

Applying physical science techniques and CERN technology to an unsolved problem in radiation treatment for cancer: the multidisciplinary ‘VoxTox’ research programme

Neil G Burnet^{1,2}, Jessica E Scaife^{1,2}, Marina Romanchikova^{1,3}, Simon J Thomas^{1,3}, Amy M Bates^{1,4}, Emma Wong^{1,4}, David J Noble^{1,4}, Leila EA Shelley^{1,5}, Simon J Bond^{1,6}, Julia R Forman^{1,6}, Andrew CF Hoole^{1,3}, Gillian C Barnett^{1,4}, Frederic M Brochu^{1,7}, Michael PD Simmons^{1,7}, Raj Jena^{1,2}, Karl Harrison^{1,7}, Ping Lin Yeap^{1,7}, Amelia Drew^{1,7}, Emma Silvester^{1,7}, Patrick Elwood^{1,7}, Hannah Pullen^{1,7}, Andrew Sultana^{1,7}, Shannon YK Seah^{1,7}, Megan Z Wilson^{1,7}, Simon G Russell^{1,4}, Richard J Benson^{1,4}, Yvonne L Rimmer^{1,4}, Sarah J Jefferies^{1,4}, Nicolette Taku^{1,2}, Mark Gurnell^{1,8}, Andrew S Powlson^{1,8}, Carola-Bibiane Schönlieb^{1,9}, Xiaohao Cai^{1,10}, Michael PF Sutcliffe^{1,7}, Michael A Parker^{1,7}

¹Cancer Research UK VoxTox Research Group, University of Cambridge Department of Oncology, Cambridge Biomedical Campus, Addenbrooke’s Hospital, Cambridge, CB2 0QQ, UK; ²University of Cambridge Department of Oncology, Cambridge Biomedical Campus, Addenbrooke’s Hospital, Hills Road, Cambridge, CB2 0QQ, UK; ³Department of Medical Physics and Clinical Engineering, Addenbrooke’s Hospital, Cambridge, CB2 0QQ, UK; ⁴Oncology Centre, Addenbrooke’s Hospital, Cambridge University Hospitals NHS Foundation Trust, Hills Road, Cambridge, CB2 0QQ, UK; ⁵Department of Engineering, University of Cambridge, Trumpington Street, Cambridge, CB2 1PZ, UK; ⁶Cambridge Clinical Trials Unit, Box 401, Cambridge University Hospitals NHS Foundation Trust, Addenbrooke’s Hospital, Hills Road, Cambridge, CB2 0QQ, UK; ⁷Cavendish Laboratory, University of Cambridge, JJ Thomson Avenue, Cambridge, CB3 0HE, UK; ⁸Metabolic Research Laboratories, Wellcome Trust-MRC Institute of Metabolic Science, University of Cambridge and National Institute for Health Research, Cambridge Biomedical Research Campus, Addenbrooke’s Hospital, Cambridge, UK, CB2 0QQ, UK; ⁹Department of Applied Mathematics and Theoretical Physics, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Road, Cambridge, CB3 0WA, UK; ¹⁰Mullard Space Science Laboratory (MSSL), University College London, Surrey, RH5 6NT, UK

*Corresponding author: email address(s): ngb21@cam.ac.uk and ccd35@medschl.cam.ac.uk

ABSTRACT

The VoxTox research programme has applied expertise from the physical sciences to the problem of radiotherapy toxicity, bringing together expertise from engineering, mathematics, high energy physics (including the Large Hadron Collider), medical physics and radiation oncology. In our initial cohort of 109 men treated with curative radiotherapy for prostate cancer, daily image guidance computed tomography (CT) scans have been used to calculate delivered dose to the rectum, as distinct from planned dose, using an automated approach. Clinical toxicity data have been collected, allowing us to address the hypothesis that delivered dose provides a better predictor of toxicity than planned dose.

Keywords: Multidisciplinary; physical sciences; radiation toxicity.

Received: February 2017. Accepted: April 2017. Available online: June 2017.

INTRODUCTION

An estimated 3.4 million Europeans were diagnosed with cancer in 2012, and globally the figure was 14.1 million (Ferlay et al., 2015). Radiotherapy (RT) is the most effective non-surgical treatment for cancer (Bentzen et al., 2005), and 60% of those who receive it are treated with the objective of cure (Möller et al., 2003, IAEA Human Health Series, 2010). This amounts to over 2 million people across Europe each year (Ferlay et al., 2015), indicating the scale and importance of RT in the curative treatment of cancer. The success of RT in eradicating tumours depends chiefly on the total radiation dose. What limits this dose is the tolerance of the normal tissues surrounding the tumour. As the dose is increased so the incidence and severity of normal tissue damage also rises, and when severe, normal tissue damage can produce

significant morbidity, which may even be life-threatening. Selection of the appropriate treatment is based on a balance between lowering the dose to keep the incidence of severe normal tissue complications at an acceptably low level, and raising the dose to increase the probability of tumour control.

Complications from RT are the result of the dose actually delivered to the patient (Jaffray et al., 2010). In many circumstances this differs from the planned dose, for example due to daily positional variation in mobile internal anatomy. There is a steep dose-cure relationship, both in experimental animal systems and in man, and a 5% increase in dose will typically achieve an increase in tumour cure in the range 5-10% (Suit, 2002). Normal tissue dose-toxicity relationships are even steeper, at least for some tissues (Barnett et al., 2009). Thus, small



increases in dose to tumour or reductions in dose to normal tissues can result in clinically valuable improvements, with the potential to improve quality of life for the individual patient and reduce society's burden of care.

We report a research study which has applied methodologies and expertise from the physical sciences to exploit 'marginal gains' that might follow from the use of delivered dose to more accurately predict toxicity in patients receiving radiotherapy treatment for cancer.

THEORETICAL BACKGROUND

Preparation of a RT plan is based on a single computed tomography (CT) scan performed some days before the start of the treatment course. However, this does not capture day-to-day differences resulting from internal organ positional change. For example, in men receiving curative RT for prostate cancer, the shape and position of the rectum, which lies immediately behind the prostate, are known to vary from one day to the next (de Crevoisier et al., 2005, Scaife et al., 2014), which alters the collateral dose delivered to the rectum. In one study using daily image guidance (IG) CT scans, accumulated delivered dose (D_A) to the rectum was different to planned dose in all patients (Scaife et al., 2015). We sought to exploit the daily IG CT scans, which are taken every day to ensure accurate and reproducible treatment being delivered to the tumour, to calculate the daily dose delivered to the rectum (Figure 1). At present it is impossible to calculate this in routine practice, because the rectum would have to be contoured manually on each daily CT scan for each patient, which is labour-intensive, requires training, and is very slow.

We hypothesised that development of specific computerised solutions might allow us to use image guidance scans to develop an individualised adaptive treatment approach, which could be used to reduce toxicity, increase tumour dose, or both. If this is to be developed for clinical use, then the processes of de-archiving the IG CT data, curating them, identifying the rectum on each scan, computing the delivered dose D_A , and reporting the results must be fully automated. Expertise has been brought together from the high energy physics community (in particular from the team at the Large Hadron Collider - LHC), engineering, applied

mathematics, medical physics and radiation oncology to develop such an automated methodology to calculate delivered dose to the rectum on IG CT scans.

METHOD AND DATA

The primary objectives of the project were to develop methods to calculate D_A in normal tissues, and to use this to compare planned and delivered doses with toxicity experienced by individual patients during, and for a minimum of two years after, treatment. Therefore a clinical study was developed, and Ethical Committee and research permissions were obtained from the hospital. Clinical toxicity data are being collected from >850 patients with prostate, head and neck, and central nervous system (CNS) tumours. We believe this is the largest such dataset anywhere in the world, linking toxicity outcome with IG imaging from which delivered doses can be calculated. This will allow us to go on to address the hypothesis that that delivered dose provides a better predictor of rectal toxicity than planned dose.

Interaction between the clinical and non-clinical teams was needed for multiple steps, three of which presented major challenges (Table 1). The first related to handling a volume of data, originally estimated to total approximately 4 TB^A which was too large to be managed within the radiotherapy centre. This task included retrieval, quality assurance, transfer and processing of the data (Figure 2). The second key challenge was to develop a method for automated contouring (auto-contouring) of the relevant critical structures, in a data set of many thousands of scans. In our research programme overall, with over 850 patients with prostate, head and neck, and CNS tumours, with different organs at risk in each anatomical site, there are 24,000 scans, containing 316,000 CT slices, with a total of 1.1 million contours required. One contour is defined as a structure outlined on 1 slice in a CT scan. Contouring of such a large number of scans would be impossible by hand: assuming that each contour takes an expert 1 minute to complete, this would take 1 million minutes, or 1.9 person-years. Moreover, the final objective is for this to be done in real time for clinical application, which cannot realistically rely on expert human operators. The third major challenge is the recalculation of the dose using the IG CT scans, which is a computationally very intensive task (see below).

^A Our initial estimate of scale was of ~60 Megabytes per scan, giving 2.4 GB of raw data per patient. Data generated during processing could double this, making a total volume of 4 TB for 850 patients. This was too large to be managed within the capacity available in the NHS radiotherapy centre at the time, but was easily accommodated in the LHC Tier 2 Centre at Cambridge. In reality, scans are down-sampled and some of the image guidance scans are quite short and therefore smaller, together reducing the total data volume by more than a factor of 10, to about 250 GB after processing.




Fig. 1. Planning kV CT scan (A) and scan-of-the day image guidance TomoTherapy HiArt™ MV CT scan (B), at the same level in the same patient receiving curative radiotherapy for prostate cancer. Note that the rectum (arrowed) is of modest size at the time of the planning scan, but dilated principally with air in the treatment scan.

Tab. 1. Areas of interaction required between the clinical and the physics environment. The major challenges are shown with '*'. The first was the general challenge of handling volumes of data, originally estimated to total approximately 4 TB, which was too large to be managed within the radiotherapy centre. Not only storage, but also curation and processing of the data, was required (Figure 2).

1	De-archiving of image guidance (IG) CT scans, and associated data, from manufacturer's proprietary archive, including alterations required with new versions of the archive software structure; extraction of planning CT scan data from separate archive, and correlation of the two data sets
2	Tokenisation & export of each patient's data set through hospital firewall
3	Data storage, curation and processing at High Energy Physics group *
4	Auto-contouring of rectum on IG CT scans *
5	Dose calculation of rectal dose on each daily IG CT scan, using auto-contouring application (Python & MatLab) *
6	Toxicity score mapping to accumulated dose (D_A)
7	Automation of control of processes 3-5

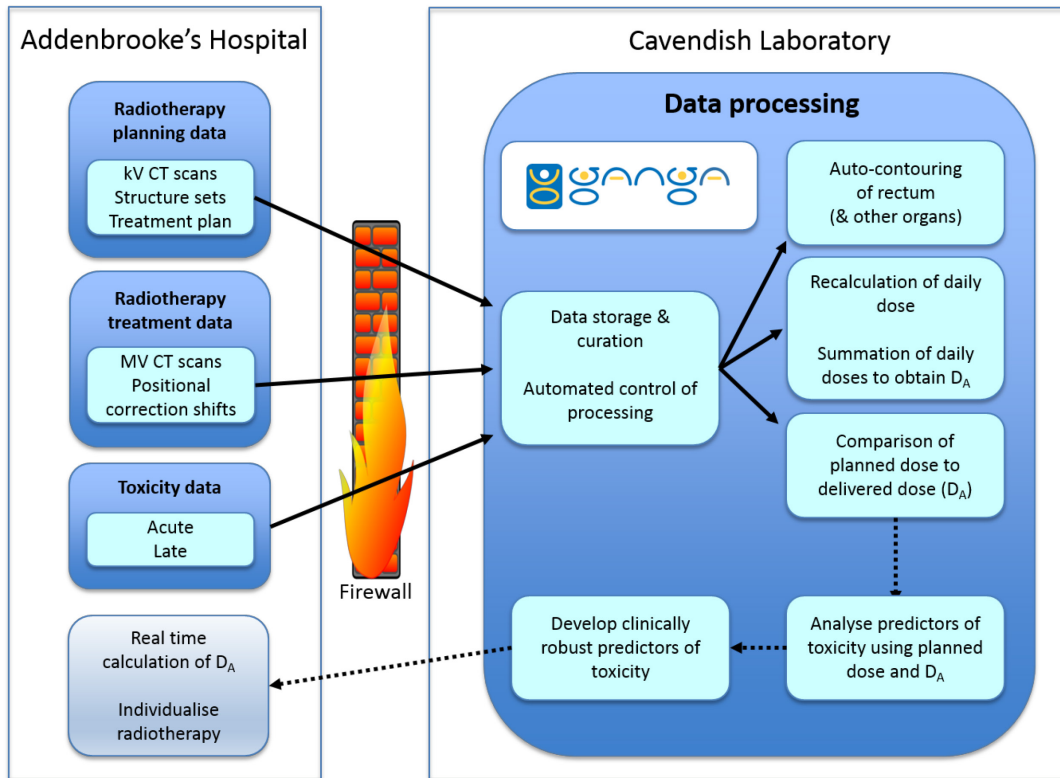


Fig. 2. VoxTox data flows. kV, kilovoltage; MV, megavoltage; D_A Accumulated dose.

Proof-of-concept patient cohort

We commenced with a cohort of patients who had received curative RT for prostate cancer, with a special focus on delivered dose to the rectum. We selected 109 patients, each treated to a median dose of 74 gray (Gy) in 37 fractions, delivered on 5 days per week over 7½ weeks. All patients were imaged with on-board CT and positional correction was made immediately prior to treatment. Treatment was delivered with intensity modulated RT (Mackie, 2006, Burnet *et al.*, 2010).

Manual contours for comparison with automated contouring algorithm

The first step in developing an automated method to auto-contour the rectum was to create a set of *manual* rectal contours to which the automated system could be compared. For 10 patients, the rectum was contoured on all 37 daily IG scans by 1 operator (JES), who had a median intra-observer Jaccard Conformity Index (JCI) for contouring of 0.87^B. Random variation in rectal position during radiotherapy for prostate cancer was seen to be two

to three times greater than that predicted from inter-fraction motion of the prostate (Scaife *et al.*, 2014), and differences between planned and accumulated dose to the rectum were seen in all 10 participants (Scaife *et al.*, 2015). These data were also used to estimate the potential value to patients (Scaife *et al.*, 2015, Barnett *et al.*, 2015). To assess inter-operator variation, one IG scan in another 6 patients was contoured by 8 consultant oncologists specialising in prostate radiotherapy, who showed a median inter-observer JCI of 0.79.

Automated processing pipeline

A system for automated control was required for all the steps in the process of calculating accumulated delivered dose, to perform automated batch processing of image retrieval from our data store, rectal auto-contouring, dose re-calculation and dose-volume reporting. Computing models for petabyte-scale data analysis in experiments at the LHC were used in developing a data-processing system for the VoxTox study. A software framework has been implemented to allow a wide range of computing tasks to be carried out efficiently (Barrand *et al.*, 2001). Although the framework

^B For 2 intersecting sets, A and B, the Jaccard Conformity Index (also known as the Jaccard index, Jaccard similarity coefficient, concordance index, and initially described by Paul Jaccard as the ‘coefficient de communauté’ (Hanna *et al.*, 2010)), is defined as the size of the intersection divided by the size of the union of A and B:

$$J(A, B) = \frac{A \cap B}{A \cup B}$$

itself is run in Python, algorithms may wrap code in other languages, such as Matlab, used for some of our processes (Thomas et al., 2011, Thomas et al., 2016). The computing jobs for a full-scale analysis are tracked using Ganga (Gaudi and Grid Alliance) (Mościcki et al., 2009). This system provides an efficient solution for the analysis work, involving just over 250GB of data, and would scale easily for use in our follow-on studies, with significantly higher data volumes.

Automation of patient imaging extraction

In order to manage the volume of data, we implemented a bespoke automated software solution to locate patients consented for the study, extract patient data directly from the electronic archives without having to use clinical pathways (Romanchikova et al., 2017), remove all personal-identifiable information, assign a study-specific patient identifier (token), and convert the extracted data to Digital Imaging and Communication in Medicine (DICOM) format. The insertion of the token is essential for later correlation with clinical toxicity data collected over time. Data on positional corrections applied during each image guidance procedure were included in the exported data to improve the accuracy of the image registration and facilitate calculation of D_A . The anonymised DICOM imaging data were then transferred to the Physics facilities for storage, curation and processing in our LHC Tier 2 Centre.

Automated contouring

Another important use for the manual contour data set was to define search boundaries to provide a starting point for the automated algorithm (Scaife et al., 2015, Cai et al., 2016). The automated contouring algorithm itself is based on the method developed by Chan and Vese in 2001 for segmenting non-medical images (Chan and Vese, 2001). This requires a good initialisation, achieved by registering the manual contour from the planning scan onto the image guidance scan. Although our own developments have improved the overall performance, we regard this as only an interim solution (Sutcliffe et al., 2015, Scaife, 2016, Simmat et al., 2012, Whitfield et al., 2013). At the level of the prostate itself, the imaging cannot distinguish the anterior border of the rectum from the posterior edge of the prostate. However, since these 2 structures are anatomically fixed, we use the anterior part of the contour from the planning scan to provide this border (Cai et al., 2016, Sutcliffe et al., 2015, Scaife, 2016).

Recalculation of accumulated delivered dose (D_A) using the IG scans

We use our in-house independent calculation system (Thomas et al., 2011) to re-calculate the D_A for each fraction, for each patient, based on the IG CT scan, and the patient position correction shifts, extracted from the archive (Thomas et al., 2016). The daily D_A is converted

to a dose-surface map (DSM) (Buettner et al., 2009), which can be summed to provide cumulative D_A and a dose difference map. The system can be run on a cluster of independent machines, without user intervention, and without tying up the clinical planning system.

RESULTS

Evaluation of automated contouring algorithm

In order to assess the effectiveness of the automated auto-contouring algorithm, in our cohort of 109 patients, all CT slices from one randomly selected megavoltage (MV) scan for each patient were visually evaluated against the objective of >70% of contours being acceptable (Lu et al., 2006). Out of 1107 slices in total, 821 (74%) were found to be acceptable.

The auto-contouring was compared to manual contours on 370 scans from 10 patients, and performed with a median JCI of 0.79 (IQR 0.74 to 0.79) (Figure 3) (Scaife et al., 2015). Auto-contouring was also compared to 8 consultant oncologists and performed with a median JCI per scan of 0.64 (IQR 0.53-0.71) (Figure 4) (Scaife, 2016).

Although performance of the automated algorithm was generally successful, there were discrepancies that we wish to remove. This is especially important since ultimately we aim to implement a method for altering a patient's treatment based on the automated contours. Work is in progress to develop this further by: firstly, segmenting the treatment scan in three dimensions rather than slice by slice; secondly, using manual segmentation from the planning scan, combined with anatomical and biomechanical knowledge of the organ to define a 'shape prior'; thirdly, rendering the computational solution feasible and robust by replacing the non-convex energy by a sequence of convex energies (Cai et al., 2013), making the auto-contour less sensitive to spurious suboptimal contours and poor initialisation; and fourthly, integrating a machine learning-framework, to achieve an automated choice of model parameters (Calatroni et al., 2015, Criminisi and Shotton, 2013). Nevertheless, we are not aware of any other fully automated solution to contour the rectum on image guidance scans that can deliver results comparable to our system, making this a ground-breaking development.

Proof-of-concept patient cohort results

The system was tested on our 109 prostate patient cohort, for which there is median 4 year follow up, with late rectal toxicity data collected prospectively. The planning scan, and 37 image guidance CTs (totalling 4033), were retrieved, automated contouring of the rectum was performed, and the image guidance scans were used to compute D_A . The entire process took 94 hours for the 109 patients, an average of 52 minutes per

patient. However, this was achieved using approximately 240 machines in parallel. The average time for the dose calculation on the MV CT scans is about 6 hours per MV scan per patient, using a single machine. Although slow, this would still allow calculation of daily accumulated dose overnight, before the next treatment is delivered.

Our results suggest a small improvement in correlation between dose and probability of rectal toxicity (bleeding) for accumulated delivered dose compared to planned dose, and details will be presented elsewhere (Shelley *et al.*, 2017).

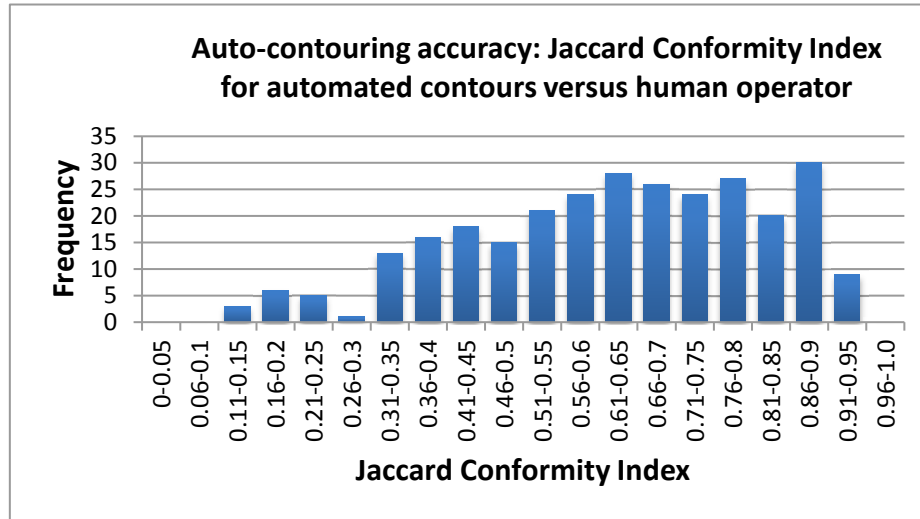


Fig. 3. Frequency distribution of JCI for 370 MV CT scans in 10 patients (~40,000 slices), compared to a single expert operator. The median JCI per scan was 0.64 (range 0.12 – 0.94).




Fig. 4. Example of an MV CT slice of the centre of the pelvis, showing rectal contours from eight senior radiation oncologists (red - median inter-observer JCI of 0.79.) and the automated algorithm (yellow). The median JCI for the automated system for this slice was 0.64. However, visually, the degree of agreement is striking.

Ant - Anterior (front of the patient), Post – posterior (back of the patient), Left – left side of the patient (shown on the right of the screen by convention), Right – right of the patient.

DISCUSSION AND CONCLUSIONS

To address the challenge of calculating accumulated delivered dose (D_A) as distinct from planned radiation dose, in order to improve toxicity predictions in patients receiving radiotherapy for cancer, we brought together a group that applied physical science expertise to medical questions.

Collaboration between oncologists and medical physicists is a routine part of clinical care for patients receiving radiotherapy. The collaboration with engineering and high energy physics, later expanded to include mathematics, was the result of finding researchers with an interest in the application of science to a wider context, including people with cancer. Having researchers collaborating from different disciplines and backgrounds stimulated novel approaches to problem-solving which would not otherwise have occurred.

The links between the groups, emphasising the multidisciplinary teamwork, are shown in Figure 5 (see also Fig 2). At the outset the Programme was divided into 5 workstreams (WS): WS0 – programme management; WS1 – extraction of imaging data, calculating delivered dose, modelling toxicity risk; WS2 – curation and storage of data, image segmentation and biomechanical modelling, integration of daily D_A calculation; WS3 – collection and analysis of clinical toxicity data; WS4 – *in silico* modelling of irradiated normal tissue cells in luminal structures. The concept

and design of these individual workstreams arose from detailed and iterative discussions between the different disciplines, focussed on a real clinical problem, and using or developing clinically applicable (i.e. not research-restricted) algorithms.

Organisationally, the group meet for 3 monthly Progress Meetings, bringing the whole group together to present and, importantly, discuss active components of the programme and their results. This has been highly effective in exchanging information between workstreams, guiding the direction of the work, and maintaining enthusiasm in the group. In addition, regular informal meetings take place (fortnightly), some of which serve as multi-disciplinary PhD supervision meetings. Informal meetings with the medical statisticians occur about every 2 months, to address specific statistics questions. The groups are housed at 3

different campuses in Cambridge, so regular, frequent, formal and informal meetings have been an important aspect of the collaboration. The programme is overseen by a Programme Board, who determine over-arching science policy and compare progress to the deliverables set at the beginning of the Programme.

For the Progress Meetings, anyone working in a similar area is encouraged to attend. This includes undergraduates and postgraduates linked to the VoxTox Programme (10 undergraduate projects, 1 PhD and 2 MPhils awarded, 2 PhDs and 1 MPhil underway). The group has also contributed to the annual 'Physics at Work' outreach event for schoolchildren (reaching > 500 students and teachers each year), and the Programme has driven 1 international and 2 national symposia.

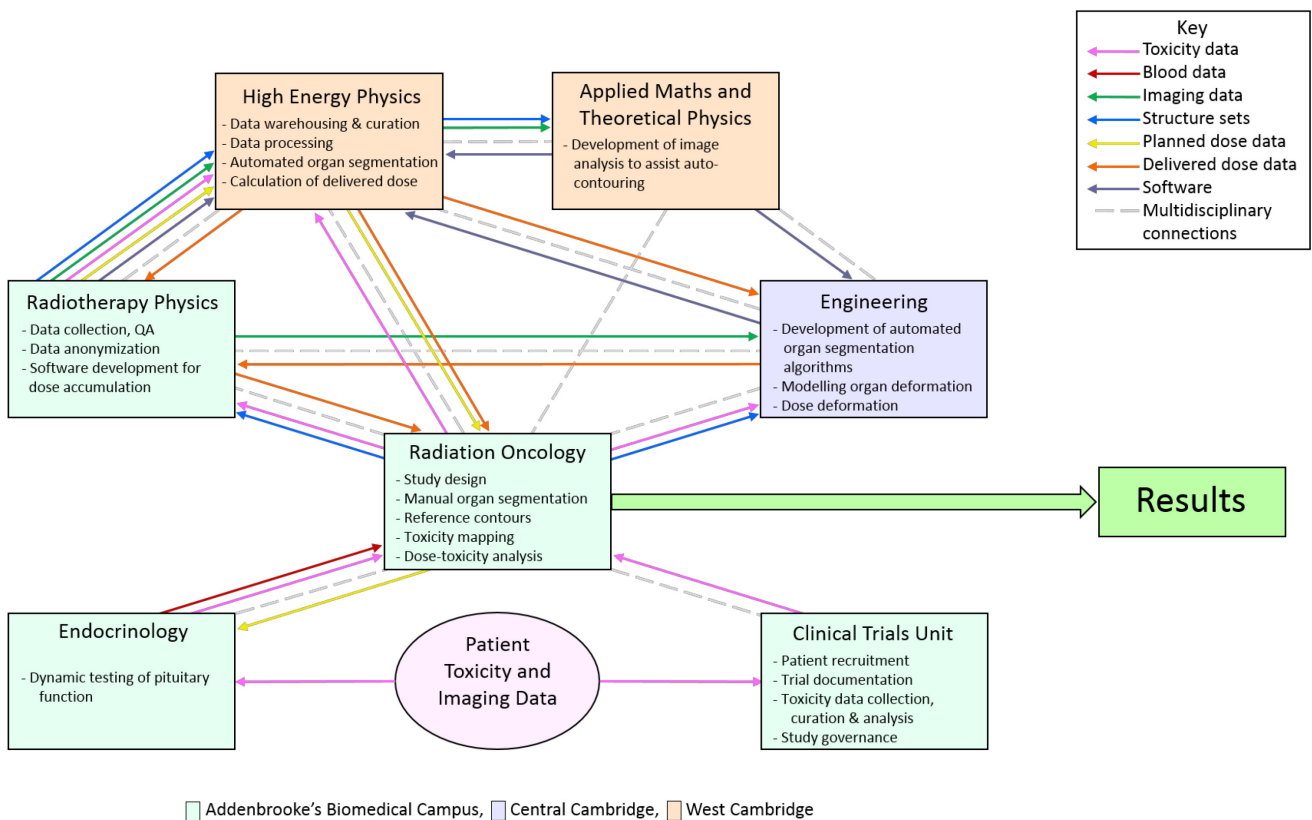


Fig. 5. Administrative and scientific multidisciplinary connections, together with data flows used in the Programme. Background colours indicate on which campus each group is located.

Automated analysis has been applied to daily IG CT scans. Accurate segmentation of soft tissues is very challenging because these scans intrinsically have a lower diagnostic quality than planning CT scans due to higher noise and low soft tissue contrast. However, IG CT scans are of good enough quality to see the boundaries of the rectum, and they are used routinely for IG in clinics around the world. Where the anterior rectal wall cannot be distinguished from the prostate, human anatomical structure permits use of the contour from the planning scan. Deformable registration is a technique that cannot reliably be used to determine rectal size and position (Simmat *et al.*, 2012, Godley *et al.*, 2013). Automation also has the potential to improve the reliability and speed of adaptive treatment.

Although it would be attractive to use an imaging modality with higher quality, for example magnetic resonance imaging (MRI), such imaging needs to be performed at the time of treatment and in the treatment position. Although MRI-linear accelerators are now being deployed, they are comparatively expensive and few in number, and there is as yet no evidence on their role in improving RT. It is unlikely that there will ever be sufficient machines to treat all European men requiring RT for prostate cancer, and CT-based systems will remain the mainstay at least for the immediate future. There is thus potential value for patients and society in developing the use of IG imaging based on CT technology.

A better understanding of the relationship between delivered dose and planned dose offers the potential to improve predictions of toxicity. If differences can be identified during a course of treatment then the possibility arises to adapt an individual patient's treatment based on the delivered dose and associated toxicity prediction. For men receiving curative RT for prostate cancer, demonstration of a difference between delivered and planned dose would be potentially valuable, to allow either dose escalation to the tumour (for lower than expected toxicity risk) or re-planning to abrogate excess risk to the rectum (for patients with risk higher than expected). Across the population this could lead to more men being cured and fewer suffering from complications of treatment, even though toxicity has been significantly reduced by the use of modern radiotherapy techniques.

The development of a methodology based on sophisticated computing rather than highly trained staff was considered essential to maximise reproducibility on a large patient cohort, to reduce inter-observer variability and to eliminate the human error associated with manual processing. It may also facilitate deployment in today's economic climate.

The collection of toxicity data from patients in our clinical study has provided us with the largest resource of patient IG imaging linked with toxicity information anywhere in the world. This has allowed us to start addressing the hypothesis that delivered dose provides a

better predictor of toxicity than planned dose (Shelley *et al.*, 2017).

Conclusions

Links between departments and effective knowledge exchange have been essential in our research programme. Although needing further refinement, our approach allows automated contouring of the rectum and calculation of the daily delivered dose, for comparison with the planned dose. Inter-disciplinary expertise has added considerable value, allowing us to tackle a major challenge related to improving the outcome of patients with cancer, by more accurately predicting the risk of a patient developing toxicity.

ACKNOWLEDGEMENT

We are grateful to Dr Mike Sharpe for helpful discussions on the importance of marginal gains, and to the many other colleagues who have encouraged and supported the work. We would especially like to thank the additional oncologists who generously provided contours for this work: Drs Alex Martin, Cathryn Woodward, Gail Horan, and Luke Hughes-Davies.

FUNDING

JES was supported by Cancer Research UK through the Cambridge Cancer Centre. NGB, ASP and MG are supported by the National Institute of Health Research Cambridge Biomedical Research Centre. KH, MR AMB, EW and SJB were supported by the VoxTox Research Programme, funded by Cancer Research UK. DJN is supported by Addenbrooke's Charitable Trust and Cancer Research UK through the Cambridge Cancer Centre. FMB was supported by the Science and Technology Facilities Council. MPDS was part supported by the VoxTox Research Programme, funded by Cancer Research UK. RJ was part supported by the VoxTox Research Programme, funded by Cancer Research UK. LS is supported by the Armstrong Trust. XC was supported by the Isaac Newton Trust. CBS acknowledges support from the EPSRC Centre for Mathematical and Statistical Analysis of Multimodal Clinical Imaging, the Leverhulme Trust, the EU-RISE project CHiPS and the Cantab Capital Institute for the Mathematics of Information. NT was supported by a Gates-Cambridge Scholarship, funded by the Bill and Melinda Gates Foundation, PLY and SYKS by the Singapore Government.

REFERENCES

- Barnett, G.C., et al., 2009. Normal tissue reactions to radiotherapy: towards tailoring treatment dose by genotype. *Nature Reviews Cancer*, 9(2): pp. 134-42.
- Barnett, G.C., et al., 2015. Incorporating genetic biomarkers into predictive models of normal tissue toxicity. *Clinical Oncology (Royal College of Radiologists)*, 27(10): pp. 579-87.
- Barrand, G., et al., 2001. GAUDI — A software architecture and framework for building HEP data processing applications. *Computer Physics Communications*, 140: pp. 45-55.
- Bentzen, S.M., et al., 2005. Towards evidence-based guidelines for radiotherapy infrastructure and staffing needs in Europe: the ESTRO QUARTS project. *Radiotherapy Oncology*, 75(3): pp. 355-65.
- Buettner, F., et al., 2009. Assessing correlations between the spatial distribution of the dose to the rectal wall and late rectal toxicity after prostate radiotherapy: an analysis of data from the MRC RT01 trial (ISRCTN 47772397). *Physics in Medicine & Biology*, 54(21): pp. 6535-48.
- Burnet, N.G., et al., 2010. Practical aspects of implementation of helical tomotherapy for intensity-modulated and image-guided radiotherapy. *Clinical Oncology (Royal College of Radiologists)*, 22(4): pp. 294-312.
- Cai X., et al., 2016. Automatic contouring of soft organs for image-guided prostate radiotherapy (Abstract E35-0631). *Radiotherapy and Oncology*, 119(Suppl 1): pp. S895-S896.
- Cai, X., Chan, R., Zeng, T., 2013. A two-stage image segmentation method using a convex variant of the Mumford-Shah model and thresholding. *SIAM Journal on Imaging Sciences*, 6(1): pp. 368–390.
- Calatroni, L., Chung, C., De Los Reyes, J.C., Schönlieb, C-B., Valkonen, T., 2015. Bilevel approaches for learning of variational imaging models, arXiv:1505.02120 [math.OC].
- Chan, T.F. and Vese, L.A., 2001. Active contours without edges. *IEEE transactions on image processing: a publication of the IEEE Signal Processing Society*, 10(2): pp. 266-77.
- Criminisi, A. and Shotton, J. (Eds). 2013. *Decision forests for computer vision and medical image analysis*, Springer-Verlag, London.
- de Crevoisier, R., et al., 2005. Increased risk of biochemical and local failure in patients with distended rectum on the planning CT for prostate cancer radiotherapy. *International Journal of Radiation Oncology Biology Physics*, 62(4): pp. 965-73.
- Ferlay, J., et al., 2015. Cancer incidence and mortality worldwide: sources, methods and major patterns in GLOBOCAN 2012. *International Journal of Cancer*, 136(5): E359-86
- Godley, A., Sheplán Olsen, L.J., Stephans, K., Zhao, A., 2013. Combining prior day contours to improve automated prostate segmentation. *Medical Physics*, 40(2): 021722.
- Hanna, G.G., Hounsell, A.R., O'Sullivan, J.M., 2010. Geometrical analysis of radiotherapy target volume delineation: a systematic review of reported comparison methods. *Clinical Oncology (Royal College of Radiologists)*, 22(7): pp. 515-25.
- IAEA Human Health Series, 2010. *Planning National Radiotherapy Services: A Practical Tool*. No. 14, International Atomic Energy Agency, Vienna.
- Jaffray, D.A., Lindsay, P.E., Brock, K.K., Deasy, J.O., Tomé, W.A., 2010. Accurate accumulation of dose for improved understanding of radiation effects in normal tissue. *International Journal of Radiation Oncology Biology Physics*, 76(3 Suppl): S135-9. Review.
- Lu, W., et al., 2006. Deformable registration of the planning image (kVCT) and the daily images (MVCT) for adaptive radiation therapy. *Physics in Medicine & Biology*, 51(17): pp. 4357-74.
- Mackie, T.R., 2006. History of tomotherapy, *Physics in Medicine & Biology*, 51(13): R427-53. Review.
- Möller, T.R., Einhorn, N., Lindholm, C., Ringborg, U., Svensson, H., 2003. Radiotherapy and cancer care in Sweden, *Acta Oncologica*, 42(5-6): pp. 366-75.
- Mościcki, J.T., et al., 2009. GANGA: a tool for computational-task management and easy access to Grid resources, *Computer Physics Communications*, 180(11): pp. 2303-2316.
- Romanchikova M., et al., 2017. Automated translation of radiotherapy imaging, planning and dose data from proprietary storage to a user-defined format using the example of TomoTherapy and DICOM RT. *Physics and Imaging in Radiation Oncology*, in press.
- Scaife, J., et al., 2014. Random variation in rectal position during radiotherapy for prostate cancer is two to three times greater than that predicted from prostate motion. *British Journal of Radiology*, 87(1042): 20140343. Erratum in: *British Journal of Radiology*, 2014, 87(1043): 20149003.
- Scaife, J., et al., 2015. Accumulated dose to the rectum, measured using dose-volume histograms and dose-surface maps, is different from planned dose in all patients treated with radiotherapy for prostate cancer. *British Journal of Radiology*, 88(1054): 20150243.
- Scaife, J., et al., 2015. Exploiting biological and physical determinants of radiotherapy toxicity to individualize treatment. *British Journal of Radiology*, 88(1051): 20150172. Review.
- Scaife, J.E., 2016. Calculating accumulated radiotherapy dose to the rectum in order to improve late toxicity prediction in patients treated for prostate cancer. PhD dissertation, University of Cambridge.
- Scaife, J.E., et al., 2015. Accuracy of manual and automated rectal contours using helical tomotherapy image guidance scans during prostate radiotherapy. *Journal of Clinical Oncology*, 33 (suppl 7; abstract 94). <http://meetinglibrary.asco.org/content/141425-159> (accessed 29.11.16).
- Shelley, L.E.A., et al., 2017. Delivered dose can be a better predictor of rectal toxicity than planned dose in prostate radiotherapy. *Radiotherapy and Oncology*, 123(3): pp. 466-471. DOI: 10.1016/j.radonc.2017.04.008
- Simmat, I., et al., 2012. Assessment of accuracy and efficiency of atlas-based autosegmentation for prostate radiotherapy in a variety of clinical conditions, *Strahlentherapie Und Onkologie*, 188(9): pp. 807-15.
- Suit, H., 2002. The Gray Lecture 2001: coming technical advances in radiation oncology, *International Journal of Radiation Oncology Biology Physics*, 53(4): pp. 798-809.
- Sutcliffe, M.P.F., Harrison, K., Scaife, J.E., Parker, M.A., Romanchikova, M., 2015. Auto-contouring of the rectum

on megavoltage computed tomography images.
Cambridge University Engineering Department Technical
Report CUED/C-MICROMECH/TR.100, 18th August
2015. <http://www.comprt.org/publications/auto-contouring-rectum-megavoltage-computed-tomography-images>.

Thomas S.J., et al., 2016. Recalculation of dose for each fraction of treatment on TomoTherapy. *British Journal of Radiology*, 89(1059): 20150770.

Thomas, S.J., Eyre K.R., Tudor G.S., Fairfoul J., 2011. An independent dose calculation program for the checking of TomoTherapy plans. *Radiotherapy and Oncology*, 99(Suppl; 1): pp. S153.

Whitfield, G.A., Price, P., Price, G.J., Moore, C.J. 2013. Automated delineation of radiotherapy volumes: are we going in the right direction? *British Journal of Radiology*, 86(1021): 20110718.