scores for anxiety and motions sickness immediately following the study were both 0. In general, anxiety and motion sickness scores were low throughout the trial. All patients except one completed the study without interruption.

Conclusions: Slow, single arc rotation in both upright and lying orientations was well tolerated in this study. These results support the need for further studies into the clinical implementation of patient rotation, which has the potential to have a major impact on the practice and cost of radiotherapy.

1. Introduction

Radiation therapy is a pillar of modern cancer treatment, with approximately 50% of all cancer patients indicated for at least one course of radiation therapy as part of their treatment [1, 2]. The fundamental goal of radiation therapy is to deliver a prescribed radiation dose to the tumour, whilst simultaneously minimising the dose to surrounding healthy tissues, thereby minimising treatment side effects. For external beam radiation therapy techniques, the most obvious and effective means of achieving this is to utilise multiple beam angles intersecting at the tumour, which requires relative rotation between the beam and the patient. This relative rotation is typically achieved in modern radiation therapy by rotating a complex and heavy beam forming apparatus around a patient positioned at the rotation centre. However, much simpler gantry engineering designs would be possible with the alternative approach of slowly rotating the patient within a fixed beam line. Simplified gantry design would be of particular benefit for emerging treatments such as proton therapy, heavy ion therapy, MRI-Linac therapy, and medical applications of synchrotron radiation [3], where rotating gantries are either far more complex than their conventional counterparts or (in the case of synchrotron radiation) completely impractical (Figure 1). Fixed beam line systems could also enable the construction of more conventional X-ray systems [4] at much lower cost to help address the growing shortfall of treatment units in low and middle income countries, which is estimated to be 22 000 by the year 2035 [5]. As well as this, for some patients upright treatments can be beneficial in terms of both dosimetry and tolerance [6]. However, there are also challenges with patient rotation, most notably (1) rotation induced anatomic deformation and (2) patient tolerance to rotation - the focus of this work.

Assumed lack of patient tolerance has been a recurring theme in our discussions with radiation therapy stake holders (oncologists, therapists and physicists). Motion tolerance of radiation therapy patients has been previously assessed in the context of continuous couch motion for adaptive radiation therapy. Prior to the collection of empirical data for this application, similar sentiments regarding patient acceptance were expressed within the radiation oncology community [7]. However, study results did not show any evidence for a lack of tolerance, inertial anatomic motion, or changes in breathing patterns [8, 9]. Limited patient rotation is currently in clinical use in some specialist centres including The Francis Burr Proton Centre at Massachusetts General Hospital, and the Heavy Ion Medical Accelerator (HIMAC) in Chiba, Japan. In other areas of medicine, patient rotation is common place [10]. The purpose of this study was therefore to collect the first cancer patient reported outcomes of slow, single arc rotation in both the upright and lying positions which could be used to deliver radiation therapy without a rotating gantry.

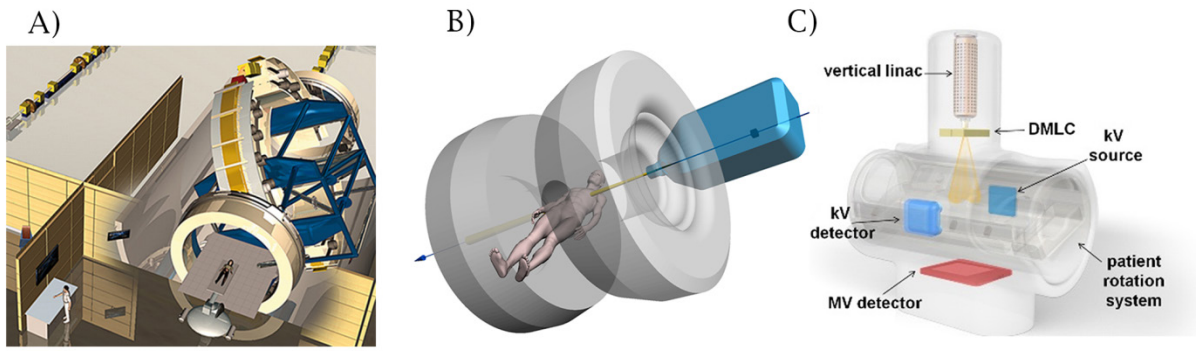

Figure 70: There are a number of treatment modalities which could derive considerable benefit from patient rotation. A) Heavy ion and proton therapy utilise much larger and heavier gantries than conventional therapy. Shown here is the rotating carbon ion gantry at the Heidelberg Ion Therapy Centre. B) Integrated MRI-Linac systems result in very complicated gantry mechanics. (Adapted with permission from Constantin et al. [11]) C) Patient rotation can help to address a devastating global short fall of linear accelerators. Show here is the Nano-X device. (Adapted with permission from Eslick et. al. [4])

2. Methods

2. 1. The Epley Omniax

In order to carry out slow, single arc patient rotation, this study utilised an existing device from the field of balance disorder therapy; the Epley Omniax (Figure 2). The Epley Omniax is most commonly used to treat benign paroxysmal positioning vertigo (BPV) [12, 13]. However, the fact that it can rotate 360 degrees around every axis also makes it is an ideal device to collect cancer patient reported outcomes of rotation.

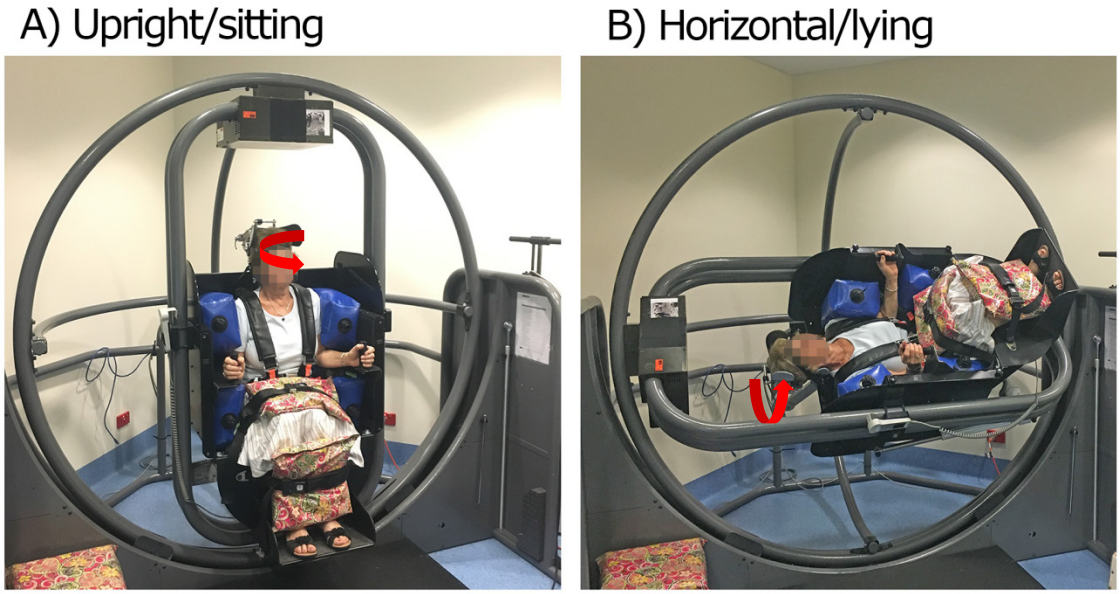

Figure 71: Patient set up in the Epley Omniax. (A) for upright rotation, (B) for lying rotation. The rotation axis is indicated by the curved arrow.

2. 2. Study design

This study was designed to assess the aspects of most concern regarding patient tolerance to rotation; (1) Would the restraint necessary to secure the patients trigger claustrophobia? (2) Would patients experience motion induced nausea? (3) To what extent would patients experience discomfort, breathing difficulties, and pain? As such, we adopted validated patient surveys to assess claustrophobia, motion sickness, and anxiety.

To quantify the claustrophobia of participants, the Claustrophobia Questionnaire (CLQ) was used. The CLQ tests for a fear of suffocation and a fear of restriction. It consists of 26 scenarios which respondents rate in terms of how anxious they would feel in each case on a scale of 0 (not at all anxious) to 4 (extremely anxious). The score is the sum of all answers [14]. To quantify the state (i.e. current) levels of anxiety among participants, the short form of the State Anxiety Test (STAI) was used. The short version was adopted to minimise the amount of paperwork participants were filling out, especially for questionnaires which would be repeatedly administered. This questionnaire has only 6 items which are scored between 1 and 4 and has been shown to correlate very well with the state section of the full form STAI [15, 16]. To derive the final score, 'pro anxiety' items (e.g. 'I feel tense') are added, and 'anti-anxiety' items ('I feel calm') are subtracted, resulting in a state anxiety score between 6 and 24. Finally, to measure motion sickness, we used the Fast Motion Sickness Survey (FMS). Again, part of the reason this survey was chosen was that it is simple and quick; patients are asked to rate their level of motion sickness or nausea on a scale from 0-20, where 0 is not sick at all, and 20 is extremely nauseous. The FMS has been shown to have high correlation to other more extensive motion sickness questionnaires [17]. For simplicity, all test scores were subsequently normalised such that the minimal possible score was 0 and the maximum possible was 100.

The study was powered based on the STAI anxiety test. Assuming a standard deviation of 20 based on previous studies of claustrophobia in MRI patients [18-21], we estimated that we would have a 90% power to detect differences in score of 16 or more on the normalised (i.e. score 0-100) STAI using a paired t-test with 15 patients. This sample size was considered sufficient for this early stage work, and so we planned to recruit 15 patients to this study. The inclusion criteria were as follows: (1) A diagnosis of abdominal, pelvic, or thoracic cancer, any stage, (2) 18 years of over, (3) Currently or previously being treated for cancer, (4) Any prior therapy allowed. The exclusion criteria were: (1) No pregnant women, (2) No mentally impaired patients or patients for whom obtaining informed consent would be difficult. Patients were recruited through a number of different cancer centres and support groups in Sydney. Ethics approval for this study was granted by the Sydney Local Health District Human Research Ethics Committee.

2. 3. Trial Procedure

At the beginning of each study, the following demographic information was collected: Age, gender, weight, primary tumour site, time since last cancer treatment (if not currently on treatment), and whether any medication was being taken which could cause nausea. Patients then completed the questionnaires for claustrophobia (CLQ), baseline anxiety (STAI), and

baseline motion sickness (FMS). Following this data collection, patients were securely set up in the Omniax using a variety of pillows, straps, a padded harness, and airbags, in an identical manner to patients undergoing clinical treatments on the Omniax for balance disorders (Figure 2). The support system is designed to hold the patient in place extremely securely and safely, and also to minimise patient discomfort by minimising any pressure points.

Figure 72: Study workflow. STAI is the anxiety questionnaire, and FMS is the motion sickness questionnaire.

Patients underwent slow, single arc rotation in two orientations; upright (sitting) and lying (Figure 2). Half the participants were rotated the upright position first, and half in the lying position first. Patients were rotated a full 360 degrees in increments of 45 degrees. The rotation was paused for 30 seconds at each 45 degree increment to simulate beam delivery. The 'very slow' setting was used for the angular velocity of the rotation, which is about 2 revolutions per minute. After completing each rotation arc, the questionnaires for motion sickness and anxiety were repeated (the claustrophobia questionnaire is designed as a predictive measure of susceptibility to feelings of claustrophobia and only needs to be administered once). After the first rotation, the participants verbally completed the questionnaires as they were still set up in the Omniax. After the second rotation, participants filled out written questionnaires. The entire study took typically took less than 40 minutes, with less than 15 minutes spent in the Omniax. In addition to the validated questionnaires, after the study we asked patients an open ended question requesting feedback about their experiences.

2. 4. Data analysis

To test for differences in either anxiety or motion sickness at any stage of the trial, test scores from before and after the rotation procedure, and between the upright and lying rotations, were compared using a two tailed paired Wilcoxon signed rank test. This test was used instead of the paired t-test because normality could not be assumed. To estimate the correlation between CLQ/STAI and CLQ/FMS Spearman's rank and Pearson's correlation coefficient was used. P-values less than 0.05 were considered statistically significant. All analyses were conducted in Matlab version 2014.

3. Results

3. 1. Study Cohort

The characteristics of the recruited patients are summarised in Table 1.

Table 10: Characteristics of the patient cohort recruited for the study.

Patient Number	Age (yrs)	Weight (kg)	Gender	Treatment Site	Time since last treatment (weeks)	Claustrophobia Score
1	NA	NA	M	Prostate	0	28
2	73	90	M	Prostate	8	44
3	65	82	M	Lymphoma	728	3
4	41	73	F	Breast	0	2
5	78	67	M	Prostate	0	53
6	65	62	F	Uterus	28	9
7	57	100	F	Uterus/ Breast	102	52
8	39	73	F	Hodgkin's Lymphoma	104	1
9	61	65	F	Breast	32	37
10	67	82	M	Prostate	96	21
11	78	70	M	Prostate	12	0
12	75	59	F	Breast	364	0
13	72	48	F	Breast	0	0
14	69	75	F	Endometrium	26	39
15	69	75	M	Prostate	3	0
Mean±STD	**65±12**	**73**	**7M, 8F**		**107±195**	**19±20**

3. 2. Patient Tolerance to slow single arc rotation

Figure 4 and Figure 5 show the percentage anxiety and motion sickness scores returned before the patient was set up in the Omniax, and after they had been rotated in both orientations. In general, the scores for both tests were very low, and we could not detect any significant difference in the scores using the signed Wilcoxon signed rank test. Similarly, Figure 6 and Figure 7 show the percentage anxiety and motion sickness between the upright and lying rotation. Again, no significant differences were detected. One patient did return the maximum anxiety score directly after the 'lying' rotation study (Figure 4, Figure 6, patient 7). Discussing their experience afterwards, they mentioned that although they felt very anxious during the rotation, having completed it they felt that they would be less anxious if they were to undertake the study again. In addition, patient 1 did not complete the study, terminating it during the 'lying' stage (interestingly, this patient still returned quite low anxiety scores). Both patients weighed over 90 kg, and both reported feeling insecure during the rotation. The Omniax harness is rated at up to 300 kg and so the participants were always safe, but this is an example of the importance of patient perceptions on patient tolerance. Very little motion sickness was reported at any point throughout the study, either via the FMS questionnaire or from discussions with participants afterwards. Overall, qualitative feedback from participants after the study was that whilst they didn't particularly enjoy being rotated, they also didn't think it was that bad. Many also made statements along the lines of "if it was necessary for my treatment, of course I'd tolerate it". Most patients found the lying rotation more challenging than the upright rotation, but found both were acceptable. This is reflected quantitatively in that no significant differences occurred in the STAI scores occurred between upright and lying. The correlation between the claustrophobia score (CLQ) and the maximum anxiety (STAI) and motion sickness score returned at any point during the study are shown in Figure 8. The anxiety scores show moderate correlation (Pearson's coefficient 0.6 ($p=0.02$), Spearman's coefficient .4 ($p=0.1$), whilst the motion sickness scores show minimal correlation (Pearson's coefficient: -0.2 ($p=0.4$), Spearman's coefficient -0.1 ($p=0.7$).

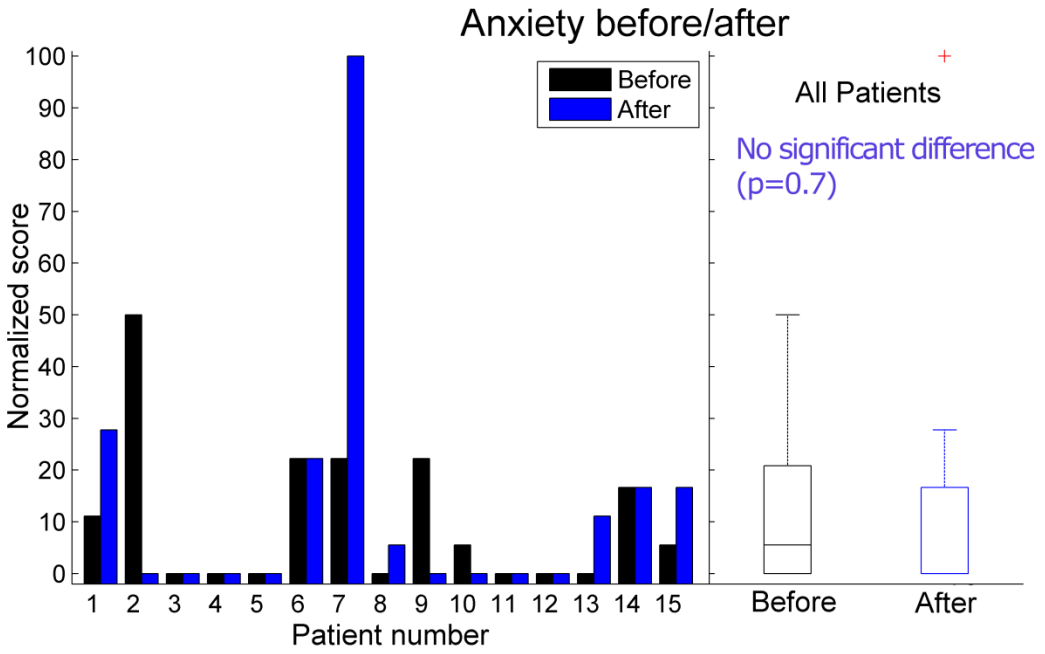

Figure 73: Percentage anxiety (STAI) scores returned before and after the study. 0 indicates the lowest possible score, and 100 indicates the maximum

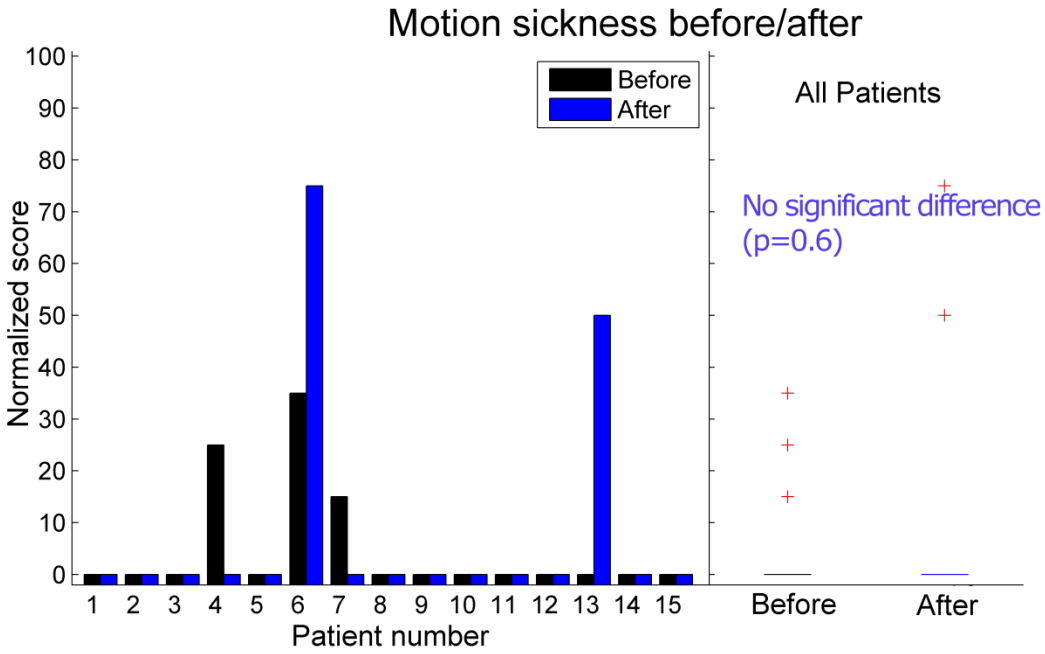

Figure 74: Percentage motion sickness (FMS) scores returned before and after the study. 0 indicates the lowest possible score, and 100 indicates the maximum

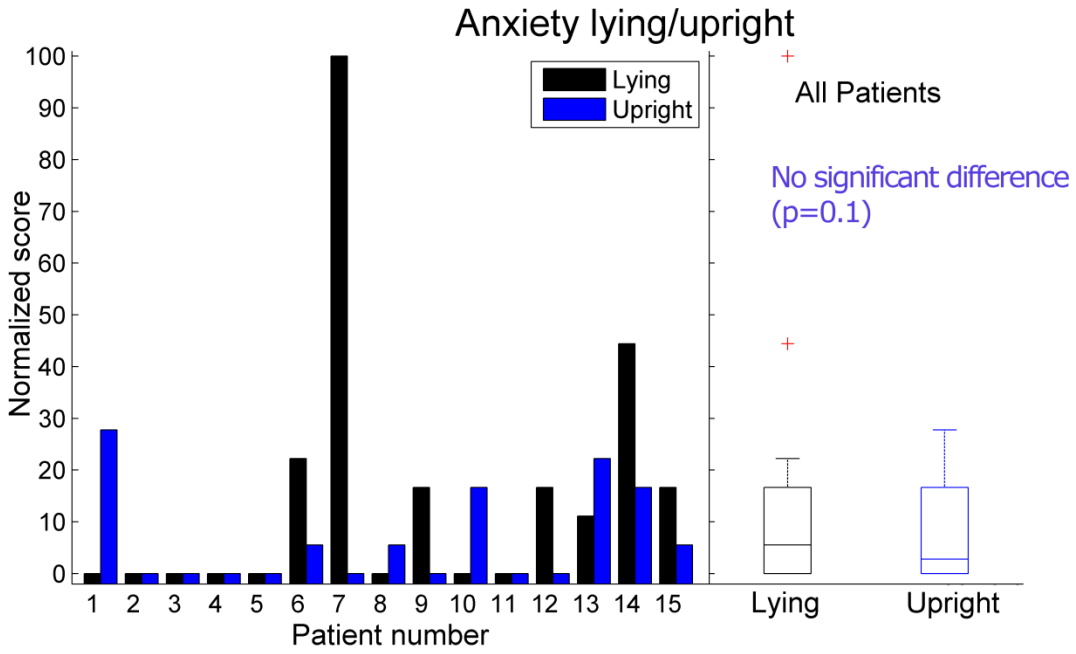

Figure 75: Percentage anxiety (STAI) scores returned for upright and lying rotation. 0 indicates the lowest possible score, and 100 indicates the maximum

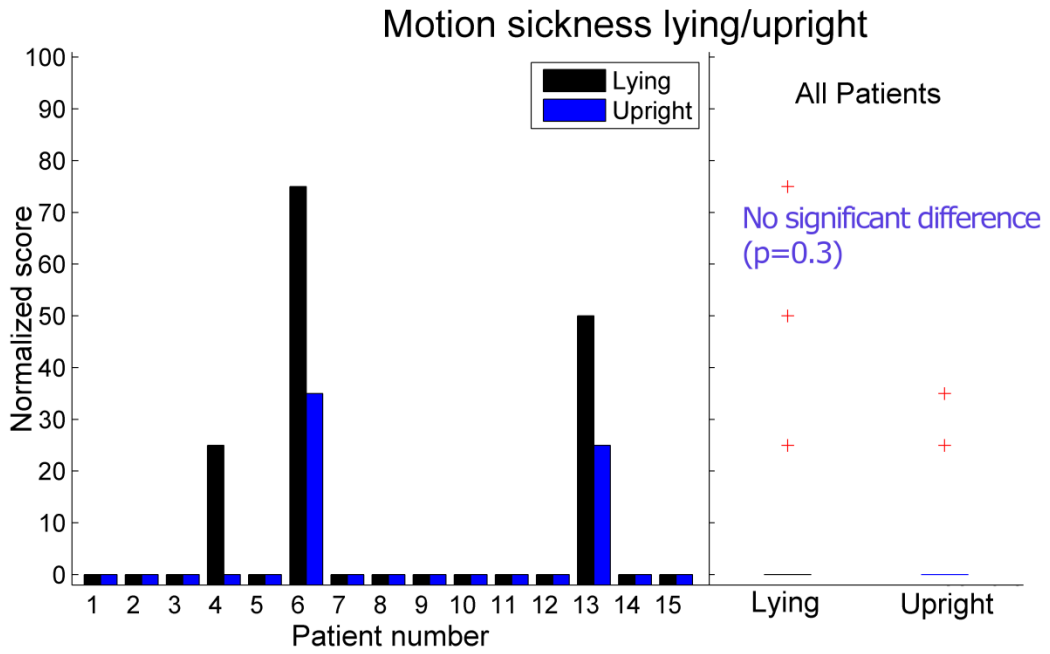

Figure 76: Percentage motion sickness (FMS) scores returned before and after the study. 0 indicates the lowest possible score, and 100 indicates the maximum

A) B)

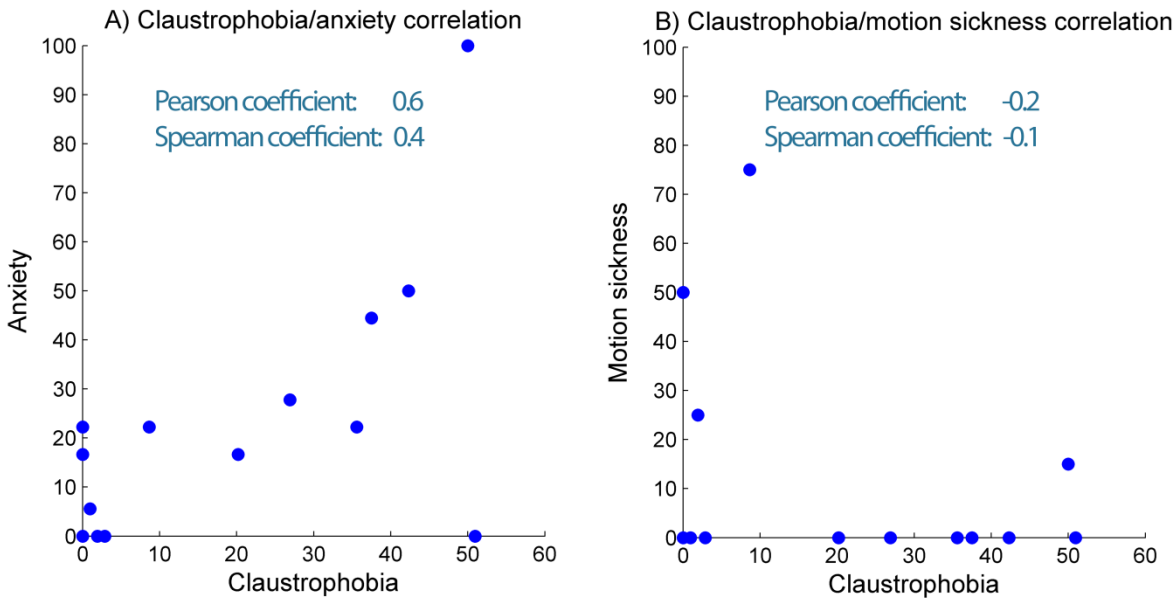

Figure 77: A) Correlation between claustrophobia score (CLQ) and the maximum anxiety (STAI) score returned at any point during the study. B) Correlation between claustrophobia score (CLQ) and the maximum motion sickness (FMS) score returned at any point during the study.

4. Discussion

In this study, we measured patient reported outcomes from 15 current or former cancer treatment patients who experienced slow, single arc rotation, as needed to deliver radiotherapy in fixed beam systems. Validated questionnaires were used to assess anxiety and motion sickness throughout the study, as well as a baseline claustrophobia score. The successful implementation of patient rotation in radiotherapy could allow the development of treatment machines without rotating gantries. Enabling patient rotation would have substantial ramifications for the cost and availability of emerging modalities such as proton therapy, heavy ion therapy, and MRI-Linac therapy. It could also enable the development of low cost linacs to service low and middle income countries, where radiotherapy utilisation falls far below established benchmarks [1, 5].

We did not detect any significant differences in anxiety or motion sickness throughout the study, and anxiety scores returned during this study were overall very low. We observed moderate correlation between claustrophobia and anxiety in this study, suggesting that minimising claustrophobia triggers may be an important means to minimise anxiety during rotation. Claustrophobia is most commonly understood to consist of two separate fears: a fear of suffocation and a fear of being trapped. As such, to minimise claustrophobia induced anxiety, it is important that patients feel that they could free themselves, and have a clear field of view around them.

There are some limitations to the results of this study. Firstly, only a small sample size was collected, and some patients had not been actively treated for cancer for some months. However, this work demonstrates that a cohort of patients who would find patient rotation an

acceptable treatment technique. Given the potential impact of patient rotation on the global practice of oncology, these results support the need for further studies on a larger and more representative patient population. If a similar methodology to the present study was adopted, we would also suggest some minor improvements to the study design. Firstly, it would have been appropriate to include a simple performance status questionnaire in the design such as ECOG Performance Status [22]. Secondly, we asked participants to circle a number between 1 and 4 to indicate their responses on the STAI test, where each number was associated with a given response, i.e. 'not at all'. Although this reflects the format of the questionnaire as published [15], we found that patients sometimes became confused about what the numbers meant when filling out the form. As such, we would suggest that all numbers are simply replaced by their relevant text (i.e. 'very much so' 'not at all'). Finally, the inclusion of a semi structured interview after the study would have likely provided more consistent and high quality qualitative data on patient experiences.

The concept of rotating patients for radiation therapy is not new - in fact in earlier incarnations of radiation therapy utilised positioning equipment to move the patient within a stationary beam was common place [23-25]. Historically, there are very good reasons why this approach was abandoned in favour of moving gantry techniques. Movement of the patient can introduce large geometric uncertainties, which if not properly compensated for can limit treatment efficacy. Furthermore, by moving to isocentric techniques facilitated by a moving gantry, the (then) difficult task of predicting dose distributions in the body was greatly simplified [26-28]. Today, there are a number of reasons why the role of patient positioning in radiation therapy should be revaluated. Vastly increased computational power and advanced imaging science mean that it is feasible to measure and adapt for patient motion – indeed, this is the exact premise upon which so much flagship medical physics research into adaptive radiotherapy is based. Whilst the problem of quantifying and adapting for rotation-induced anatomic deformation was not assessed in this study, in modern radiotherapy, it is becoming increasingly common to measure and adapt for patient motion, and there is already a great deal of research focused on increasing this capacity [29-33]. As such, it not unreasonable to hypothesise that anatomic deformation resulting from rotation could be adapted for. Based on the results presented in this work, future studies are planned to collect the data necessary to develop techniques to deliver high quality radiotherapy in the presence of such deformation.

5. Conclusion

Slow, single arc rotation in both the upright and lying orientations was well tolerated by the patients in this study. This work supports the need for further research into the clinical implementation of patient rotation, which could globally impact the practice of radiation oncology.

Acknowledgements

Brendan Whelan would like to acknowledge the Centre for Oncology Education and Research Translation (CONCERT) and Cancer Institute NSW for scholarship support. This study received funding support from the NHMRC (program grant APP1036075).

Conflicts of Interest

Authors Feain and Keall are stakeholders in Nano-X Pty. Ltd., a radiation therapy machine which incorporates patient rotation. Authors Feain and Keall are inventors on several pending patents patient rotation during radiotherapy.

References

1. Barton, Jacob, Shafiq, Wong, Thompson, Hanna, and Delaney, *Estimating the demand for radiotherapy from the evidence: a review of changes from 2003 to 2012.* Radiotherapy and Oncology, 2014. **112**(1): p. 140-144.
2. Delaney, Jacob, Featherstone, and Barton, *The role of radiotherapy in cancer treatment.* Cancer, 2005. **104**(6): p. 1129-1137.
3. Suortti and Thomlinson, *Medical applications of synchrotron radiation.* Physics in medicine and biology, 2003. **48**(13): p. R1.
4. Eslick and Keall, *The Nano-X Linear Accelerator A Compact and Economical Cancer Radiotherapy System Incorporating Patient Rotation.* Technology in cancer research & treatment, 2014: p. tcrt. 2012.500436.
5. Atun, Jaffray, Barton, Bray, Baumann, Vikram, Hanna, Knaul, Lievens, and Lui, *Expanding global access to radiotherapy.* The Lancet Oncology, 2015. **16**(10): p. 1153-1186.
6. Yang, Chu, Dong, and Court, *Advantages of simulating thoracic cancer patients in an upright position.* Practical radiation oncology, 2014. **4**(1): p. e53-e58.
7. Murphy. *Tracking moving organs in real time*. in *Seminars in radiation oncology*. 2004. Elsevier.
8. Sweeney, Arnold, Steixner, Nevinny-Stickel, and Lukas, *Compensating for tumor motion by a 6-degree-of-freedom treatment couch: Is patient tolerance an issue?* International Journal of Radiation Oncology* Biology* Physics, 2009. **74**(1): p. 168-171.
9. D'Souza, Malinowski, Van Liew, D'Souza, Asbury, McAvoy, Suntharalingam, and Regine, *Investigation of motion sickness and inertial stability on a moving couch for intra-fraction motion compensation.* Acta Oncologica, 2009. **48**(8): p. 1198-1203.
10. West, Hansen, Møller, Bloch, and Klokker, *Repositioning chairs in benign paroxysmal positional vertigo: implications and clinical outcome.* European Archives of Oto-Rhino-Laryngology, 2015: p. 1-8.
11. Constantin, Fahrig, and Keall, *A study of the effect of in-line and perpendicular magnetic fields on beam characteristics of electron guns in medical linear accelerators.* Medical physics, 2011. **38**(7): p. 4174-4185.
12. Vesticon Webpage. *The Epley Omniax*. 2016; Available from: http://vesticon.com/epley-omniax-system.
13. Lechner, Taylor, Todd, MacDougall, Yavor, Halmagyi, and Welgampola, *Causes and characteristics of horizontal positional nystagmus.* Journal of neurology, 2014. **261**(5): p. 1009-1017.
14. Radomsky, Rachman, Thordarson, McIsaac, and Teachman, *The claustrophobia questionnaire.* Journal of anxiety disorders, 2001. **15**(4): p. 287-297.
15. Marteau and Bekker, *The development of a six-item short-form of the state scale of the Spielberger State-Trait Anxiety Inventory (STAI).* Br J Clin Psychol, 1992. **31**(Pt 3): p. 301-306.
16. Tilton, *Review of the state-trait anxiety inventory (STAI).* News Notes, 2008. **48**(2): p. 1-3.
17. Keshavarz and Hecht, *Validating an efficient method to quantify motion sickness.* Human Factors: The Journal of the Human Factors and Ergonomics Society, 2011. **53**(4): p. 415-426.
18. Harris, Cumming, and Menzies, *Predicting anxiety in magnetic resonance imaging scans.* International journal of behavioral medicine, 2004. **11**(1): p. 1-7.
19. Katz, Wilson, and Frazer, *Anxiety and its determinants in patients undergoing magnetic resonance imaging.* Journal of behavior therapy and experimental psychiatry, 1994. **25**(2): p. 131-134.
20. Enders, Zimmermann, Rief, Martus, Klingebiel, Asbach, Klessen, Diederichs, Wagner, and Teichgräber, *Reduction of claustrophobia with short-bore versus open magnetic resonance imaging: a randomized controlled trial.* PloS one, 2011. **6**(8): p. e23494.
21. Bangard, Paszek, Berg, Eyl, Kessler, Lackner, and Gossmann, *MR imaging of claustrophobic patients in an open 1.0 T scanner: motion artifacts and patient acceptability compared with closed bore magnets.* European journal of radiology, 2007. **64**(1): p. 152-157.

22. Oken, Creech, Tormey, Horton, Davis, McFadden, and Carbone, *Toxicity and response criteria of the Eastern Cooperative Oncology Group.* American journal of clinical oncology, 1982. **5**(6): p. 649-656.
23. Wachsmann, *Various forms of moving field therapy and its possibilities.* Acta Radiologica, 1954. **41**(sup116): p. 524-531.
24. Smithers, *Rotation therapy.* Journal of the Faculty of Radiologists, 1954. **6**(2): p. 73-83.
25. Bagshaw, Kaplan, and Sagerman, *Linear Accelerator Supervoltage Radiotherapy: VII. Carcinoma of the Prostate 1.* Radiology, 1965. **85**(1): p. 121-129.
26. Braestrup and Mooney, *Physical Aspects of Rotating Telecobalt Equipment 1.* Radiology, 1955. **64**(1): p. 17-28.
27. Braestrup, Green, and Snarr, *Convergent Beam Irradiator: Apparatus for Convergent Beam Irradiation with Stationary or Moving Source 1.* Radiology, 1953. **61**(4): p. 614-624.
28. GARDNER, Bagshaw, PAGE, and Karzmark, *Tumor localization, dosimetry, simulation and treatment procedures in radiotherapy: the isocentric technique.* American Journal of Roentgenology, 1972. **114**(1): p. 163-171.
29. Ménard and van der Heide. *Introduction: Systems for Magnetic Resonance Image Guided Radiation Therapy*. in *Seminars in radiation oncology*. 2014. Elsevier.
30. Yan. *Adaptive radiotherapy: merging principle into clinical practice*. in *Seminars in radiation oncology*. 2010. Elsevier.
31. Keall, Colvill, O'Brien, Ng, Poulsen, Eade, Kneebone, and Booth, *The first clinical implementation of electromagnetic transponder-guided MLC tracking.* Medical physics, 2014. **41**(2): p. 020702.
32. Keall, Ng, O'Brien, Colvill, Huang, Poulsen, Fledelius, Juneja, Simpson, and Bell, *The first clinical treatment with kilovoltage intrafraction monitoring (KIM): A real-time image guidance method.* Medical physics, 2015. **42**(1): p. 354-358.
33. Shieh, Keall, Kuncic, Huang, and Feain, *Markerless tumor tracking using short kilovoltage imaging arcs for lung image-guided radiotherapy.* Physics in Medicine and Biology, 2015. **60**(24): p. 9437.

Chapter 8
MRI-Compatible patient rotation system

An MRI-Compatible patient rotation system – Design, construction, and first organ deformation results

Brendan Whelan[1, 2, 3] *brendan.whelan@sydney.edu.au, Room 495, Blackburn Building D06 | The University of Sydney | NSW | 2006*

Gary Liney[2]

Jason A. Dowling[4]

Robba Rai[2]

Lois Holloway[2]

Leigh McGarvie[3]

Ilana Feain[1, 3]

Megan Berry[2]

Michael Barton[2]

Rob Wilkins[5]

Paul Keall[1, 2, 3]

7. *Radiation Physics Laboratory, University of Sydney, Sydney (NSW), 2006, Australia*
8. *Ingham Institute for Applied Medical Research and Liverpool and Macarthur Cancer Therapy Centres, Liverpool (NSW), 2170, Australia*
9. *Central Clinical School, University of Sydney, Sydney (NSW), 2006, Australia*
10. *Australian e-Health Research Centre, CSIRO Computational Informatics, Australia*
11. *Biomech Engineering, Australia*

Abstract

Purpose: Conventionally in radiotherapy, a very heavy beam forming apparatus (gantry) is rotated around a patient. From a mechanical perspective, a more elegant approach is to rotate the patient within a stationary beam. Key obstacles to this approach are patient tolerance and anatomical deformation. Very little information on either aspect is available in the literature. The purpose of this work was therefore to design and test an MRI compatible patient rotation system such that the feasibility of a patient rotation workflow could be tested.

Methods: A patient rotation system (PRS) was designed to fit inside the bore of a 3T MRI scanner (Skyra, Siemens) such that 3D images could be acquired at different rotation angles. Once constructed, a pelvic imaging study was carried out on a healthy volunteer. T2 weighted MRI images were taken every 45° between 0° and 360°, (with 0° equivalent to supine). The prostate, bladder, and rectum were segmented using atlas-based auto contouring. The images from each angle were registered back to the 0° image in three steps: (1) Rigid registration was based on MRI visible markers on the couch. (2) Rigid registration based on the prostate contour (equivalent to a rigid shift to the prostate). (3) Non-rigid registration. The Dice similarity coefficient (DSC) and mean average surface distance (MASD) was calculated for each organ at each step.

Results: The PRS met all design constraints and was successfully integrated with the MRI scanner. Phantom images showed minimal difference in signal or noise with or without the PRS in the MRI scanner. For the MRI images the DSC (mean±standard deviation) over all angles in the prostate, rectum, and bladder was 0.60±0.11, 0.56±0.15, and 0.76±0.06 after rigid couch registration, 0.88±0.03, 0.81±0.08, and 0.86±0.03 after rigid prostate guided registration, and 0.85±0.03, 0.88±0.02, 0.87±0.02 after non-rigid registration.

Conclusions: An MRI compatible patient rotation system has been designed, constructed, and tested. A pelvic study was carried out on a healthy volunteer. Rigid registration based on the prostate contour yielded DSC overlap statistics in the prostate superior than interobserver contouring variability reported in the literature.

1. Introduction

In external beam radiation therapy, relative rotation between the patient and treatment beam is required in order to achieve conformal dose distributions. Relative rotation can be achieved by rotating the treatment head, the patient, or both. In modern radiotherapy it is most common to rotate the treatment head. However, there are a number situations where substantial benefit could be derived if instead the patient was rotated, resulting in greatly simplified design. Examples of applications where patient rotation should prove particularly beneficial include hadron therapy in which rotating gantries can weigh up to 600 tonnes, MRI-Linac therapy in which the presence of an MRI magnet considerably complicates gantry design [1], medical applications of synchrotron radiation [2, 3], and the development of low cost treatment delivery systems (it is estimated that 22 000 new Linacs are needed by the year 2035 to meet global demand) [4].

There are two distinct orientations in which patients could be rotated; upright or lying. Whilst both orientations address the applications outlined above, lying rotation has an important advantage compared to upright rotation: unlike upright rotation, lying rotation workflows are compatible with common medical imaging devices (CT and MRI) which are essential for radiotherapy treatment planning. As such, a lying rotation approach could utilise conventional imaging equipment which either already exists within hospitals, or is much easier to access or acquire second hand than specialized upright scanners. The obvious problem with lying rotation is that gravitationally induced anatomic deformation will occur due to rotation between the gravitational field and the patient. If not taken into account during planning and delivery, this deformation will introduce uncertainties into the treatment. In order to unlock the potential benefits of lying rotation for modern radiation therapy, it is necessary to quantify the mechanics of deformation and adopt strategies to sculpt high quality dose distributions in the presence of such deformation, which will almost certainly require in-room adaption. To date, there is very little published data which addresses this issue. To address this gap in scientific knowledge, the goal of this study was to design and investigate an MRI compatible patient rotation system (PRS) to be used in conjunction with our 3 Tesla MRI scanner. In this work, the design and construction of the PRS is reported, and the first images and quantification of organ deformation from a volunteer study in the prostate region are presented.

2. Methods

2.A. Patient Rotation System design

The PRS was designed to meet the following criteria: rotate 360 degrees lockable at 1 degree increments, safely and comfortably support patients up to 100 kg and 190 cm at any angle with no more than 3 mm of flexion, ensure patient is safely restrained such that no strain occurs on the neck or spine, minimise torque forces on the PRS (i.e. centre of gravity of

patient near centre of rotation), combined weight of patient and PRS must be less than 160 kg. In addition, the device had to be compatible with our 3 Tesla MRI scanner (Skyra, Siemens, Germany) which added two more design constraints; the device had to be small enough to fit within the 70 cm bore of the MRI scanner, and be entirely constructed from MRI compatible materials. It was also decided that the device would be manually positioned by operators in order to avoid electromagnetic interference between motors and the MRI scanner. Another consideration was whether to incorporate additional MR receive coils within the PRS, or acquire images using only the in-built body coil of the scanner. After some preliminary testing, it was decided to proceed without imaging coils, but to design such that they could be retrospectively added if necessary.

To ensure MRI compatibility and mechanical tolerance within the MRI bore, a commercial bioengineering company was engaged. To maximise the amount of space available for the patient rotation system, the standard MRI couch was fully withdrawn from the scanner bore. This also meant that the existing track structure of the scanner could be utilised inside the bore. Outside the bore, the standard couch was used to support the patient rotation device. A drawback of this approach is that the conventional couch must be manually positioned at the correct height before use. The maximum cylinder diameter which fits into the MRI bore with the couch withdrawn is 600 mm (Figure 1B) – however, when the required support structure of the PRS was taken into account, 567 mm was available for the outer cylinder of the PRS system, resulting in an internal 'bed' width of 438 mm. Based on the NHANES anthropometric database of US civilians, it was anticipated that the PRS would be of sufficient dimensions to accommodate around 60% of males and 80% of females [5].

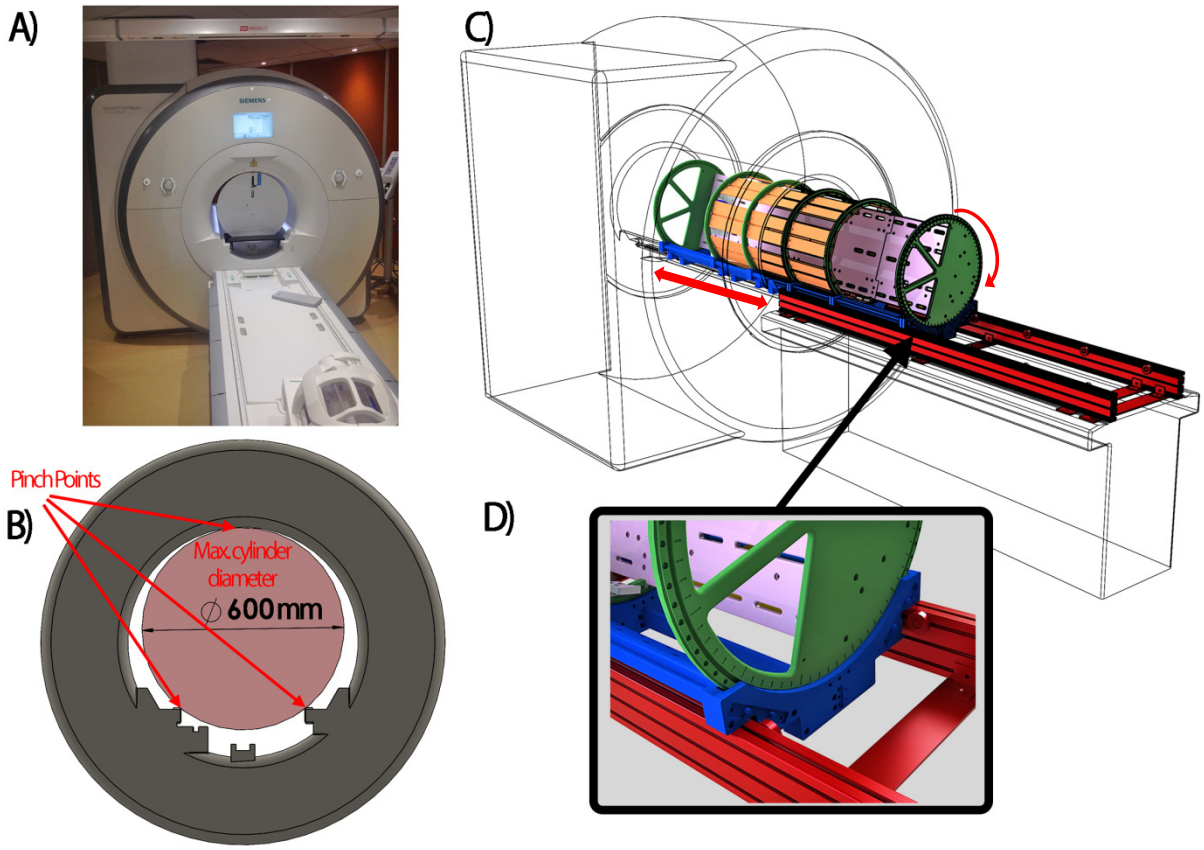

Figure 78: A) The Skyra MRI scanner with the couch retracted. In the bore there is a measuring device used to measure clearances. B) Diagram of the bore structure, showing the maximum cylinder which can fit inside, and where it is restricted (pinch points) C) Diagram of the PRS mounted onto the MRI scanner. The scanner superstructure is sketched in black and white, whilst the PRS is rendered in colour. For clarity, the patient support of airbags and strapping is not shown D) Close up of the rotation support structure, showing the external rails which mirror the rails built into the bore (red), the rotation support structure which moves axially in and out of the bore (blue), and the supporting structure which rotates (green).

To facilitate axial movement in and out of the scanner, the track structure of the bore was replicated into a track structure which could be mounted onto the conventional couch. The standard couch can be completely undocked from the scanner, meaning the rotating couch and track structure could be mounted on the standard couch external to the MRI room, before being wheeled back in and attached to the scanner. The section of the PRS which entered the bore of the magnet was mainly constructed from acrylic and acetal, with smaller amounts of polyester, polyurethane foam, PVC, vinyl, and nylon. All materials which were to enter the bore were tested for MRI compatibility prior to manufacture. The externally mounted rails which sit outside the bore were constructed from aluminium, and polymer and glass ball bearings were used throughout.

To ensure the device was as patient friendly as possible, a number of experts were consulted during the preliminary design stage, including radiation therapists, physicists, and oncologists. Importantly for the final design, clinical researchers from the field of balance disorder therapy (in which it is common to rotate patients) were consulted. On the basis of their advice and experiences, three inflatable airbags inserted between the bore of the PRS

and the patient torso were chosen as the primary patient support. This facilitates even loading across the patient surface, minimising pressure points. The PRS was also designed for patients to have as large field of visibility as possible while they were being imaged, to minimise claustrophobia.

2.B. Phantom study to assess MRI compatibility

To test the MRI computability of the device, a spherical 25 cm diameter oil filled QA phantom was imaged using a spin echo QA sequence (TE 30 ms, TR 1000 ms, pixel size 1.76x1.76 mm^2,slice thickness 5 mm, FOV 45 cm) both with and without the PRS. For the latter case, the phantom was positioned on the conventional couch. For both cases, only the body coil was used to acquire the images. As the MRI scanner is used for radiation therapy simulation, it is already equipped with a laser positioning bridge (Figure 1). These lasers were used to reproduce the phantom positioning as accurately as possible between scans. The signal to noise Ratio (SNR) was assessed as the ratio of the mean signal in the phantom to the standard deviation of the signal outside the phantom [6], using regions of interest which are shown in section 3.B.

2.C. Volunteer study

A healthy male volunteer was scanned using a modified version of the prostate protocol used for MRI simulation in our clinic. T2 weighted MRI images were taken at 45 degree increments over 360 degrees. A turbo spin weighted echo sequence was used (TE=96 ms, TR=6910 ms, turbo factor 15, pixel size 1.4x1.4 mm^2, slice thickness 3mm, FOV 45 cm). The bandwidth was 400 Hz/pixel and 3D distortion correction was applied. Each image took approximately 2 minutes to acquire and the total study time was 45 minutes. MRI visible markers were placed on the underside of the PRS couch (Figure 4) to aid in registration. To extract deformation information, the MRI images were processed as follows:

(1) Each non-0 degree MR data set was rigidly registered manually back to the 0 degree data using the couch markers using Slicer3D [7]. This essentially puts all data back into the same coordinate system. For each data set, the prostate, bladder, rectum, body and bones were automatically contoured using a previously published multi-atlas local weighted voting method [8].

(2) Structure guided registration based on the prostate contours was used to register each data set back to the 0 degree data. Only translation was used in this step to avoid rotation errors due to spherical shape of the prostate. This code was written in C++ using classes from the Insight Toolkit[27] (Metric: Mean Squared Error; Optimizer: Regular Step Gradient Descent (maximum step length 4.00, Minimum step length 0.01, 200 iterations); Transform: Translation). The transform from each prostate registration was then applied to globally the parent data (including contours.)

(3) To quantify residual deformation after rigid prostate guided registration, the output from each angled data set from step 2 was registered to the 0-degree data using demons based non-rigid registration [9]. The deformation field was applied to the contours output from step 3.

(4) The structures from each steps 2 and 3 were compared using the Dice Similarity Coefficient (DSC), and the Mean Absolute Surface Distance (MASD). The results from step 2 are representative of the geometric errors which would occur if a simple rigid shift adaption was used, as is common practice before treatment in conventional radiotherapy [10]. The deformation vector fields were analysed within each contour.

[27] https://itk.org/

3. Results

3.A. MRI compatible patient rotator

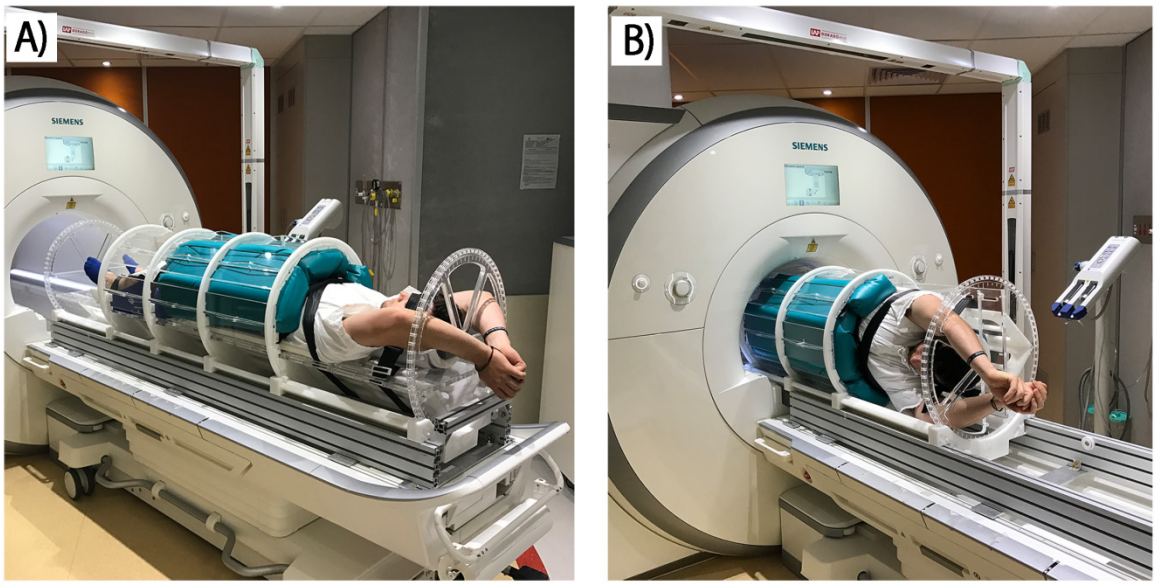

Figure 79: A volunteer in the patient rotation system (A) 0 degree position (B) 90 degree position, partially inside the scanner bore.

Figure 2 shows a volunteer in the PRS at two different angles. It can be seen that the PRS is mounted on the standard couch with the use of custom built rails. These rails mirror the rails within the bore of the MRI scanner, allowing the PRS to translate into the scanner. Also note the use of airbags to produce a distributed load across the patient surface. The delivered device met or exceeded all design constraints.

3.B. Assessment of MRI compatibility

Figure 3 shows a comparison of the phantom images taken with and without the PRS. It can be seen that only minimal differences exist between the two cases, indicating the MRI compatibility of the PRS. It can also be seen in Figure 3C that there was not perfect alignment of the phantom between the two scans – however, this is not entirely surprising as it is quite difficult to line up the object exactly, particularly on the MRI scanner couch which has a curved bed. When examining the pixel values in the two indicated circles we found there was a small but statistically significant ($p<0.05$) difference in both the signal and noise regions of the two images. The SNR was actually slightly higher with the PRS inside the couch, although we believe this is simply due to normal experimental variation.

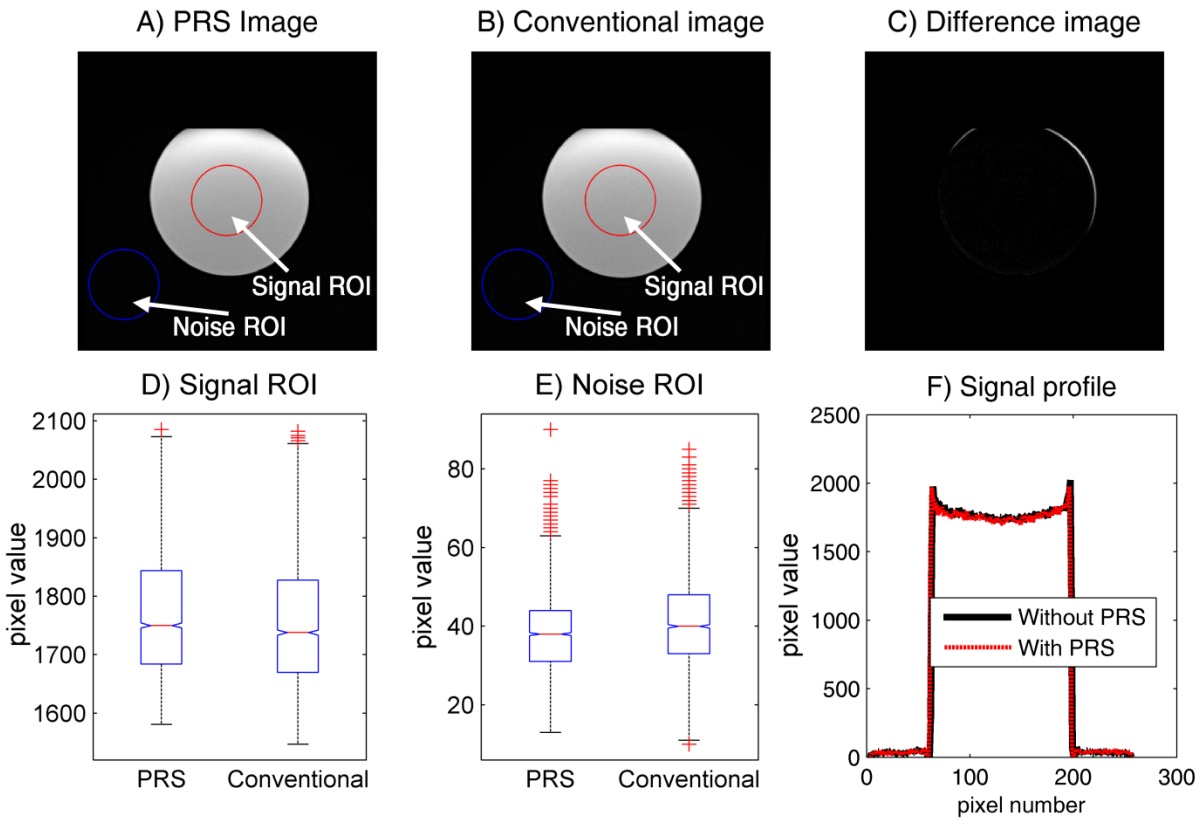

Figure 80: Assessment of image quality both with and without the presence of the PRS. A) Phantom with PRS. The two circles indicate the regions used to assess the SNR. B) Phantom image with the conventional couch. C) Difference of images A and B. D) Distribution of pixel values within the signal region of interest indicated on A and B. E) Distribution of pixel values within the noise region of interest F) Profile across the centre of images A and B.

3.C. Anatomical deformation

Figure 4 shows examples of the MRI images before and after image registration. The top row (a-d) shows the original images taken at each angle. The second row (e-g) shows the same images after they have been registered back to the 0 degree image (a) using prostate guided rigid registration. The third row (h-j) shows each image after being registered back to the zero degree image using non-rigid registration with the resultant vector field overlayed. The vector fields show non rigid deformation within the prostate region is quite low. It also appears that there is sometimes a torque-like force around the exterior part of the body, which could be a result of interplay between gravitational force and the patient support system. It can be seen that good image quality was maintained over all imaging angles, with clear anatomical detail. This implies that the volunteer was able to remain quite still throughout the study (approximately 45 minutes), as larger image artefacts would be expected otherwise. Unfortunately, truncation of the bladder did occur in some of images, which should be taken into consideration when interpreting the contour analysis presented below.

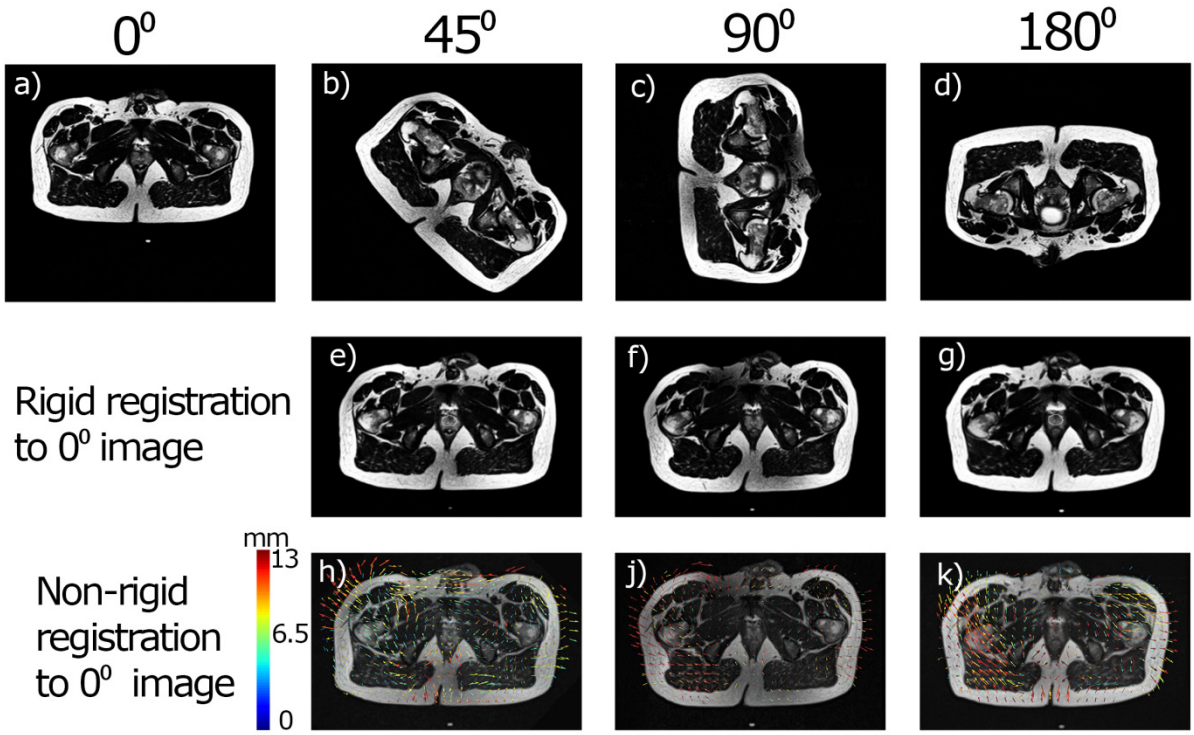

Figure 81: Representative MRI images. Top row shows the original images taken at a) 0, b) 45, c) 90, and d) 180 degrees. Second row shows each image after rigid contour guided registration back to the 0 degree image. Third row shows each image after non rigid registration back to the 0 degree image. The DVF from this step is overlayed on each image. (Note to reviewers: a high resolution version of this figure will be linked to from the caption)

The data in this study were registered in a three stage process (section 2.C). The first stage was a rigid registration based on an MRI visible marker placed on the couch, which essentially puts all data back into the same coordinate system. After this step, the mean DSC over all angles in prostate, rectum, and bladder was 0.60±0.11, 0.56±0.15, and 0.76±.06 respectively, and MASD differences were 3.76±1.22, 4.55±1.90, and 4.83±1.22 mm. The next stage was a contour guided rigid registration based on the prostate contour. After this step, the same results were: 0.88±.03, 0.81±0.03, and 0.86±0.03 (DSC), and 0.87±0.25, 1.75±0.77 and 2.81±1.12 mm (MASD). Finally, a non-rigid registration was applied. After this step, the same results were 0.85±0.03, 0.88±0.02, and 0.87±0.02 (DSC) and 1.14±0.34, 1.07±0.18, and 2.55±0.88 mm (MASD). Table 1 gives a summary of DSC values for prostate, rectum and bladder over all angles for each registration strategy. Note that both the DSC overlap and MASD offset of the prostate are slightly higher after non-rigid registration as this step is based on the entire image rather than just the prostate.

Table 11: DSC values for prostate, rectum, and bladder for the three registration strategies. Couch registration refers to the rigid rotation alignment based on the couch marker (i.e. this step simply puts all data back into the same coordinate system). Rigid prostate registration refers to the contour guided rigid registration step. The bottom row shows the mean overlap across all angles.

Angle	**Couch registration**			**Rigid prostate registration**			**Non-rigid registration**		
	Prostate	Rectum	Bladder	Prostate	Rectum	Bladder	Prostate	Rectum	Bladder
45°	0.56	0.52	0.73	0.83	0.80	0.80	0.80	0.88	0.89
90°	0.42	0.26	0.66	0.83	0.70	0.80	0.82	0.89	0.89
135°	0.67	0.53	0.83	0.87	0.65	0.88	0.82	0.86	0.91
180°	0.62	0.66	0.81	0.92	0.86	0.89	0.86	0.89	0.84
225°	0.67	0.70	0.83	0.92	0.85	0.89	0.87	0.85	0.86
270°	0.49	0.42	0.75	0.86	0.86	0.88	0.82	0.85	0.86
315°	0.53	0.53	0.75	0.90	0.86	0.87	0.88	0.88	0.83
360°	0.80	0.83	0.84	0.90	0.88	0.88	0.90	0.90	0.87
MEAN ±STD	**0.60 ±0.11**	**0.56 ±0.15**	**0.76 ±0.06**	**0.88 ±0.03**	**0.81 ±0.08**	**0.86 ±0.03**	**0.85 ±0.03**	**0.88 ±0.02**	**0.87 ±0.02**

4. Discussion

In this work the design and construction of an MRI compatible Patient Rotation System was presented, which to the best of our knowledge is the first of its kind. Using this system, prostate images of a healthy male volunteer at imaging angles between 0° and 360° were acquired, and deformable image registration was used to quantify the resulting anatomical deformation.

The first registration performed was a rigid couch registration step, which simply rotates all data back into the same coordinate system. Based on the results from this step (mean prostate DSC overlap of 0.60±.11 and MASD offset 3.76±1.22 mm) it is clear that in order to achieve this accuracy in a treatment approach based on patient rotation, an adaptive radiotherapy technique would have to be adopted. The results from the second stage of the registration steps (prostate contour guided rigid registration) are representative of the geometric accuracy which could be achieved with a simple rigid alignment to the prostate before each beam. This kind of adaption is commonly used for pre-treatment alignment in current clinical practice [10], and as such it is reasonable to believe it could be incorporated into a patient rotation workflow. After this step, the mean DSC coefficient for the prostate was 0.88±.08, and mean

MASD differences 0.87±0.25 mm. To put these numbers in context, 0.88 DSC is higher than reported interobserver contouring variability on prostate [8, 11] and 0.87 mm is less than the levels of intrafraction motion which can commonly occur for prostate cancer patients treated with conventional techniques [12, 13]. The agreement for rectum and bladder was also quite high after the prostate guided rigid registration step (DSC>0.8, MASD<3 mm) however there is very little information in the literature regarding geometric uncertainty in organs at risk, and so these results are more difficult to contextually interpret. Based on the prostate data alone, these results suggest that even the simple adaption strategy of rigidly shifting beams to the prostate may be effective in adapting for the deformation observed in this study. More advanced adaptive strategies could also be applied to compensate for the anatomic deformation, which could bring the geometric accuracy of the therapy to the levels seen after the non-rigid registration step. One interesting approach to adaptive radiotherapy was recently proposed in ref. [14]. In this work, the plan was adaptively updated in response to the real time position of the organs, whilst also taking into account the dose which has already been delivered. Such an approach may prove ideally suited for a rotation based radiotherapy workflow. A limitation of the data in this study is that it comes from only one healthy volunteer. To properly test the efficacy of different adaption strategies and extend the results to a population level, more data from a more representative patient cohort needs to be obtained, and appropriate treatment planning studies carried out.

It is expected that different treatment sites will show different levels of gravitationally induced motion; for instance we would expect that the lung region would exhibit greater motion than the prostate, as the latter organ is much more constrained by surrounding structures. Different imaging and analysis approaches will also be required for different regions, especially where there is already large motion such as lung. It is also likely that the exact mechanism used to secure the patient has a bearing on the deformation results, especially around the outer parts of the body. This was not explicitly investigated in this study; however from the authors experience, the combined airbag and nylon strapping support structure does do a very good job of supporting the patient securely in place. A potential issue with the implantation of a patient rotation based workflow in radiation therapy is patient comfort and tolerance. The healthy volunteer investigated this study did not report any substantial discomfort during or after being imaged in the PRS. However, this was not the focus of the present study and was not quantitatively investigated. Quantitative data addressing these issues in a larger patient cohort is needed, and will be collected in future studies via validated questionnaires to assess anxiety, claustrophobia, and motion sickness, e.g. see refs. [15-17].

We initially had some concerns about the image quality which could be obtained using only the in-built ‘body’ coil of the MRI scanner, which is why the device was designed such that it could be retrofitted with additional receive coils if required. The images taken thus far have been of sufficiently high quality that this may not be necessary; however this may be a pathway to improved imaging performance in future versions of the system. Similarly, MRI imaging performance does not appear adversely affected by the presence of the PRS in the bore (section 3.B). Due to the fact that imaging often occurs close to the edge of the MRI

field of view, some distortion is anticipated around the periphery of the images. Based on previous work [18], it is expected that around the periphery of the images machine based distortion on the order of 2mm may be occurring, although when looking at the images no obvious distortion artefacts can be seen. Machine based distortion for the PRS will be quantified in the future using a similar approach as our previous work [18]. Other limitations of the current version of the MRI compatible PRS are the restrictions on patient size and weight, the fact that it must be manually positioned, and the difficulty that low performance status patients may have getting in and out of the system.

The potential benefits of patient rotation in radiotherapy are obvious: greatly simplified gantry engineering, more cost effective treatments, and better patient access to radiation therapy. As such, a number of proposals for patient rotation have been made previously (for example, refs [19-23]). Typically, these have focused on upright rotation – however, this creates difficulties in terms of patient treatment planning as most medical imaging equipment cannot image patients in an upright position. A lying rotation workflow could be made compatible with existing imaging machines, but introduces obvious difficulties into the treatment workflow in terms of anatomic deformation. As such, an intriguing research question is whether high quality radiotherapy can be planned and delivered in the presence of such deformation. The increasing capacity in modern radiotherapy to quantify and mitigate anatomic uncertainty through motion modelling, robust treatment planning [24], and in room adaptation, gives reason to believe that this may indeed be possible. In particular, first order strategies to monitor and adapt for tumour motion using electromagnetic transponders, kV x-rays, and MRI already exist [25-27]. As such, and given the potential benefits of lying rotation, we believe that the challenge of delivering high quality radiotherapy via lying rotation is worthy of further investigation. To do this, data on the mechanics of gravity induced organ deformation is needed, and this is precisely the data the MRI compatible PRS has been built to capture.

Conclusion

An MRI compatible patient rotation system has been designed, constructed, and tested. A prostate study was carried out on a healthy volunteer. Rigid registration based on the prostate contour yielded DSC overlap statistics in the prostate superior than reported intra observer contouring variability taken from the literature.

Acknowledgements

Brendan Whelan would like to acknowledge the Centre for Oncology Education and Research Translation (CONCERT) and Cancer Institute NSW for scholarship support. This study received funding support from the NHMRC (program grant APP1036075).

Conflict of interest

Authors Feain and Keall are shareholders and directors of Nano-X Pty. Ltd., a radiation therapy machine which incorporates patient rotation. Authors Feain and Keall are inventors on several pending patents patient rotation during radiotherapy.

References

1. Ménard and van der Heide. *Introduction: Systems for Magnetic Resonance Image Guided Radiation Therapy*. in *Seminars in radiation oncology*. 2014. Elsevier.
2. Blattmann, Gebbers, Bräuer-Krisch, Bravin, Le Duc, Burkard, Di Michiel, Djonov, Slatkin, and Stepanek, *Applications of synchrotron X-rays to radiotherapy.* Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 2005. **548**(1): p. 17-22.
3. Suortti and Thomlinson, *Medical applications of synchrotron radiation.* Physics in medicine and biology, 2003. **48**(13): p. R1.
4. Atun, Jaffray, Barton, Bray, Baumann, Vikram, Hanna, Knaul, Lievens, and Lui, *Expanding global access to radiotherapy.* The Lancet Oncology, 2015. **16**(10): p. 1153-1186.
5. McDowell and Statistics, *Anthropometric reference data for children and adults: United States, 2003-2006.* 2008.
6. Kaufman, Kramer, Crooks, and Ortendahl, *Measuring signal-to-noise ratios in MR imaging.* Radiology, 1989. **173**(1): p. 265-267.
7. Fedorov, Beichel, Kalpathy-Cramer, Finet, Fillion-Robin, Pujol, Bauer, Jennings, Fennessy, and Sonka, *3D Slicer as an image computing platform for the Quantitative Imaging Network.* Magnetic resonance imaging, 2012. **30**(9): p. 1323-1341.
8. Dowling, Sun, Pichler, Rivest-Hénault, Ghose, Richardson, Wratten, Martin, Arm, and Best, *Automatic Substitute Computed Tomography Generation and Contouring for Magnetic Resonance Imaging (MRI)-Alone External Beam Radiation Therapy From Standard MRI Sequences.* International Journal of Radiation Oncology* Biology* Physics, 2015. **93**(5): p. 1144-1153.
9. Vercauteren, Pennec, Perchant, and Ayache, *Diffeomorphic demons: Efficient non-parametric image registration.* NeuroImage, 2009. **45**(1): p. S61-S72.
10. Simpson, Lawson, Nath, Rose, Mundt, and Mell, *A survey of image guided radiation therapy use in the United States.* Cancer, 2010. **116**(16): p. 3953-3960.
11. Gao, Wilkins, Eapen, Morash, Wassef, and Gerig, *A study of prostate delineation referenced against a gold standard created from the visible human data.* Radiotherapy and oncology, 2007. **85**(2): p. 239-246.
12. Ng, Booth, Poulsen, Fledelius, Worm, Eade, Hegi, Kneebone, Kuncic, and Keall, *Kilovoltage intrafraction monitoring for prostate intensity modulated arc therapy: first clinical results.* International Journal of Radiation Oncology* Biology* Physics, 2012. **84**(5): p. e655-e661.
13. Kupelian, Willoughby, Mahadevan, Djemil, Weinstein, Jani, Enke, Solberg, Flores, and Liu, *Multi-institutional clinical experience with the Calypso System in localization and continuous, real-time monitoring of the prostate gland during external radiotherapy.* International Journal of Radiation Oncology* Biology* Physics, 2007. **67**(4): p. 1088-1098.
14. Kontaxis, Bol, Lagendijk, and Raaymakers, *A new methodology for inter-and intrafraction plan adaptation for the MR-linac.* Physics in medicine and biology, 2015. **60**(19): p. 7485.

15. Marteau and Bekker, *The development of a six-item short-form of the state scale of the Spielberger State-Trait Anxiety Inventory (STAI).* Br J Clin Psychol, 1992. **31**(Pt 3): p. 301-306.
16. Radomsky, Rachman, Thordarson, McIsaac, and Teachman, *The claustrophobia questionnaire.* Journal of anxiety disorders, 2001. **15**(4): p. 287-297.
17. Keshavarz and Hecht, *Validating an efficient method to quantify motion sickness.* Human Factors: The Journal of the Human Factors and Ergonomics Society, 2011. **53**(4): p. 415-426.
18. Walker, Liney, Holloway, Dowling, Rivest-Henault, and Metcalfe, *Continuous table acquisition MRI for radiotherapy treatment planning: Distortion assessment with a new extended 3D volumetric phantom.* Medical physics, 2015. **42**(4): p. 1982-1991.
19. Boag and Hodt, *Adjustable chair for radiotherapy of head and neck cancer.* The British journal of radiology, 1971. **44**(520): p. 316-317.
20. Goga and Horton Jr, *Modular radiotherapy treatment chair and methods of treatment*. 1992, Google Patents.
21. Karzmark, Bagshaw, Huisman, and Lawson, *A versatile radiotherapy treatment chair.* The British journal of radiology, 1980. **53**(636): p. 1190-1194.
22. Klein, Wasserman, and Ermer, *Clinical introduction of a commercial treatment chair to facilitate thorax irradiation.* Medical Dosimetry, 1995. **20**(3): p. 171-176.
23. Yang, Chu, Dong, and Court, *Advantages of simulating thoracic cancer patients in an upright position.* Practical radiation oncology, 2014. **4**(1): p. e53-e58.
24. Bortfeld and Jeraj, *The physical basis and future of radiation therapy.* The British journal of radiology, 2014.
25. Keall, Colvill, O'Brien, Ng, Poulsen, Eade, Kneebone, and Booth, *The first clinical implementation of electromagnetic transponder-guided MLC tracking.* Medical physics, 2014. **41**(2): p. 020702.
26. McPartlin, Li, Kershaw, Heide, Kerkmeijer, Lawton, Mahmood, Pos, van As, and van Herk, *MRI-guided prostate adaptive radiotherapy–A systematic review.* Radiotherapy and Oncology, 2016.
27. Keall, Ng, Juneja, O'Brien, Huang, Colvill, Caillet, Simpson, Poulsen, and Kneebone, *Real-Time 3D Image Guidance Using a Standard LINAC: Measured Motion, Accuracy, and Precision of the First Prospective Clinical Trial of Kilovoltage Intrafraction Monitoring–Guided Gating for Prostate Cancer Radiation Therapy.* International Journal of Radiation Oncology* Biology* Physics, 2016. **94**(5): p. 1015-1021.

Chapter 9
Conclusions

The work undertaken in this thesis was part of the Australian MRI-Linac research program, which aims to develop an experimental MRI-Linac system. The work in this thesis can be split into two distinct sections, both related to the development of the MRI-Linac system. In the first section (Chapters 3-6) electromagnetic coupling effects between the MRI and Linac were explored, with a particular emphasis on electron transport in the linear accelerator.

In Chapter 3, an FEM model of a clinical gridded electron gun was developed based on accurate 3D laser scanning and electrical measurements. This is the first model of a clinical gridded electron gun to be presented in the literature, and the first electron gun to be characterised in magnetic fields which was actually based off the measured geometry of a gun which is in current clinical use. We found that this gun was over two times more sensitive to axial magnetic fields than the less realistic models which had previously been presented.

In Chapter 4, the principles of magnetic shielding were explored and applied to the 1.0 Tesla magnetic of the Australian MRI Linac program. We were able to design shields which would enable a standard linac to operate in an MRI environment for both the in-line and perpendicular MRI-Linac configurations, despite purposefully choosing a challenging source to Isocentre of distance of 1.5 meters. In addition, the shield introduced only moderate distortion into the MRI scanner; substantially below the upper limit of 300 PPM.

In Chapter 5, an alternative approach to operating an electron accelerator in an MRI environment was presented: to redesign the accelerator. We showed that using an RF-gun concept, an accelerator could be designed which could operate in a wide range of axial fields with minimal current loss. We also showed that axial fields could affect the spot size of the electron beam at the target, which potentially has serious ramifications for target heating and lifetime. We also quantified back bombardment power of the in-silico models for both the RF-gun accelerator and a standard accelerator, showing that one limitation with the RF-gun concept may be back bombardment power at the cathode.

In Chapter 6, an experimental beam line was developed in collaboration with the Stanford Linear accelerator facility (SLAC). The goal of this work was to validate and benchmark the in-silico findings of Chapter 5. At the time of writing, this work was ongoing; however the first beam on has been achieved, and preliminary results demonstrating the effect of magnetic fields on the electron spot size collected.

The second part of this thesis studied the problem of patient rotation in radiotherapy (Chapter 7 and chapter 8). Whilst patient rotation has potential to benefit a broad range of radiotherapy treatments, it could be of particular utility in MRI-Linac systems. This is because (1) the reduction in gantry engineering complexity is expected to be particularly beneficial for MRI linac systems, and (2) the MRI can be used to compensate for anatomic deformation in real time. In Chapter 7, a clinical study was undertaken showing that patient acceptance of slow, single arc rotation was quite high. This is the first quantitative data addressing this important issue and is supportive of the need for further studies in a larger cohort of patients.

The second major issue with patient rotation is the anatomical deformation which will result. There is very little data available to address this problem, however in Chapter 8, a unique medical device was constructed and tested which will allow the collection of high quality MRI data of patients under rotation. We also used this device to carry out a volunteer study and demonstrated the use of deformable image registration to quantify anatomic deformation.

Finally, appendix one details a data mining approach which was carried out as part of my engagement with the American Association of Medical Physicists (AAPM). This work details an algorithm which links NIH records back to AAPM membership records, allowing the creation of the AAPM research database. We used the resultant data to perform a top level analysis of AAPM funding dating back to 2002, showing that inflation adjusted funding of AAPM members is stagnating even whilst membership increases.

Appendix 1
AAPM research database

Development and testing of a database of NIH research funding of AAPM members

Brendan Whelan[1, 2]

Eduardo Moros[3]

Rebecca Fahrig[4]

James Deye[5]

Thomas Yi[6]

Michael Woodward[7]

Paul Keall[1]

Jeff Siewerdsen[6]

13. Radiation Physics Laboratory, University of Sydney, Sydney (NSW), 2006, Australia
14. Ingham Institute for Applied Medical Research, Liverpool (NSW), 2170, Australia
15. H. Lee Moffitt Cancer Center, Tampa (Florida), 33612, USA
16. Stanford University, Palo Alto (CA), 94305, USA
17. National Cancer Institute, Bethesda (MD), 20892, USA
18. I-Star Lab, John Hopkins University, Baltimore (MD), 21205, USA
19. American Association of Physicists in Medicine, Alexandria (VA), 22314, USA

Abstract

Purpose: To produce and maintain database of National Institutes of Health (NIH) funded research projects held by American Association of Physicists in Medicine (AAPM), and to make this data (hereafter referred to as the AAPM research database) available for the use of AAPM and its members.

Methods: NIH funded research dating back to 1985 is available for public download through the NIH exporter website, and AAPM membership information dating back to 2002 was supplied by the AAPM. To link these two sources of data together, a data mining algorithm which matched records based on investigator name, institution, and record date was developed in Matlab. The goal of this process is to link the NIH unique investigator number 'PI_ID' to AAPM investigators. Accuracy was assessed in terms of false positive and negative rates; false positives were manually estimated based on a random sample of 100 records, while false negatives were assessed by comparing against 99 member supplied PI_ID numbers. The AAPM research database was queried to produce an analysis of trends and demographics in research funding dating from 2002 to 2015.

Results: A total of 566 PI_ID numbers were matched to AAPM members. False positive and negative rates were respectively 4% (95% CI: 1-10%, N=100) and 10% (955 CI:5-18%, N=99). Analysis of the AAPM research database showed that in 2015 the NIH awarded $USD 110 to members of the AAPM. The data can be quickly updated each year as more data becomes available. The four NIH institutes most likely to fund AAPM members were the National Cancer Institute, National Institute of Biomedical Imaging and Bioengineering, and National Institute of Neurological Disorders and Stroke. These four institutes awarded over 85% of the total NIH research funding awarded to AAPM members in 2015; 1.1% of the combined budget of these four institutes was awarded to AAPM members. Members of the AAPM received less research funding in 2015 than in 2003 once inflation was adjusted for. In the same period of time, the number of AAPM members nearly doubled.

Conclusions: A database of NIH funded research awarded to AAPM members has been developed and tested using a novel data mining approach. The database will be maintained on annual basis, and is available to the AAPM, its committees, working groups, and members for download. A wide range of questions regarding financial and demographic funding trends can be addressed by this data.

1. Introduction and background

The American Association of Physicists in Medicine (AAPM) is the representative body of most medical physicists in the USA, and is the principal organization promoting the professional practice, educational activities, and research endeavors in the field of Medical Physics. At time of writing, the membership of the AAPM consisted of approximately 8350 members working across hospitals, universities, and industry. Broadly speaking, medical physicists perform two important roles in a modern health care system: firstly, to ensure the optimal and safe performance of a variety sophisticated therapeutic and diagnostic machines, systems and processes; secondly, to perform research ranging from basic science and technology development to the invention of new techniques, procedures, and translation of new scientific findings into clinical practice. With research presenting a critical aspect of the state and future of medical physics, and with the National Institutes of Health (NIH) presenting a major source of research support in these areas, this paper reports on the development of a database of NIH research funding of AAPM members.

Performing scientific research requires funding sufficient to support resources and personnel. In areas like healthcare, the innovation and improvements that result from research are in the national interest [1]; as such, many nations dedicate a substantial proportion of their annual budget to medical research [2]. In the United States, approximately 2.5% to 3% of the federal budget is dedicated to scientific and medical research. The main federal funding body of healthcare in the USA is the National Institute of Health (NIH) that attracted approximately 0.75% of the total US budget in 2016 (equating to approx. 32 billion USD) [3]. This money is allocated to medical researchers through a competitive peer review process [4]. Through this process, researchers (including many AAPM members) apply for and are awarded research funding of various amounts and with various degrees of success. Exactly how much a researcher receives, whether this proportion is changing with time, which research areas are funded each year, and which AAPM members are receiving funding are all questions that have been difficult to answer in the past, even though such information is in principle available in the public domain.

Obtaining this information is important to both to the AAPM, who must ensure that its members continue to perform clinically relevant and nationally competitive research, and to the NIH, which should aim to ensure the discipline that gave the world CT, MRI, ultrasound, radiation therapy, medical lasers, etc. (to name but a few [5]) continues to attract funding commensurate with the potential of this research to positively impact human health. The purpose of this paper is twofold; (1) present the development and testing of a research database of AAPM members awarded research funding by the NIH, and (2) present a top level analysis of trends in this funding. Although the NIH records dating back to 1985 are publicly available, it is not straightforward to extract the data associated with a given field or organization. This work presents a data mining approach which can extract funding records associated with an input list of AAMP members from data available from the NIH RePORTER tool [6]. This project was undertaken as part of the activities of the AAPM Working Group for the Development of a Research Database (WGDRD) with support of the AAPM Science Council, and the resultant data will be made available to AAPM members.

2. Methods

The process for creating the AAPM research database is shown in Figure 1. Each stage involved in creating the database is described in detail below.

2.1. Input Data

This work utilizes two databases: NIH RePORTER records and AAPM membership records. The AAPM data extends back to 2002, as well as containing a 'member since' field that extends back to 1962. Each AAPM member is identified by a unique 'STATUS_ID' number, and First, Last, and Middle names are stored in separate fields (note that throughout this manuscript, searchable fields are indicated by THIS FORMATING). All fields in the AAPM data are shown in Table 1.

NIH RePORTER funding records are publicly available for download and can be queried online [6, 7]. A brief outline of this data is shown in Table 2; a more detailed description of the contents of these records is available online [8].The data extends back to 1985, although funding amount is only recorded from 2000 onwards. Each principal investigator (PI) is assigned a seven or eight digit 'PI_ID' number, which typically remains constant throughout their career (occasional exceptions do exist as discussed in Section 4). PI names are stored in a single field, in the format 'last, first'. The last name is straightforward to interpret, however the first can include multiple names as well as middle initials. In this work we separate the 'first' string by spaces, and take the first space separated string to be the first name of that investigator. From 2006 onwards, the NIH introduced a multi PI application model. In these cases, each record in the NIH data can have more than one PI. Each PI name and number is stored in the same field separated by a semicolon. The PI who submits the grant and is responsible for communication with the NIH is assigned as 'contact'. In such a case, the PI name field might read: 'PIname1; PIname2 (contact); PIname3', where each name follows the 'Last, First' convention described above. The same structure is applied to the PI_ID field.

In order to extract records relating to research grants held by AAPM members from the NIH funding records, the unique AAPM identifier 'STATUS_ID' must be associated with the unique NIH identifier 'PI_ID'. However, there is no straightforward means to do this: therefore a data mining approach was developed and is described below.

Table 12: Structure of the AAPM data. The shaded entries represent the fields that are queried at least once in the present work. Note that the last two fields show the time period each record is valid (e.g. 2002-2006)

Field	Example
Status_ID (internal unique identifier)	4134
Title prefix	Dr
First name	Jane
Middle name/ initial	A
Last name	Doe
Title Suffix	Jnr.
Date of Birth	12-30-1969
Job title	Asst. Professor
Organization	Ivory Tower University
Department	Medical Physics
Highest academic degree	PhD
Gender	F
Email	Jane.doe@institute.edu
Member since	2001
Dues category	Full
Dues sub category	Full
Phone Number	555-555-5555
Extension?	None
Fax number	555-555-5556
Fellow	Y
Charter member	N
Active from	2002
Active till	2015

Table 13: Relevant fields from the NIH records. Fields that are queried at least once in this work are shaded grey. Other fields shown here were not queried, but may be useful in future work. Note that a complete description of NIH records can be found online [8].

Field	Example
Application_ID	8913171
Activity	R01
ARRA_Funded?	No
IC_Name	National Cancer Institute
NIH_spending_Cats	Cardiovascular; Heart Disease; Lung; Neurosciences; Rare Diseases;
Org_City	Dubbo
Org_Country	USA
Org_Department	Biomedical Engineering
Org_name	The University of Dubbo
Org_state	CA
PI_IDs	7682827
PI_Names	Doe, Jane
ProgramOfficerName	Steven Stevens
Suffix	S1
Direct_cost_amt	100
Indirect_cost_amt	100
Total_cost	200

2.2. Data mining algorithm

While multiple common fields exist in each data set, there is no unique field that would allow simple, reliable, and unambiguous information linkage. Therefore, the approach taken was to query multiple non-unique fields and consider the combined net evidence before making a decision. For example, querying according to LAST NAME and INSTITUTION provides far better discriminatory evidence than LAST NAME alone. Extraction of grants is a three stage process, broadly outlined in Figure 1 and described in detail below.

I. Initial filtration and processing of NIH data.

The first stage is an initial query of the NIH data, in which all grants from PIs with the same LAST NAME and FIRST INITIAL as an AAPM member (e.g. 'smith, j') are extracted, processed, and written to a separate excel spreadsheet. The purpose of this is threefold: (1) it cuts down the amount of data that must be subjected to more detailed analysis downstream; (2) it preprocesses the (occasionally messy) NIH data into a consistent and easy to read format to avoid downstream errors; and (3) it splits up multi PI grants into records that can be individually queried. This last step warrants further description. For each multi PI grant, the 'contact' PI_ID and PI_NAMES are extracted, and other PIs are deleted (see section 2A for a description of the data format). Note that this does not discard any useful data because we rely on the institution field downstream, and the NIH stores this data only for the 'contact' PI. Also, at this stage we are only trying to find PI_ID numbers. Once found, they are still used to query multi PI grants including non 'contact' PIs, as described in section III. At the end of this step, the data is reduced from ~900 MB to ~10 MB, and comprises ~222,000 records in the format: LAST NAME, FIRST NAME, PI_ID, FINANCIAL YEAR OF RECORD, PI INSTITUTE, and PROJECT NAME.

II. Associate AAPM members with their PI_ID

The next step is to apply a series of more detailed tests to the data from part I to associate AAPM members with their NIH PI_ID. To do this, we consider the net evidence provided by comparisons between the NIH record and the AAPM record, querying: FIRST NAME, FIRST INITIAL, LAST NAME, AND INSTITUTE. In cases for which the PI changes institution during the period of a grant project, both institutions are listed within the NIH record, and the total cost that year is the sum of the records from each institution. As well as this, the funding body from the NIH data is checked. Explanation of the tests used is given in Table 3, and the combined test results that trigger a match are described in Table 4. This process was undertaken for each year of NIH records from 1985 to 2015. For each year, the algorithm reads all recorded NIH funding, and AAPM members who were active within one year of the year being tested. For example, if the year 2000 was being tested, AAPM members active from 1999 to 2001 (inclusive) would be included in the testing process. Because the PI_ID does not (usually) change from year to year, this approach gives multiple opportunities to make a positive match. In any single year, errors may occur in either data set that can confound the test results; however, when multiple years are analyzed it can substantially reduce the impact of these errors.

The results from each year are stored in separate sheets of a Microsoft Excel workbook, along with the binary indicators for each of the test results, as outlined in Table 4. Based on these test results, the PI_ID associated with each record is either included or discarded. The inclusion criteria are shown in Table 4. After the final year, all results are analyzed and consolidated into a list of unique PI_ID numbers. When multiple matches for the same PI_ID number are found, the binary flags shown in Table 4 are updated, and the best case (determined by the binary flags) is kept. Also at this stage, a list of PI_ID corrections from previous iterations is read – known false positives are discounted, while previously identified 'manual review' numbers are un-flagged.

Two lists must be defined before running this part of the code. The first is a list of words not to include in the institute match test. The second is a list of 'danger names' (Table 5). This is a list of very common last names that were observed to cause a large number of false positives. If the last name is a 'danger name', then inclusion instance 2 (which allows a match if the first name does not match) is not used, and the entry is automatically flagged for manual review. Both lists have been developed based on repeated running of the code and identification of failure points; both can also be very quickly updated as required for future use.

Before continuing to the next step, the returned list of PI_IDs is subjected to manual QA. Firstly, any entries flagged for manual review are checked. Entries can be flagged for manual review based on several reasons: the entry was included based on criteria 3, the last name was a danger name, the PI_ID was matched to more than one person, or the investigator had multiple PI_IDs assigned to them. When errors are found in the PI_ID list, they are stored on a separate worksheet so they don't appear again, and deleted from the final list. Similarly, when entries that were flagged for manual review pass inspection, they are added to a 'un-flag' list so they are not flagged again in subsequent iterations. In addition to the review of flagged entries, additional QA is carried out on this list to estimate false positive and negative rates. This is described in section 2C.

III. Using the PI_ID list to extract AAPM grant records

Using the list of PI_IDs described above, the NIH data was again queried, and all grants from those PIs were extracted and stored in an Excel workbook, with grants from each stored on a separate sheet. Note that we now query the full NIH data set, and not the filtered intermediate data that was used to obtain the list of PI_IDs in step II. When a positive match is found, the NIH data is appended with additional data which is only available in the AAPM records: Date of Birth, Gender, and the unique AAPM identifier STATUS_ID, which allows each grant to be linked back to AAPM records. Grants are recorded by the expenditure each year, so for example, a grant running from 2001 to 2003 will have a record in each of those years. Multiyear projects can easily be consolidated if desired, since each NIH project has a unique ID, APPLICATION_ID.

Occasionally, a single project has multiple entries in the NIH database in one year. There are a number of reasons why this occurs: most commonly, the grant has been appended or supplemented, which means that some change was made to the project. In this work,

duplicate records are treated as follows: if funding information exists for both records, then records are combined and total funding updated. If funding information exists for one record but not for the other, the empty record is deleted. If funding information exists for neither, then one is deleted and one kept. In this way, both overall funding and number of grants awarded are as accurate as possible.

In the case of multi PI grants, two additional databases are created; one for the case where the 'contact' PI is an AAPM member, and another where the 'contact' PI is not an AAPM member but an AAPM member is listed elsewhere as an investigator. Unless otherwise stated, the results presented below are the combined data from single PI grants and multi-PI grants where the 'contact' investigator is an AAPM member. In other words, multi PI grants where the lead investigator is not an AAPM member are excluded from most of the analysis below (To get a sense of the amount of funding attributed to each case, see Figure 6).

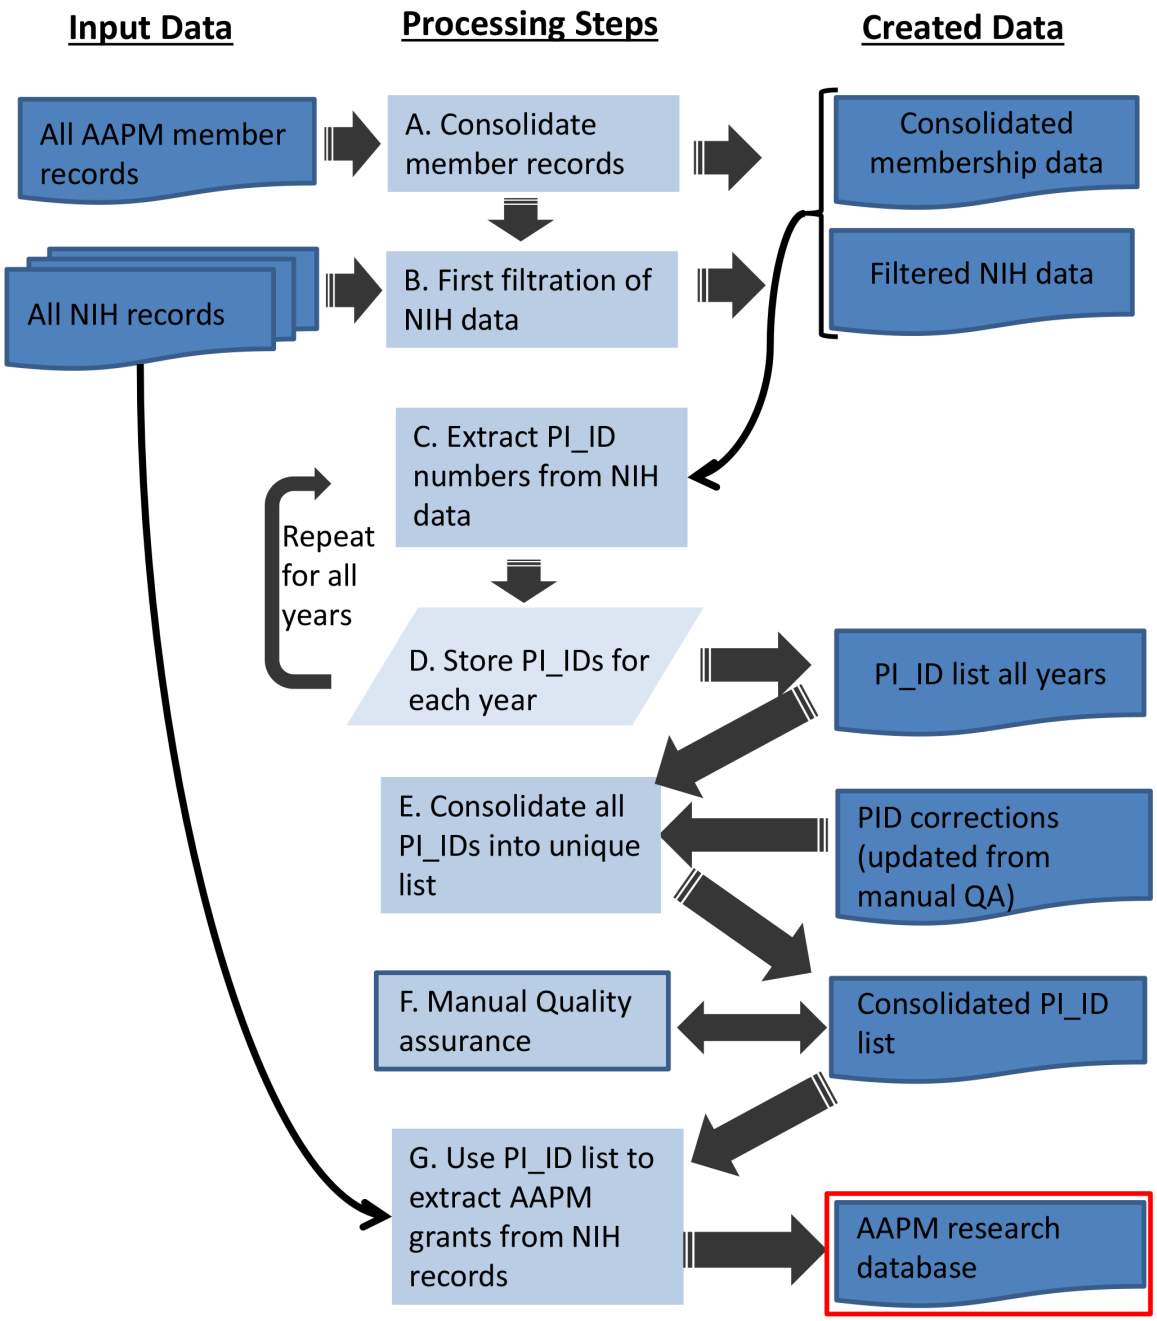

Figure 82: The basic process used to generate the AAPM research database. This process is described in detail in the methods section 2B

Table 14: Description of the tests used to decide whether each record should be matched to an AAPM member

Test	Pre processing		Criteria for positive result ('T')
	AAPM data	**NIH data**	
Last name	Last name already has its own field. Spaces, dashes etc. are removed from string.	Names are stored as Last, First Middle. Last name extracted, then spaces, dashes etc. removed from string.	Case insensitive string match
First name	First name already has its own field.	First name is taken as the first space separated string after the comma (see above)	Case insensitive string match
First initial	First character of first name.	First character of first name.	Case insensitive string match
Institute	1. Common institute abbreviation are 'unfolded' , e.g. 'UCLA' becomes 'UCLA University of California Los Angeles" 2. All common words such as 'university' removed from name Table 5) 3. Remaining string separated into words using the space character	1. Common institute abbreviation are 'unfolded' , e.g. 'UCLA' becomes 'UCLA University of California Los Angeles" 2. All common words such as 'university' removed from name (Table 5) 3. Remaining string separated into words using the space character	All words between the two data sets are compared using a lower case string match. Only one match between the two sets of words is required for a positive result.
Funding Body	NA	Funding body stored in the "IC_NAME" field of NIH data.	If funding body is one of: 'National Cancer Institute', 'National Heart, Lung and Blood Institute', 'National Institute of Biomedical Imaging and Bioengineering', 'National Institute of Neurological disorders and stroke.'

Table 15: The test results required to identify a positive match between an NIH record and an AAPM member. Three inclusion instances are currently supported; these could be extended.

Inclusion instance	Last Name	First Name	First Initial	Institute	Funding Body	Manual review
1	T	T	T	T	T	F
2	T	F	T	T	T	F
3	T	T	T	T	F	T

Table 16: The list of common words removed before institution matching and a list of last names that have been identified as producing a high number of false positives. For these names, inclusion instance 3 (Table 4) is excluded

Words removed from institution matching	university, univ, univ., of, the, institute, hospital, medical, center, college, school, cancer, therapy, centre, center, NaN, na, for, city, physics, program, national, health, inc., specialists, centers, ctr, ctr., oncology, royal, and, &, medicine, at, hospitals, clinic, inc, inc., state
Danger names	kim, lee, chen, zhang, huang, wang

2.3. False positives and false negatives

The above process is by no means infallible; some records will be missed (false negatives) while others will be incorrectly included (false positives). For the data to be useful, it is important to have a method for estimating the false positive and false negative rates. As both false positives and false negatives are binary quantities (true/ false), confidence intervals for the results were assessed using binomial statistics [9].

To assess the false positive rate, a list of 100 PI_IDs was randomly extracted with replacement from the total list of ~600, and each of these was entered into the 'Principal Investigator (PI) / Project Leader' field of the NIH reporter website with all available years selected [6]. Based on the returned grants, PI name, and institution, each record was examined manually to determine whether or not this PI_ID should legitimately be associated with an AAPM member. This process was repeated numerous times throughout the development of the algorithm to identify failure points and improve the algorithm. Each time this was performed, the false negatives were stored on a separate spreadsheet and excluded from the next iteration of the algorithm. Although it could be argued that such a process biases the final false positive results presented below, in the future less than 20 new PI_ID numbers are expected each year, and these can and should be quickly checked manually in the manner described above to ensure the ongoing fidelity of the database.

False negatives are somewhat harder to detect than false positives, as by definition one is dealing with data that has not been detected. Therefore, a secondary source of data is required against which comparisons can be made. Fortunately, we had at our disposal a list of 99 AAPM member supplied PI numbers, which were obtained as part of an alternative approach to obtain the NIH funding data. By assessing how many of these PI IDs our algorithm could detect we were able to obtain a measure of the false negative rate. This was also repeated throughout the algorithm development, and the results were used to fine tune the algorithm.

3. Results

3.1. Accuracy of PI_ID Data

The data mining algorithm identified 560 PI_ID numbers associated with the AAPM since 1985, of which 272 were active members at the time of writing. Manual testing of this list of numbers as described above resulted in a false positive rate of 4%. A binomial fit to this data

results in a 95 % confidence interval of 1-10 %, with N=100. Comparison with the member supplied PI_ID numbers yielded a false negative rate of 10 %, with a 95% confidence interval of 5-18 %, N=99. Note that after testing, the 10 false negatives were added to the list of PI_IDs used to generate the rest of the data in this report, while the 4 false positives identified were removed, such that 566 PI_IDs were used to generate the results reported below.

Note that while the exportable NIH data extends back to 1985, the online records only extend back to 2000 – also, funding amount is only available from 2000 onwards. As such, data prior to the year 2000 cannot be easily manually tested for false positives. Also, recall that the AAPM data only extends back to the year 2002, so data prior to this cannot easily be tested for false negatives. Therefore, the above results can only be considered valid for data after the year 2002. Although we present data before this, we caution that data prior to 2002 should be interpreted with the above information in mind.

3.2. Analysis of funding data

A top level analysis of NIH funding of members of the AAPM is presented in the following section. Firstly, some general points: Wherever a funding amount is shown, it is in USD, and no adjustment for inflation has been made unless otherwise stated. Where adjustment for inflation is made, it is to 2015 USD, based on USD using the Biomedical Research and Development Price Index [10]. Where box plots are used, the 'box' represents the 25th and 75th percentiles of the data, and points lying outside q3+1.5*(q3-q1) are classed as outliers, where 'q' represents data quartiles. Although no error bars are shown on the plots, they should be interpreted bearing in mind the false positive and negative rates outlined above (4% and 10%, respectively).

Figure 2 shows the yearly expenditure of grants held by AAPM members each year. In this plot, a grant running for three years would have one data point in each year, corresponding to the 'total cost' column of the NIH data in that year. The red line indicates the mean value. In 2015, the mean expenditure each year was $447k. In the year 2000 the mean expenditure each year was $328k, or $522k in 2015 dollars, suggesting this metric has for the most part kept pace with inflation.

Figure 3 shows the age of members receiving funding in each year; no strong trend is immediately apparent. We note that there do appear to be some errors in this member entered data. For instance some grants were apparently awarded to PIs at the age of 14. In the data shown, any grants whose PI age was less than 18 have been discarded on the assumption that this is a data entry artifact.

The number of grants awarded to members of the AAPM each year, as well as the number awarded to male and female members is shown in Figure 4. In 2015, 11% of grants held by AAPM members were held by female members. In the same year, 22% of all AAPM members were female, suggesting that males are twice as likely to hold research funding compared to females.

Figure 4 shows the total funding awarded to the AAPM, compared to the 'funding pool'. This is defined as the total budget of the top four funding agencies providing research grants to members of the AAPM over all years, outlined in Figure 7 and Table 6. These four funding agencies (NCI, NIBIB, NHLBI, and NINDS) represent approximately 81% of the total funding granted to members of the AAPM since such records were available. Some features of this graph can be explained by the political and economic climate at that time. In 2009 and 2010, a large spike appears in the total AAPM funds. This is due to the American Recovery and Reinvestment Act (ARRA) [11]. Note that the ARRA allocated funding is not included in the NIH budget, which explains why the NIH budget does not show the same trend during these years. In 2013, there is a sharp dip in the NIH funding pool. This is due to a budget sequestration [12]. At the time of writing, the NIH had not published its 2015 budget online and so this data point is missing.

To give a better insight into grant funds held by members of the AAPM as a percentage of the available funds, Figure 6 shows the proportion of the 'funding pool' that was allocated to members of the AAPM in each year. The horizontal red line shows the proportion at the latest available data point, in 2014. Again, the apparent spike in 2009 and 2010 is due to the ARRA funding, which does not appear in the NIH budgets.

Figure 7 shows the amount to which the various national institutes comprising the NIH have awarded grant funding to members of the AAPM. This graph shows funding data across all available years; the top four funding agencies have been the National Cancer Institute, the National Institute of Biomedical Engineering and Biomedical Engineering, the National Heart, Lung and Blood Institute, and the National Institute of neurological disorders and stroke. The budget for these four funding bodies is used to define the 'funding pool' in Figure 5 Figure 6. For information on the abbreviations and for the amount of funding awarded to AAPM members, see Table 6.

Figure 8 shows the different grant types which AAPM members have been awarded; we see that funding has most prevalently been awarded through the R01 project grant mechanism. An explanation of the different funding types can be found in ref. [13]

Finally, Figure 9 shows the inflation adjusted 'funds per member' in year 2000 USD, compared to the membership of the AAPM. It can be seen that, while the cumulative research funds awarded to members of the AAPM has stayed approximately constant, membership has steadily grown, meaning less research is funding is available on a 'per member' basis over time.

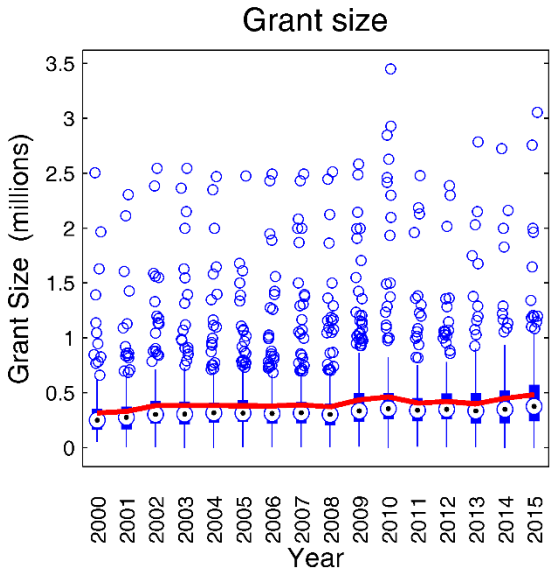

Figure 83: Box plots of AAPM member grant amount each year. Note that two grants which were larger than 3.5 million are not shown in this figure, occurring in 2014 ($7.86 mill) and 2015 ($7 mill).

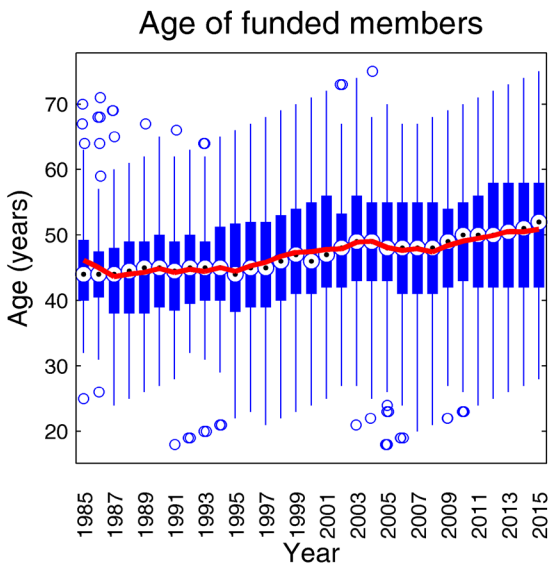

Figure 84: Box plots of the age of NIH funded members versus time. The line plot represents the mean. Note that some of the outliers in this data are a likely a result of incorrectly entered member data.

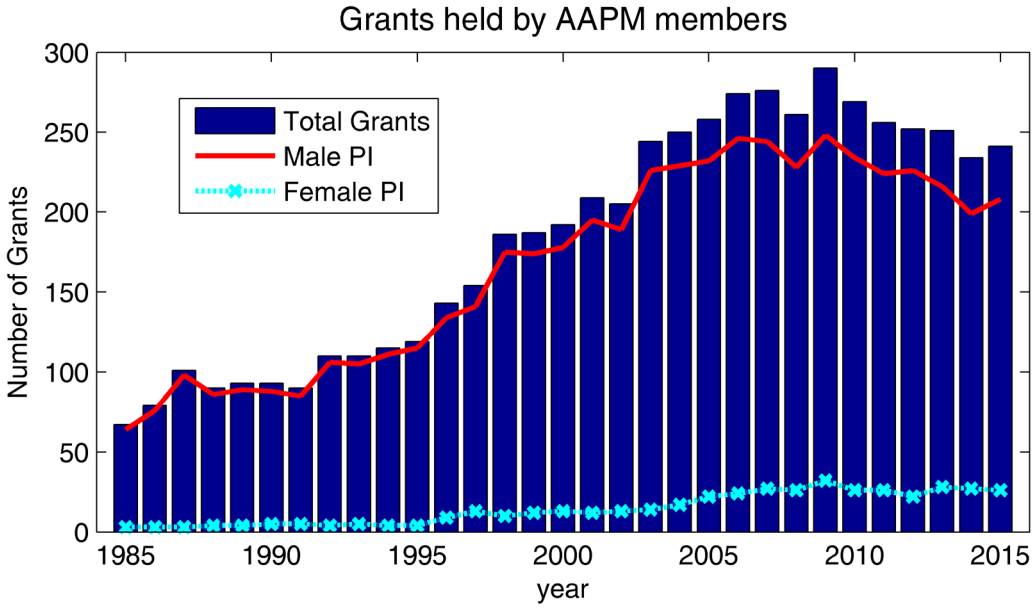

Figure 85: The total number of grants awarded to AAPM members each year, and total grants awarded to male and female PIs.

7

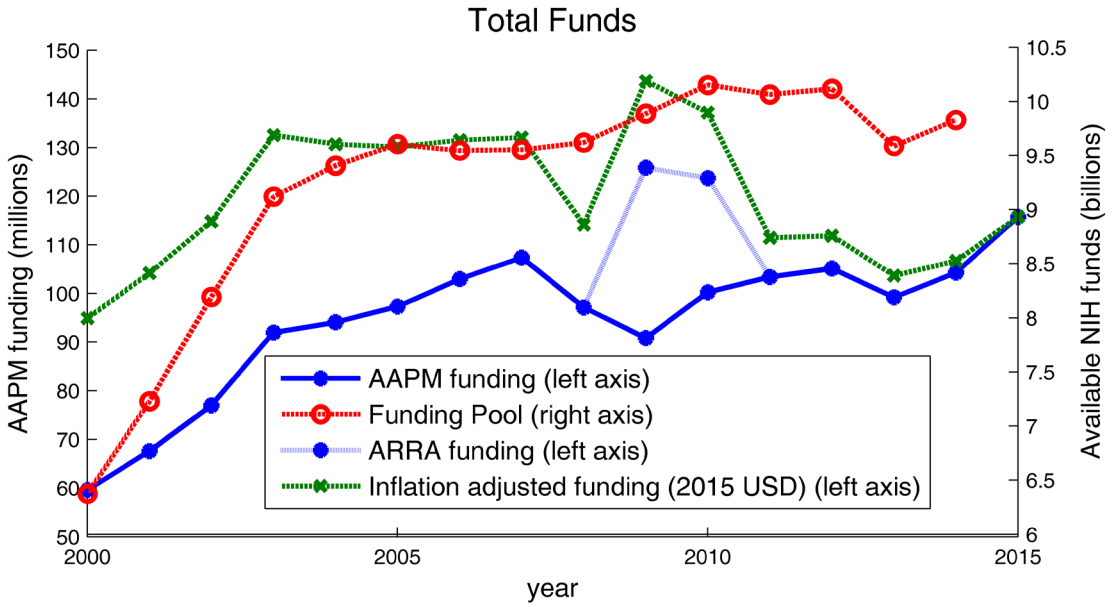

Figure 86: Total funds allocated to members of the AAPM, and 'funding pool' (defined as the total budget of the top four funding agencies for AAPM members – the National Cancer Institute, the National Heart, Lung, and Blood Institute, National Institute of Biomedical Engineering and Biomedical Imaging, and the National Institute of Neurological Diseases and Stroke (Figure 7). Also shown is inflation adjusted funding in 2015 dollars.

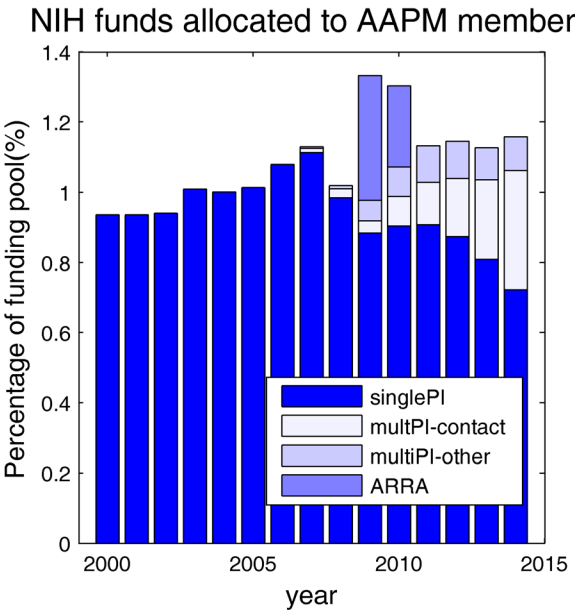

Figure 87: Percentage of the available funding pool awarded to members of the AAPM each year. From 2006 multi PI grants are also plotted, separated by cases where the AAPM member was the 'contact' PI and when they were not ('multiPI-other')

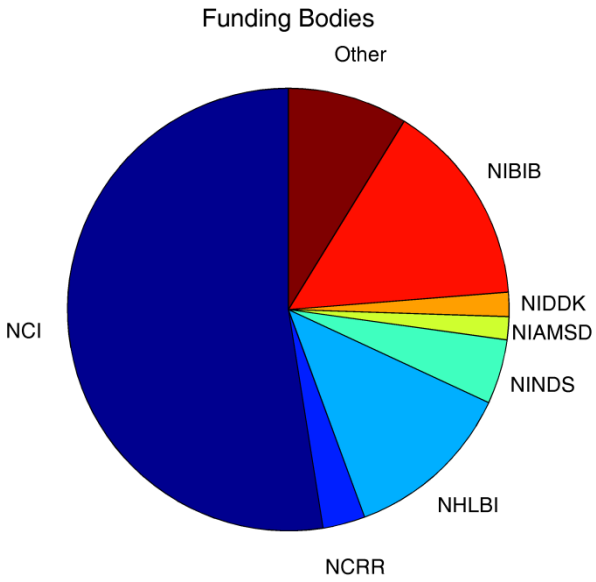

Figure 88: Relative amount that different national institutes have awarded grants to members of the AAPM. For abbreviations, see Table 6: The extent to which different organizations awarded funding to AAPM

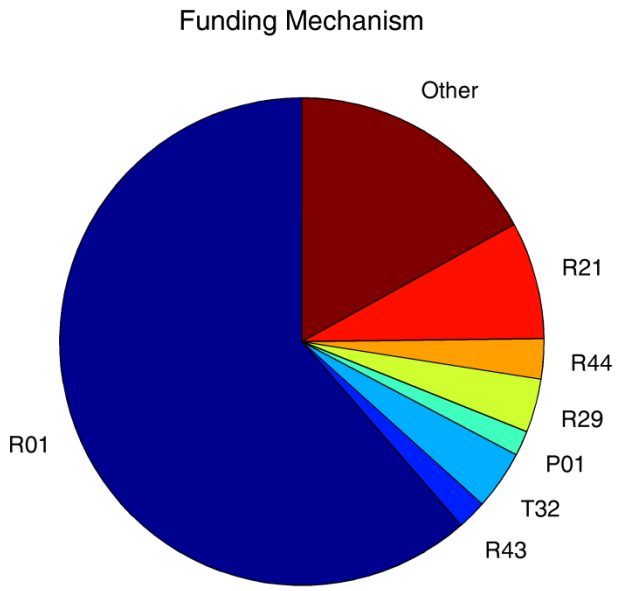

Figure 89: Funding mechanisms (grant types) for AAPM members. For explanation of abbreviations, see ref 12.

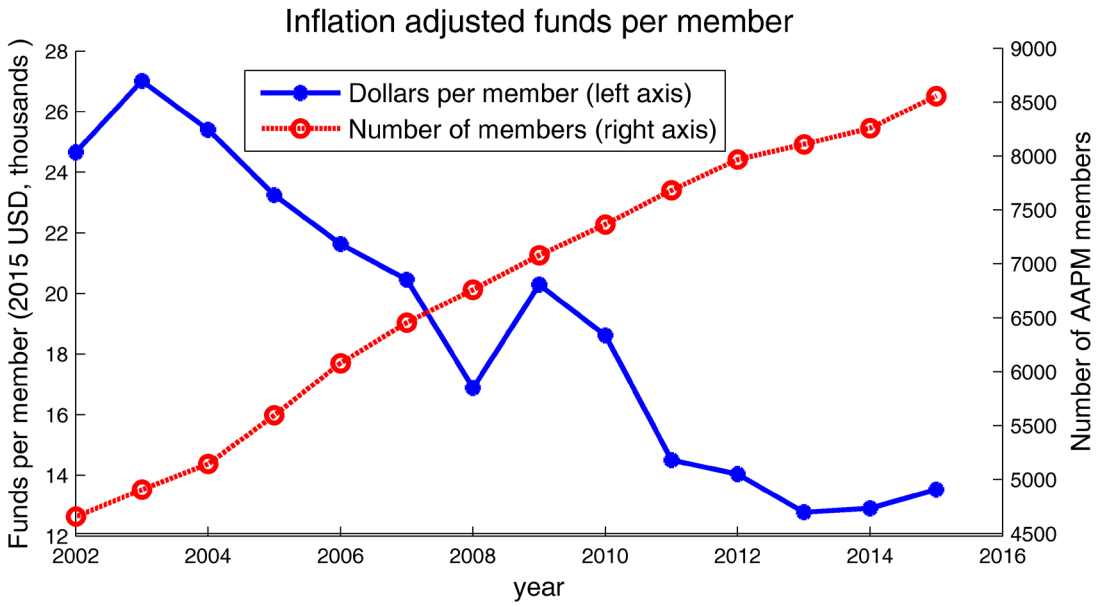

Figure 90: Dollars per member adjusted for inflation (USD year 2000, in thousands) compared to the number of members from 2002 to 2015.

3

Table 17: The extent to which different organizations awarded funding to AAPM members in 2015. *For NIAMSD, the 2015 budget was not available at time of writing, so the 2014 figure was used.

Institute abbreviation	Full name	Number of AAPM member grants funded in 2015 (% total)	Total funding of AAPM member grants in 2015 in USD millions (% total)	2015 budget (billions)	Percent of budget allocated to AAPM member grants
NCI	National Cancer Institute	124 (52%)	60.9 (53%)	4.95	1.2
NIBIB	National Institute of Biomedical Imaging and Biomedical Engineering	60 (25%)	26.1 (23%)	0.3	8.7
NHLBI	National Heart Lung and Blood Institute	19 (8%)	8.4 (7%)	2.9	0.7
NINDS	National Institute of neurological Disorders and Stroke	12 (5%)	3.6 (3%)	1.6	0.2
NIDDK	National Institute of Diabetes and Digestive and Kidney Diseases	2 (1%)	0.6 (1%)	1.7	0.04
NIAMSD	National Institute of Arthritis and Musculoskeletal and Skin Diseases	3 (1%)	1.0 (1%)	0.5*	0.2

4. Discussion

In this work, we have developed and tested a data mining algorithm that extracts PI_ID numbers from NIH records based on tests of name, institution, funding body, and temporal correlation. The utility of the algorithm has been demonstrated using AAPM membership data. We have tested the resultant PI_ID list for false positives and false negatives, obtaining

rates of 4% and 10%, respectively. Using this PI_ID list, we have extracted an estimate of NIH funding granted to AAPM members going back to 1985, although the accuracy of the data is likely to be degraded for years prior to 2002. We then presented a top level analysis of the resultant 'AAPM member research grant funding database'. Importantly, the input data required to extract this information is quite simple. This means the code could be quickly adjusted to produce similar databases for other groups of medical scientists, whom would presumably also be interested in such data. The highlighted columns in Table 1 show the data which is needed to adapt this algorithm to other groups.

The data generated in this work will be made available to AAPM members through the AAPM website. There are a number of uses of this database. The information could facilitate the AAPM board, councils, working groups and task groups and its members to:

(1) Provide general information about research activities
(2) Understand the magnitude and breadth of AAPM member research activities
(3) Determine trends of overall funding, and trends within specific research areas
(4) Identify funding opportunities where members have been successful in the past
(5) Track the success of AAPM members to specific research initiatives of the NIH
(6) Understand the demographics of successful researchers and identify and address areas with disparities
(7) Identify speakers for the AAPM annual meeting and chapter meetings
(8) Select reviewers for the AAPM annual meeting
(9) Lobby grant-funding bodies for increased consideration of medical physics applications
(10) Perform strategic planning

For non-AAPM members the database could allow:

(1) Media personnel to contact domain experts
(2) Grant funding agencies to seek appropriate reviewers
(3) Related societies (e.g. the American Institute of Physics) to solicit speakers for meetings

The approach we have taken appears novel with respect to analysis of NIH records. Although multiple authors have published on the overall NIH budget, we could not find any other publications where funding records associated with a given national (or international) cohort of investigators. Ref [14] attempted to extract information on radiation oncology funding based on the 'department' field of the NIH records, identifying 26 medical physics grants in 2013. A significant limitation of this approach, as acknowledged by the authors, is that the 'department' field is often left blank – for instance, in the year 2013, PERCENT of all NIH records had no recorded department.

There are several limitations of the current algorithm and initial results, and a variety of means by which the analysis could be improved. Probably the foremost of these is the point that the data used to develop the algorithm and the data used to detect false negatives was the same. It is important to note that the algorithm described in this paper is not data driven, but logic driven. As such, the data used in the development phase will not have as strong an

impact on the results as, say, a machine learning approach. Nevertheless, it is likely that the true false negative rate is somewhat higher than the 10% identified in this work, as there are probably failure points that have not been identified yet. However, identifying 89 of 99 member supplied PI_IDs gives us confidence that the obtained results are representative of the truth. The detected false negative rate of 4% is also encouraging. Although this was achieved after multiple iterations of manual curation of the data, we believe it is possible to maintain or improve this in future iterations. Regarding future improvements to the algorithm: there are a number of data fields we have not queried that could provide additional discriminatory evidence. These include the department name, study section name, and the grant title itself. However, given that the false positive and negative rates outlined in this work are fairly good, further improvements to the algorithm at this point, while certainly worthwhile, may be pursuing diminishing returns.

Regarding the use of the PI_ID number to extract grants: according to the NIH, the PI_ID is a unique number that remains constant throughout a PI's career [8]. While we did not find any exceptions to the 'uniqueness' claim, there are several instances in which one PI does have multiple PI_IDs. In this work, both PI_IDs were kept in such instances, although the records are flagged for review. These cases all occurred in the early NIH records before funding information was available (i.e. before year 2000). Another limitation is that in this work, once a PI_ID number is associated with a member of the AAPM, it stays that way for all time. So for instance, if an investigator at one point was an AAPM member before changing fields, all subsequent grants would be incorrectly attributed to membership of the AAPM. We suspect that such cases are rare, but it is an area that can be improved upon in future work. Finally, at the present time, the code cannot process names with nonstandard alphanumeric characters (such as umlauts and commas), which means a number of names are skipped each year. This is another clear area in which to improve, although it did not result in any detected false negatives in the current analysis.

The original motivation for this work was to capture data relating to NIH funding awarded to AAPM members as PIs. Figure 6 shows how much funding members of the AAPM are awarded each year as a proportion of total available funding. For the total funding, we have used the combined budget of the top 4 funding institutes associated with the AAPM research community (Figure 8), since this should provide a more relevant baseline than the overall NIH budget. From 2000 to 2007, the proportion of funding captured by AAPM members was trending upwards. Between 2007 and 2014, this proportion has either dropped or stayed about the same (depending on whether you include multi PI grants on which AAPM members were not the contact PI). This analysis is complicated somewhat by the presence of the ARRA funding in 2009 and 2010, and the multi PI grant model introduced in 2006. Regarding the latter, it can be seen that the multi PI model is being utilized with increasing frequency by AAPM members, while the proportion of single PI grants is steadily falling. Figure 5 shows the overall amount of funding attributed to members of the AAPM between 2000 and 2015. It can be seen that the general trend here is upwards, and in 2015, a total of 115 million USD was awarded to AAPM members. When inflation is taken into account, however, it is clear that funding for AAPM researchers is in real terms lower in 2015 than it was in 2007. In this

time, the membership of the AAPM increased 33%, which means that the research funds per member in real terms has reduced by almost a factor of 2 since 2000 (Figure 9). The members of the AAPM are likely not alone in suffering this dilution of research support; the NIH budget itself has barely kept pace with inflation, a fact that has attracted much discussion [15-17].

Another important finding of this work is that among AAPM members, males were found to be about twice as likely to hold research funding as females. This inequality is likely not unique to the AAPM, and the data evident in this analysis is representative of more widespread gender disparity in research grant funding. Reference [18] gives an excellent overview of issues surrounding women's application and success rates with the NIH funding. The root cause deserves further consideration - for example: are females less successful in the NIH review process, or do they not submit as many grants applications? The data in this work cannot definitively answer with respect to the AAPM membership, but based on other studies conducted in the field of biomedical research, it appears that the number of applications submitted by females is a major factor [18]. The analysis therefore suggests that to address this inequality the AAPM consider developing mechanisms that better encourage and support female members applying for research funding.

It is important to note that this analysis captures only one source of research funding of AAPM members. Other important sources include federal funding agencies (NSF, DOE etc.) and industry research collaboration. The latter is likely to be particularly important, and particularly difficult to capture. It would appear that private industry invests on the order of twice as much money as the NIH into biotechnology research [19, 20], however it not clear how much of this is outsourced as opposed to performed in house. Capturing meaningful data in this regard is likely to prove challenging, but may be possible in part by parsing conflict of interest declarations required on journal publications and abstracts submitted to the AAPM Annual Meeting each year.

Our analysis shows that the amount of funding per AAPM member is decreasing at a rapid rate (Figure 9) - a result of increasing membership and a flat or declining NIH budget. The AAPM is hardly alone suffering such effects; however, there should still be cause for concern. Research drives scientific advances, and with science among the three pillars of the AAPM mission, the Association should consider the ramifications of its membership engaged less and less in research. If science is to remain a pillar of the AAPM, then the Association should recognize the downward trend in research leadership among its membership and prioritize means by which to better promote and sustain research as a vital aspect of medical physics. One aspect of this is to ensure that scientists who perform research in our core areas find an appropriate home in the AAPM and are not lost to other organizations. Another is to ensure that the AAPM remains aligned with important and competitive research fields pertaining to physics in medicine, and not limit itself to traditional areas. Finally, it is important that medical physicists participate in the definition of NIH funding priorities and its grant review process to stimulate funding opportunities for medical physics researchers and to recognize high quality research grants in medical physics.

As a final point, AAPM members with interest in research and/or analytics pertaining to our Association should recognize the important role that they play in improving the accuracy of the data presented here. AAPM members who have at any point received research funding from the NIH should register their PI_ID within their AAPM member profile information. The PI_ID tag is the single most important identifier in correlating member data to NIH funding data. Note that one's PI_ID is not the same as one's eRA Commons name. To look up one's PI_ID, a person needs simply to login to their eRA Commons account, go to Personal Profile, then View 'Name and ID'. The PI_ID can be added to one's AAPM member profile by logging into the AAPM website, going to one's member profile page, and click the link entitled 'funding'.

In this work we have presented an algorithm that can extract records associated with an input list of researcher. This resultant database is a sustainable resource that can be queried by the AAPM board, councils, working groups task groups, and members to address a wide range of questions pertaining broadly to research funding and grantmanship. The APM research database is available for AAPM members to download through the AAPM website. (reviewers note: Will be made available before report is published)

Acknowledgments

The authors would like to acknowledge all previous and current members of the Working Group for the Development of Research Database (WGDRD). In addition, we would like to acknowledge Ivaylo Mihaylov (University of Miami Health System), Paul Kinahan (University of Washington), and Shayna Knazic, Farhana Khan, and Tammy Conquest (AAPM). The authors would like to acknowledge financial support from the American Association of Physicist in Medicine (AAPM).

Conflicts of Interest

The authors are members or employees of the AAPM, and financial support was provided by the AAPM.

References

1. Buxton, M., S. Hanney, and T. Jones, *Estimating the economic value to societies of the impact of health research: a critical review.* Bulletin of the World Health Organization, 2004. **82**(10): p. 733-739.
2. OECD (2016). *Research and development (R&D) (indicator)*. 2016 17 February 2016]; Available from: https://data.oecd.org/rd/gross-domestic-spending-on-r-d.htm.
3. NIH (2016). *NIH budget*. 2016; Available from: http://www.nih.gov/about-nih/what-we-do/budget.
4. NIH (2016). *NIH peer review process*. 2016; Available from: http://grants.nih.gov/grants/peer_review_process.htm.
5. Keevil, S.F., *Physics and medicine: a historical perspective.* The Lancet, 2012. **379**(9825): p. 1517-1524.
6. NIH (2016). *RePORTER - NIH RePORTER*. 2016; Available from: https://projectreporter.nih.gov/reporter.cfm.
7. NIH (2016). *ExPORTER - NIH RePORTER Database Download*. 2016; Available from: http://exporter.nih.gov/.
8. NIH (2016). *Glossary of terms*. 2016; Available from: http://exporter.nih.gov/about.aspx.
9. LeBlanc, D.C., *Statistics: concepts and applications for science*. Vol. 2. 2004: Jones & Bartlett Learning.
10. NIH (2016). *National Institutes of Health Office of Budget Price Indexes| nih.gov*. 2016; Available from: https://officeofbudget.od.nih.gov/gbipriceindexes.html.
11. NIH (2016). *NIH and the American Recovery Investment Act (ARRA) | nih.gov*. 2016; Available from: https://recovery.nih.gov/.
12. NIH (2016). *Impact of sequestration on NIH | nih.gov*. 2016; Available from: http://www.nih.gov/news-events/news-releases/fact-sheet-impact-sequestration-national-institutes-health.
13. NIH (2016). *Types of Grant Programs | grants.nih.gov*. 2016; Available from: http://grants.nih.gov/grants/funding/funding_program.htm.
14. Steinberg, M., et al., *National Institutes of Health funding in radiation oncology: A snapshot.* International Journal of Radiation Oncology* Biology* Physics, 2013. **86**(2): p. 234-240.
15. Moses III, H. and E.R. Dorsey, *Biomedical research in an age of austerity.* JAMA, 2012. **308**(22): p. 2341-2342.
16. Dzau, V.J. and H.V. Fineberg, *Restore the US lead in biomedical research.* JAMA, 2015. **313**(2): p. 143-144.
17. Alberts, B., et al., *Rescuing US biomedical research from its systemic flaws.* Proceedings of the National Academy of Sciences, 2014. **111**(16): p. 5773-5777.
18. NIH (2016). *Women in Biomedical Science | nih.gov*. 2016; Available from: https://nexus.od.nih.gov/all/2014/08/08/women-in-biomedical-research/.
19. Moses, H., et al., *Financial anatomy of biomedical research.* Jama, 2005. **294**(11): p. 1333-1342.
20. Moses, H., et al., *The anatomy of medical research: US and international comparisons.* Jama, 2015. **313**(2): p. 174-189.

Appendix 2
Discussion of simulation uncertainty

Chapters 3, 4, 5, and 6 of this thesis all rely on some form of simulated data. An important question one must ask when working with simulated data is: how closely are the simulated results likely to match real world results?

The potential sources of disagreement between simulation and experiment can be split into three categories, described below.

1. **Simulation uncertainties inherent to the simulation.** These uncertainties occur from sources fundamental to the solving methods; for instance the discretisation of space into a number of finite elements. In general these are relatively small, and can be controlled to a desired accuracy. A table outlining the inherent sources of error in the simulations is included below.
2. **Simulation uncertainties relating to unknown inputs**. These uncertainties arise when one is unsure of the exact geometry of the component being simulated (to a lesser extent, uncertainty in material properties also plays a part). Since this is the general case when simulating proprietary linac components, it is a very important uncertainty for MRI-linac systems. However, the extent to which these inputs are incorrect and the impact of this will be situation dependant.
3. **Uncertainties inherent to the experimental setup.** These errors differ from case 2 in that the (ideal) component geometry may be perfectly known, but unknown errors occur during manufacture or assembly. These will be unique to each experimental setup. This is a very important source of uncertainty in chapter 6, where we believe that experimental alignment is substantially affecting the experimental results.

Table 18: Uncertainties inherent to the simulations carried out in this thesis (Case 1 from above)

Error source	**Comments**	**Models effected in this thesis**
Discretisation errors.	Errors arising from splitting continuous aspects of a system into discrete elements. Two such cases exist in the simulation in this thesis: Geometry is discretised into a number of discrete elements (finite element method) Electron beams are discretised 'macro particle' bins, e.g. 100 identical electron trajectories are simply simulated as one macro particle having the properties of 100 electrons. These errors can be quantified and controlled to a desired accuracy which with modern computers should generally be very low. In the models in thesis this error is always	All

	well within 1%.	
Emission model errors	As discussed in chapter 3, the emission models appear to be very sensitive to input parameters and to differ between programs. Without extensive experimental characterisation, these errors are difficult to quantify and control. On the other hand (and as discussed in our response to chapter 3), they do not have a strong impact on other results reported in this thesis.	Particularly DC electron guns. To a lesser extent, RF electron guns
Material property errors	Often simplified material models are used in the simulated models. For instance; perfect conductivity is assumed at metallic walls, perfect vacuum is assumed inside accelerator systems, and average BH curves are used to represent magnetic materials. It is often difficult to quantify the extent to which these simplifications affect the solution. However, in general it would be expected to be low	all
Missing physics	In general, there are multiple sets of physics at play in the performance of a given component. However, usually only one set of physics is simulated at a time. To give an example of this: in the RF electron gun simulation (chapter 5) the cathode is at a temperature of 1245 K. In real life, this would heat up the nearby copper, causing thermal deformation, in turn slightly changing the resonant frequency of that cavity. Another example would be in solving the RF fields; the loss of energy to the walls will also result in heating. Also, if the field becomes too strong in local areas the RF breakdown will occur. These are effects which could be investigated via simulation, but so far has not been. Overall, these are expected to be second order effects; the general behaviour of a system is characterised by the primary physics, but corrections may have to be made in the real world to compensate for e.g. thermal deformation.	All

Appendix 3
Statements of author contributions

Statements of authorship contributions

I hereby declare that the below authorship attribution statements are correct. In addition, I am the corresponding author on all papers which have been published or submitted.

Brendan Whelan Date: 30/09/2016

As supervisor for the candidature upon which this thesis is based, I confirm that the authorship attribution statements below are correct.

Paul Keall Date: 30/09/2016

Chapter 3 – Triode electron gun

The project was conceived by Dragos Constantin, Lois Holloway, Paul Keall, Rebecca Fahrig and Brendan Whelan. Brendan Whelan led the project, performed all electromagnetic modeling, and data analysis. Lois Holloway and Dragos Constantin provided expert advice into the use of Opera software and supplied code from previous projects. Brad Oborn supplied an electron gun through his institution, and arranged for it to be dismantled at their local workshop such that each individual electrode could be studied in detail. Magdalena Bazalova-Carter performed the experimental measurements on the true beam linacs. Rebecca Fahrig and Paul Keall provided scientific guidance and oversight throughout the study. Brendan Whelan drafted the manuscript; all authors contributed to revising the finalized manuscript.

Chapter 4 – Magnetic shielding

The project was conceived by Brendan Whelan. Brendan Whelan led the project and performed all electromagnetic modeling and data analysis. Brad Oborn supplied the comsol MRI magnet model which was used. Paul Keall provided scientific oversight and guidance. Brendan Whelan wrote the chapter.

Chapter 5 –RF gun in-silico

John Schmerge had the initial idea that an RF gun could be used in radiation therapy. Rebecca Fahrig, Lois Holloway, John Schmerge, Paul Keall and Brendan Whelan were all involved in the early project conception. Brendan Whelan performed all 3D modeling and analysis. Stephen Gierman contributed expert knowledge and advice about the practical workings of RF guns, and more generally about accelerator and beam physics. Paul Keall and Rebecca Fahrig provided

scientific oversight and guidance. Brendan Whelan drafted the manuscript; all authors contributed to revising the finalized manuscript.

Chapter 6- Experimental beam line

Rebecca Fahrig initiated this project. Stephen Gierman, John Schmerge, Aaraon Tremaine, Ben Scott, Sami Tantawi, and Brendan Whelan were also all involved in the early stages of the project synthesis. All these people were also involved in the design of the beam line. Brendan Whelan performed 3D electromagnetic simulations to inform the design stage. A large number of technical staff at SLAC contributed to the construction and installation of the beam line and are listed below. Once the beamline was installed, Stephen Gierman led the process of commissioning the beam and collection of data, with assistance by Brendan Whelan. Brendan Whelan carried out the data analysis.

CAD schematics: Ann Trautwein, Ben Scott, and Abdullah P Rashed Ahmed

Vacuum systems: Mike Nalls, Mark Jacobson, Joe Vargas and Jon Dickey,

Installation and mechanical support: Robert DiMattie and Jose Eduardo Guerra

Radiation Safety: Shanjie Xiao

Power systems: Antonio de Lira, Greg Johnson and Tony Beukers

Machine protection: Luc Lessard

Chapter 7 – Omniax trial

Paul Keall conceived the initial study idea. Brendan Whelan, Paul Keall, Lois Holloway, Michael Barton, Robin Turner, Kuldeep Makhija, Leigh McGarvie and Miriam Welgampola contributed to the study design. Brendan Whelan coordinated patient recruitment with assistance from Michael Jackson and Ilana Feain. Brendan Whelan, Kuldeep Makhija, Leigh McGarvie and Miriam Welgampola ran the clinical sessions. Miriam Welgampola and Leigh McGarvie ran the Omniax device. Robin Turner provided expert advice into the statistical analysis of the data. Brendan Whelan performed the data analysis. Brendan Whelan drafted the manuscript; all authors contributed to revising the finalized manuscript.

Chapter 8 – MRI compatible patient rotation device

Paul Keall and Michael Barton conceived the project idea. Brendan Whelan, Paul Keall, Michael Barton, Lois Holloway, Leigh McGarvie, Gary Liney and Robba Rai all made substantial contributions to the design process. Rob Wilkins led the design and construction of the PRS. Project management during the design and construction phase was carried out by Brendan Whelan. Ilana Feain took over this role when Brendan Whelan was overseas. Once the device was made, Gary Liney, Robba Rai, and Brendan Whelan performed the MRI imaging studies.

Gary Liney, Brendan Whelan, and Lois Holloway carried out the phantom imaging measurements. Brendan Whelan, Jason Dowling, Gary Liney, Lois Holloway and Paul Keall designed the data analysis process. Jason Dowling carried performed all image registration and contour deformation analysis. Brendan Whelan drafted the manuscript including all figures and tables. All authors provided input into the manuscript draft.

Appendix 1 – NIH data mining

The initial concept of developing a research database for the AAPM was developed by Paul Keall, Rebecca Fahrig, Jeff Siewerdsen, and Eduardo Moros. For a number of years, the Working group attempted to implement this project by using AAPM coders, to little success. Jeff Siewerdsen and Thomas Yi first demonstrated a proof of concept of downloading NIH data and parsing it with Matlab. Thomas Yi wrote the code for this proof of concept stage. Brendan Whelan then conceived and implemented the data mining concept of testing multiple fields and using the combined evidence to make a decision. Brendan Whelan wrote all the code used in the final algorithm, and repeatedly tested and updated the algorithm to fine tune it. Brendan Whelan performed the analysis of trends in research funding which is presented in the report. Brendan Whelan drafted the report; all authors contributed to revising the report. James Deye provided expert knowledge about the NIH system and data; Michael Woodward performed the same role for the AAPM.