An analysis of the effectiveness of automated pre-, post-, and intra-treatment auditing of electronic health records

A studentship Project
Funded in part by the Canadian Patient Safety Institute

2013-2014
December 31, 2014

Submitted by:
Ackeem Joseph
Dr. John Kildea

McGill University Health Centre

Table of Contents

Acknowledgements **3**
Project Summary **3**
Purpose / Objectives / Research questions **4**
Methodology **5**
Results **6**
Limitations **17**
Recommendations **18**
Student Learning **18**
References **21**
Appendices **21**

Acknowledgements

This studentship project was made possible by a financial contribution from the Canadian Patient Safety Institute. I would also like to acknowledge the support and assistance of doctors and staff from the Medical Physics Unit and Radiation Oncology Department at the McGill University Health Centre, including Dr. Jan Seuntjens, Dr. Carolyn Freeman, Dr. John Kildea, and Mr. William Parker. I acknowledge partial support by the CREATE Medical Physics Research Training Network grant of the Natural Sciences and Engineering Research Council (Grant number: 432290).

Project Summary

Electronic health records (EHRs) are fast becoming the standard in clinical practice. Compared to paper records, they offer numerous advantages to the medical team, including improved efficiency, streamlined workflow and ease-of-access to documents and data. They also offer the potential for major data-driven improvements in patient safety and in the quality of patient care. In terms of safety, an EHR offers the possibility of automated chart audits to ensure complete documentation, compliance with standards, and a means to search for anomalous and potentially dangerous practice.

The goal of this project was to provide a simple software system to automatically audit patient charts in a manner that can be replicated by any expert user of an EHR system. We have developed web-based software to regularly audit the electronic charts of all patients receiving treatment in our paperless radiotherapy clinic. By facilitating access to multiple databases within

the clinic, each patient's EHR is audited out-of-hours, prior to treatment, regularly during treatment, and post treatment. Anomalies such as missing documentation, non-compliant workflow, and practice that stands out significantly from the norm are monitored and flagged. Our software determines historical trends using existing patient data and audits new patient data by comparing them with the historical and by flagging significant deviations.

Purpose / Objectives / Research questions

Many situations that may compromise patient safety arise from incomplete patient records or instances where the standard-of-care was not adhered to. At our radiotherapy centre, all patient documents are created and archived electronically. Furthermore, all the steps that comprise a patient's treatment journey are kept in order using an electronic schedule. Each planning step is considered a task that must be completed by an assigned individual before the following task, assigned to another individual, may begin. Each treatment step is an appointment at which the patient is present and for which all treatment parameters are recorded and verified electronically by the radiotherapy treatment machines (linear accelerators). Occasionally, tasks remain open or are not completed by the assigned individual, causing inaccurate or unnecessarily long planning delay times. Likewise, appointments may be late or incomplete and treatments may be inconsistent from day to day or patient to patient. Software that compares the consistency of patient treatments has the potential to find errors before they occur.

For treatment planning of infrequently-utilized radiotherapy techniques, planners and physicians alike often refer to previously-treated plans for guidance. Often this is done ad-hoc

and using a list of previously-treated patients that is either personal to the planner or to the physician. Indeed, the patient list and the procedure to prepare the treatment may often exist only in the head of the experienced planner or physician. A robust, electronic and automated method for radiation oncology staff to refer to previously-treated patient plans has the potential not only to improve the efficiency of the overall planning process but also to minimize the effects of inexperience and lapses in memory, thereby improving patient safety.

With the above in mind we attempted to address the following two questions in this project:

1. Can automated auditing of EHRs improve compliance and reduce errors in patient treatments?
2. Can the provision of previously-treated data to radiation oncology staff improve patient safety?

Methodology

At our centre, we have developed software that provides a connection between our record-and-verify database (ARIA, Varian Medical Systems, Palo Alto, California) and a treatment registry (custom MySQL database that we designed in-house for this project) through the use of a web-interface. Our software uses HTML, CSS and JavaScript for the graphical user interface, and PHP and PERL to communicate between the two databases. This communication allows us to do the following:

1. Gather, in our treatment registry, relevant parameters from the treatment records of all patients who have undergone or who are undergoing radiotherapy at our centre.

2. Compare the patient's prescription against the prescriptions of similar patients already in our database to determine treatment standards and monitor outliers that may represent potentially dangerous errors or deviations from the standard-of-care.
3. Audit the patient's chart at the end of the treatment planning process and prior to the start of treatment to ensure agreement between intended prescription and planned medication (complex radiation dose distribution). This audit does not replace existing manual checks performed by clinical personnel but serves to augment such checks.
4. Log all errors (representing near misses) found by the system. These valuable data allow for both an analysis of the effectiveness of our electronic chart auditing system and they form the basis of incident learning to reduce future harm.

Results

Software overview

Our software is profile-based so that individual users can setup a profile to login, register, and monitor their own studies and configurations. Following the connection put in place between the ARIA database and our treatment registry using several custom PERL modules as a back-end tool, a front-end graphical user interface was built using AngularJS, a web-application framework written in JavaScript. This web-based front-end allows users to configure settings including:

1. Selection of relevant patient data from the ARIA database using a "control" web page,
2. Selection of treatment parameters to build historical trends on a "study" web page, and

3. Specification of parameters to benchmark audit new patient data against these historical trends on an “analysis” platform.

Custom PHP scripts were written to register user settings as a middle-end between the treatment registry and web-interfaces. Figure 1 summarizes the infrastructure of our software.

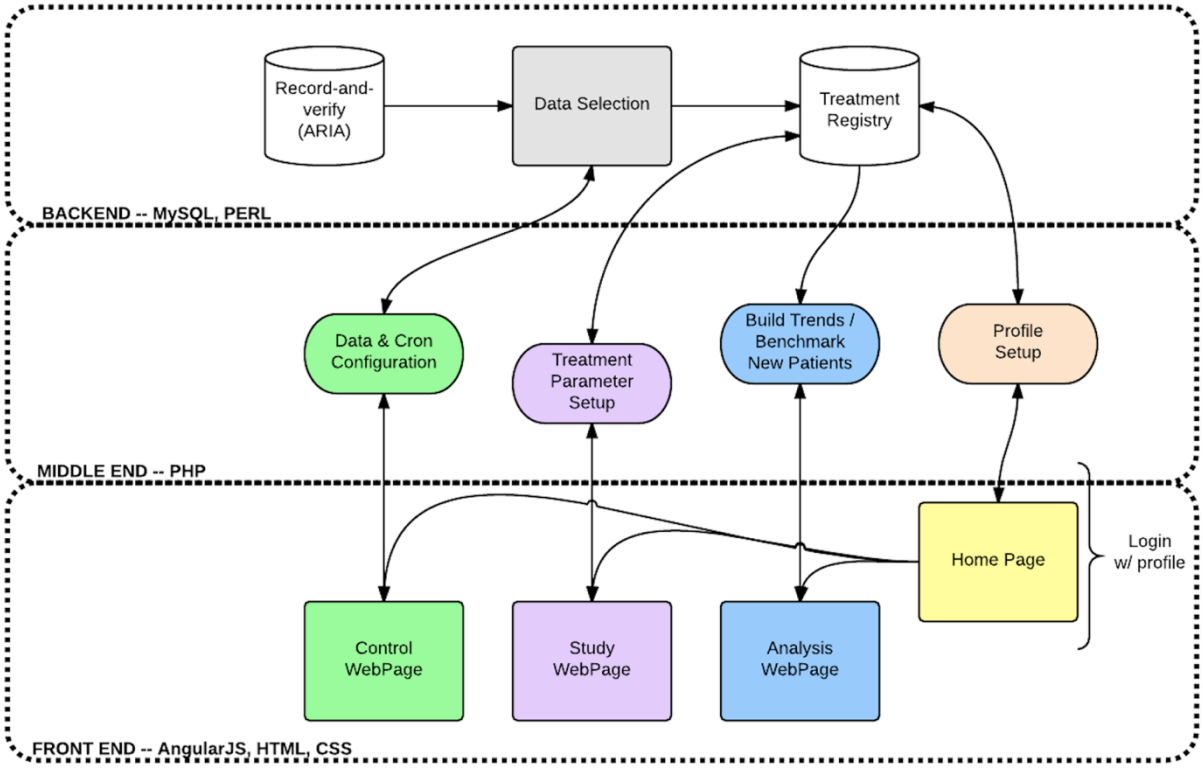

Figure 1: General infrastructure for our software. Profile-based so users can log-in and register their own settings. Relevant patient data are selected from the ARIA database and collected in our treatment registry on a time-based scheduler through the use of a “control” web-page. Treatment parameters are configured to build trends through the use of a “study” web-page. Trends are constructed using existing patient information and new patient data are benchmarked against these trends through the use of an “analysis” web-page.

Data Selection

In choosing which data to copy from Aria to the treatment registry, it is often useful to group similar data together using an alias. In many circumstances, tasks, appointments, and documents are labelled using an immense permutation of uppercase and lowercase characters, hyphens and spaces. For example, we might want to group together all consult appointments, regardless of whether the consult is named “Consult”, “CONSULT”, or “consult appt.” The data control interface allows individual users to manage aliases and provides a facility for them to populate the treatment registry using data pertaining to their own aliases. The following screenshots show the steps involved when configuring an alias.

1. Assigning an alias title.

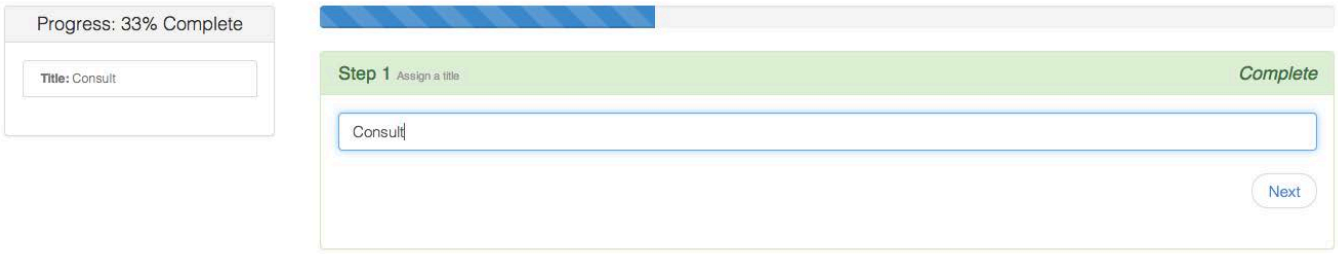

2. Assigning the type of alias (task, appointment, document, or plan).

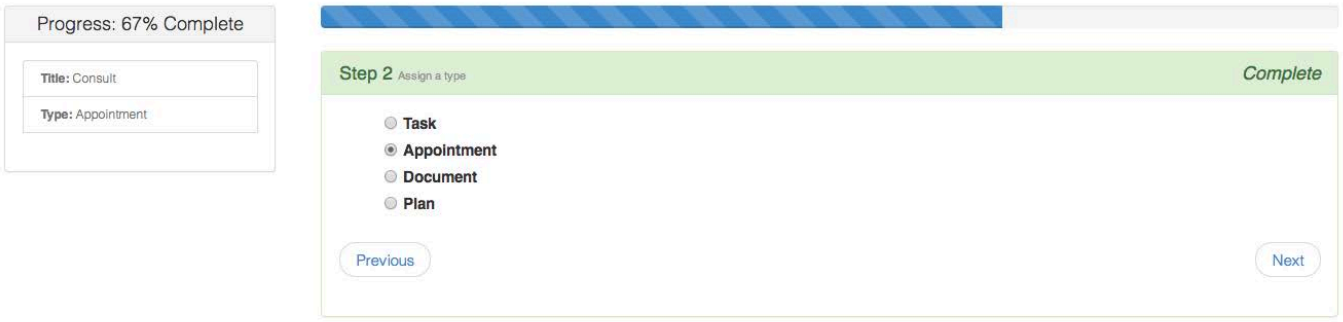

3. Grouping similar data together.

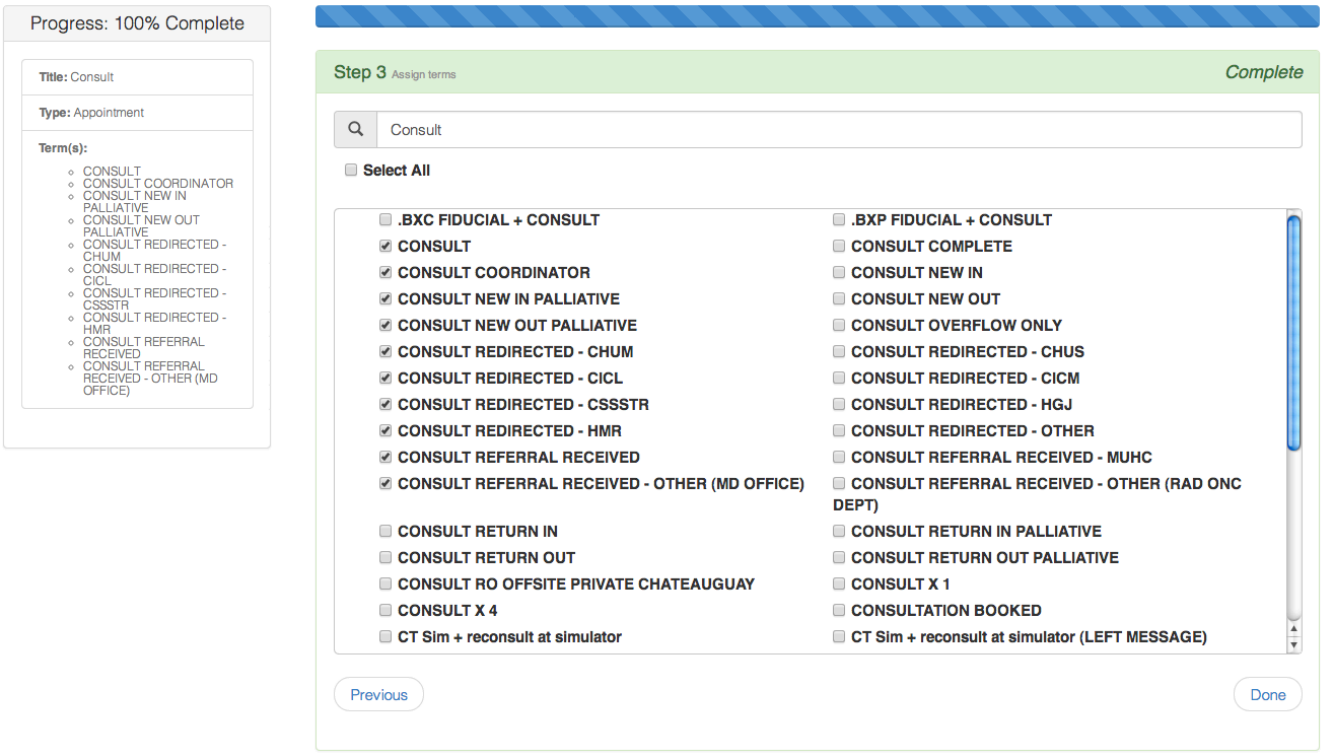

Cron Control

Data may be copied in large quantities, requiring a considerable amount of processing power that can result in slowing down clinical systems during regular working hours. It is thus essential to incorporate an automated and controlled feature to transfer data from ARIA to our treatment registry on a timely basis called a “cron”. The cron control interface allows individual users to schedule data transfer. As an example, the screenshot shown in figure 2 demonstrates a configuration to automatically run back-end PERL modules to update our treatment registry with data from ARIA each morning at 6 am.

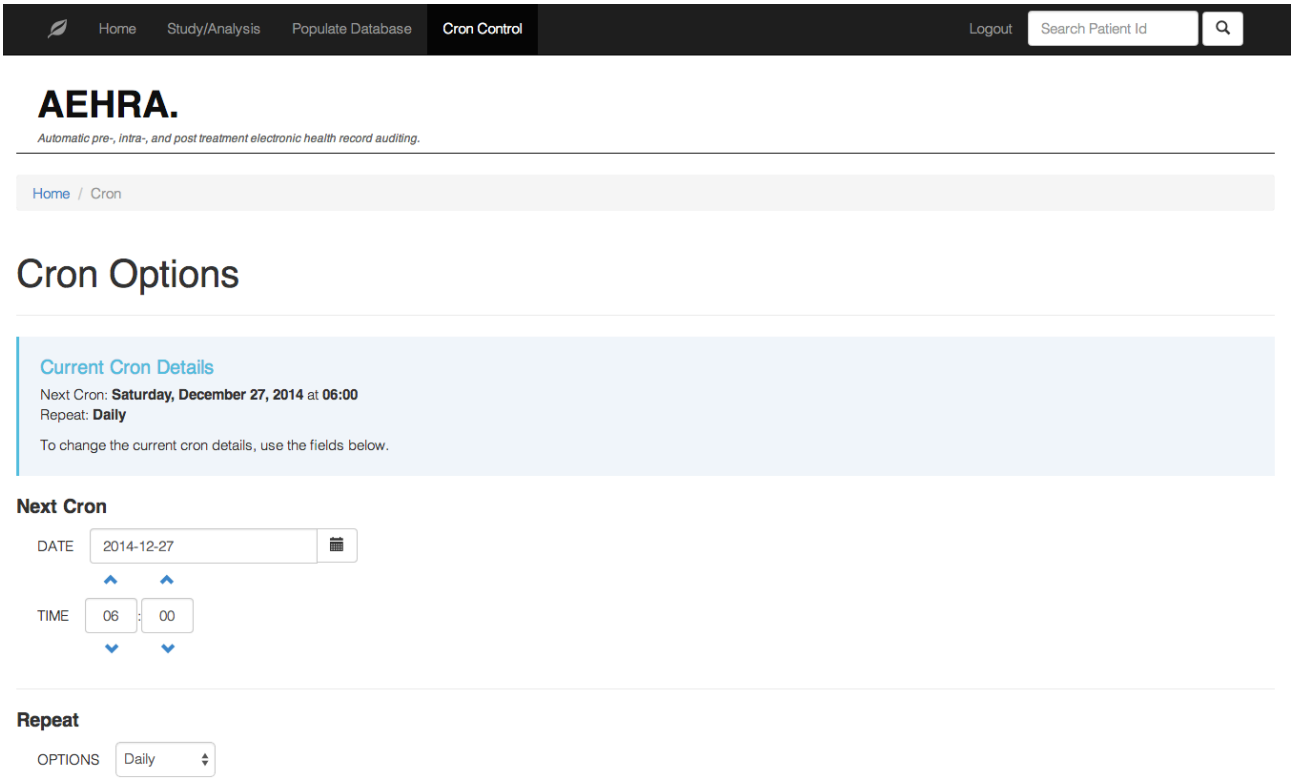

Figure 2: An example of our software showing the cron control web-page configured to establish data selection daily at 6 am.

Treatment Parameter Setup

The infrastructure put in place to populate our treatment registry with data from ARIA currently holds over 30,000 patient records. Using the data copied to our treatment registry, users can conduct studies by choosing relevant treatment parameters to build historical trends. For example, we might want to examine the treatment planning delays between a patient's initial consultation appointment and their end-of-treatment task. The study web-interface allows users to build custom studies by configuring the necessary treatment parameters. The following screenshots of our software show the steps involved when constructing a treatment planning delay study.

1. Assigning a title.

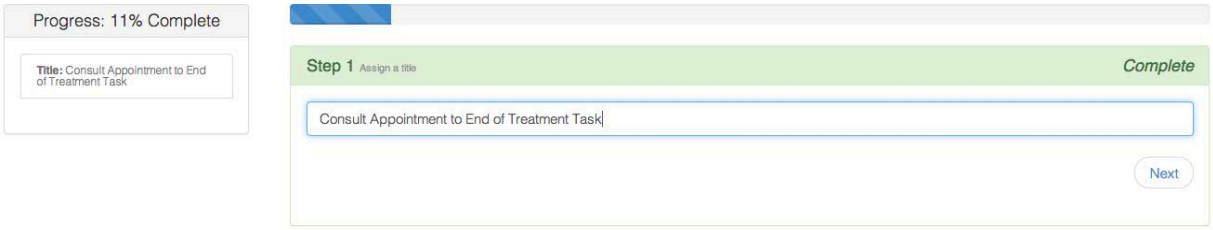

2. Choosing a starting parameter. In this case, the scheduled start timestamp of a consult appointment.

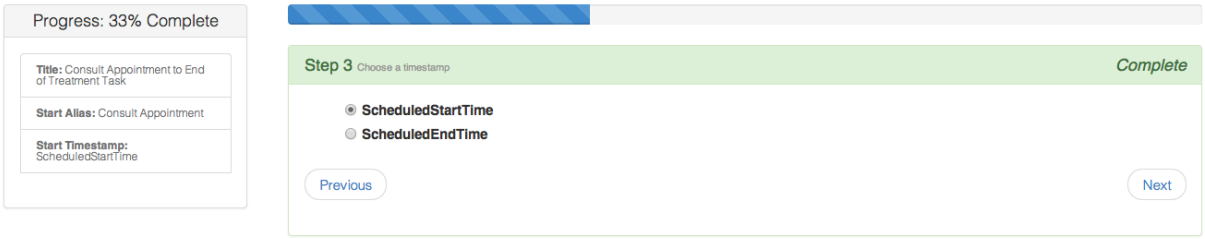

3. Choosing status conditions for the starting parameter. In this case, manually completed consult appointments are considered.

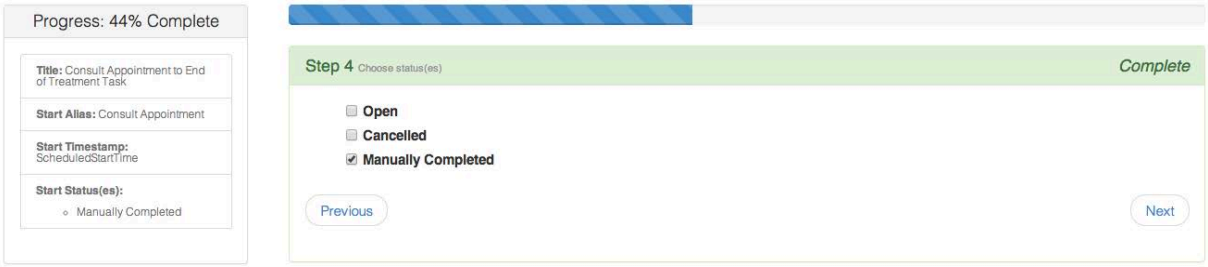

4. Choosing an ending parameter. In this case, the due date time of an end-of-treatment task.

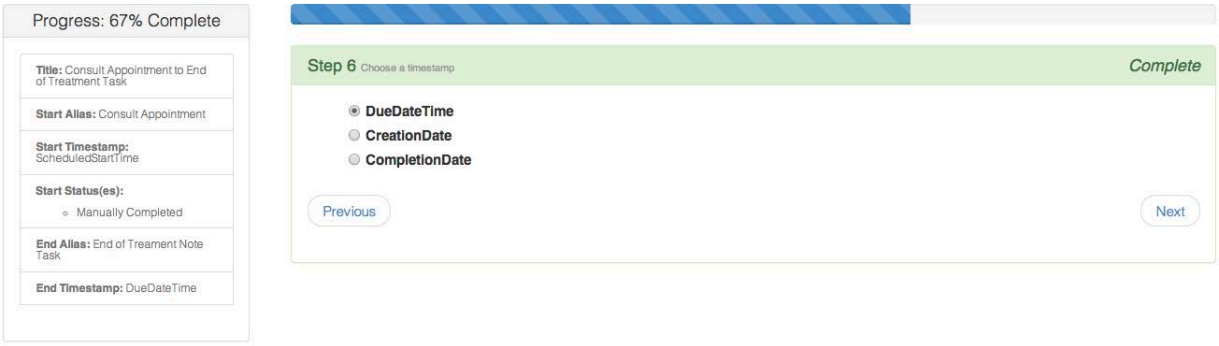

5. Choosing status conditions for the ending parameter. In this case, completed end-of-treatment tasks are considered.

Trend Determination & Benchmark Auditing

Using the relevant treatment parameters configured to build studies, we determine the statistical standard of care for previously-treated patients at our centre. The standard of care for a cohort of statistically similar patients is the average (or median) value determined for a particular treatment parameter, or the percentage compliance with a particular metric, in their medical records.

To benchmark-audit new patients we compare their EHR data to our historical standards and search for unexpected deviations. Our web-based "analysis" interface facilitates this benchmark auditing.

Treatment planning delays

Using the timestamp meta-data attached to each patient's tasking schedule, we provide treatment planning delays to determine the statistical standard-of-care for previously-treated patients at our centre. Figure 3 presents an example chart from our software showing a treatment planning delay distribution.

Figure 3: An example of our software showing the treatment planning delay standard determined by calculating the time between each patient's "Ready for MD contouring" task and the signature timestamp of their Rx prescription document. Delays are shown in 5-day bins, such that the historical and current delays appear side-by-side for the same bin.

From figure 3, the black histogram represents historical data collected from the records of patients who have previously undergone treatment. The blue histogram represents the collection of data of new patients who are undergoing treatment. Statistical thresholds representing 95% of the historical data are shaded in green while the remaining 5% of the historical data are shaded in red. Currently-treated patients who fall within the outlier region are marked as potentially anomalous and needing investigation.

DVH Registry

A dose-volume histogram (DVH) is a graph relating radiation dose to the volume of tissue receiving that dose in radiotherapy treatment planning. It is used to assess whether a plan meets desired constraints for particular volumes of interest. Despite their popularity, DVH curves are traditionally used in the context of individual patient plans and the experience of planners and physicians is usually encapsulated in point-specific dose-volume constraints. To better incorporate experience into the planning process, we used our EHR auditing software to develop a DVH registry that provides standard DVH curves based upon the plans of our previously-treated patients. Our DVH registry proved particularly useful for the intensity-modulated cranio-spinal irradiation (IMRT-CSI) technique.

Treatment planning for infrequently-utilized techniques such as paediatric IMRT-CSI is often complicated by the first-principles manner in which it is achieved. Owing to the low frequency of the technique, planners and physicians alike often refer back to previously-treated plans for guidance. Usually the list of previously-treated plans is personal to the

planner/physician and each plan is analyzed individually rather than collectively. By utilizing our DVH registry, we recorded the DVHs from ten recently-treated paediatric IMRT-CSI plans to define standards for patients treated to 36 Gy. Figure 4 presents a chart from our software showing the standard (median of the cohort) DVHs for a number of organs at risk used in our IMRT-CSI technique.

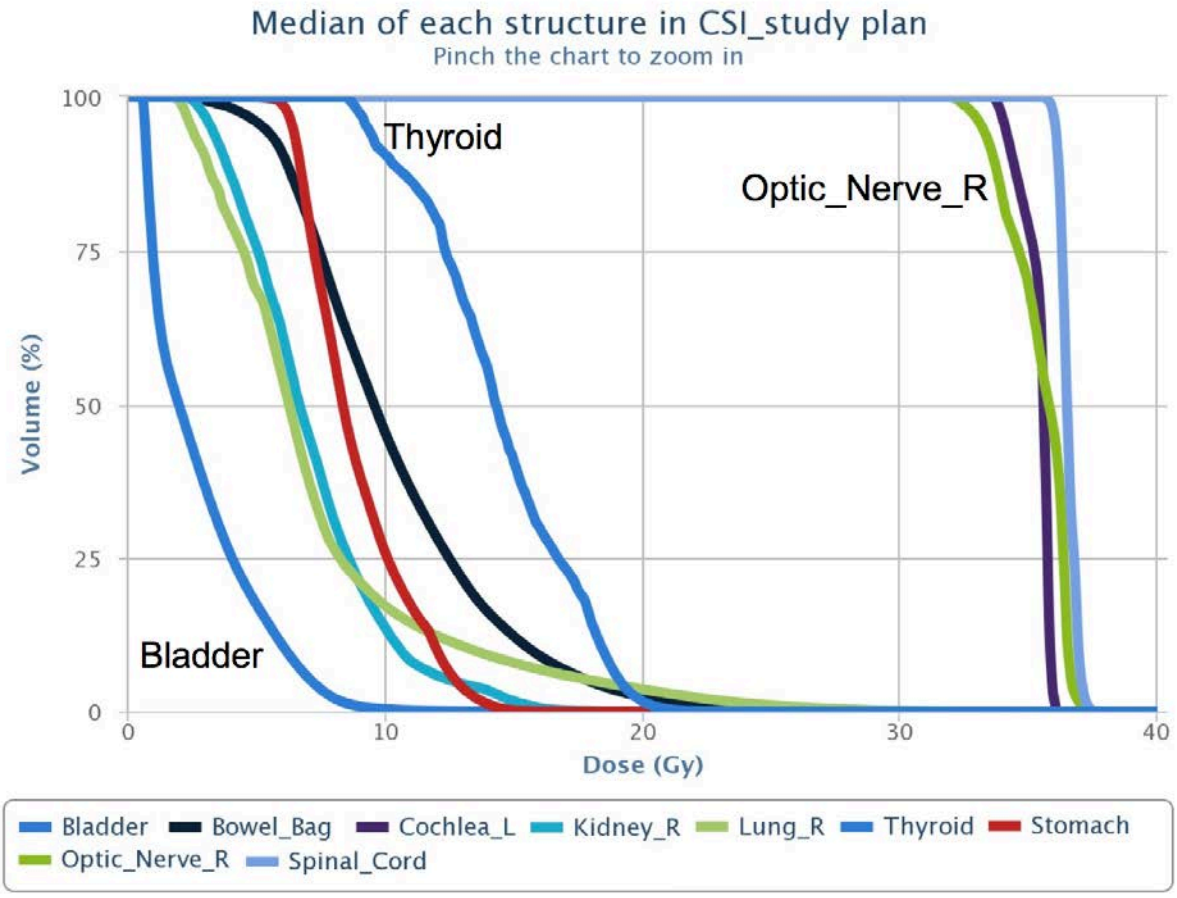

Figure 4: Standard (median of the cohort) DVHs derived for a number of structures used in our pediatric IMRT-CSI technique. Ten patient plans comprised the cohort.

Use of our DVH registry has been expanded to include other types of treatment plan, including techniques to treat prostate and breast cancer. Our software is used regularly by our clinical staff and allows rapid comparison of new treatment plans with historical standards.

Gantry Rotation Angles

By extracting the treatment field parameters attached to patient plan setups for our radiation therapy machines, we have used our treatment registry to record the gantry rotation beam angles so as to define standard gantry angles for patients receiving tangential breast treatment. Figure 5 presents a polar-chart distribution from our software comparing current and historical gantry angles from our tangential breast plans.

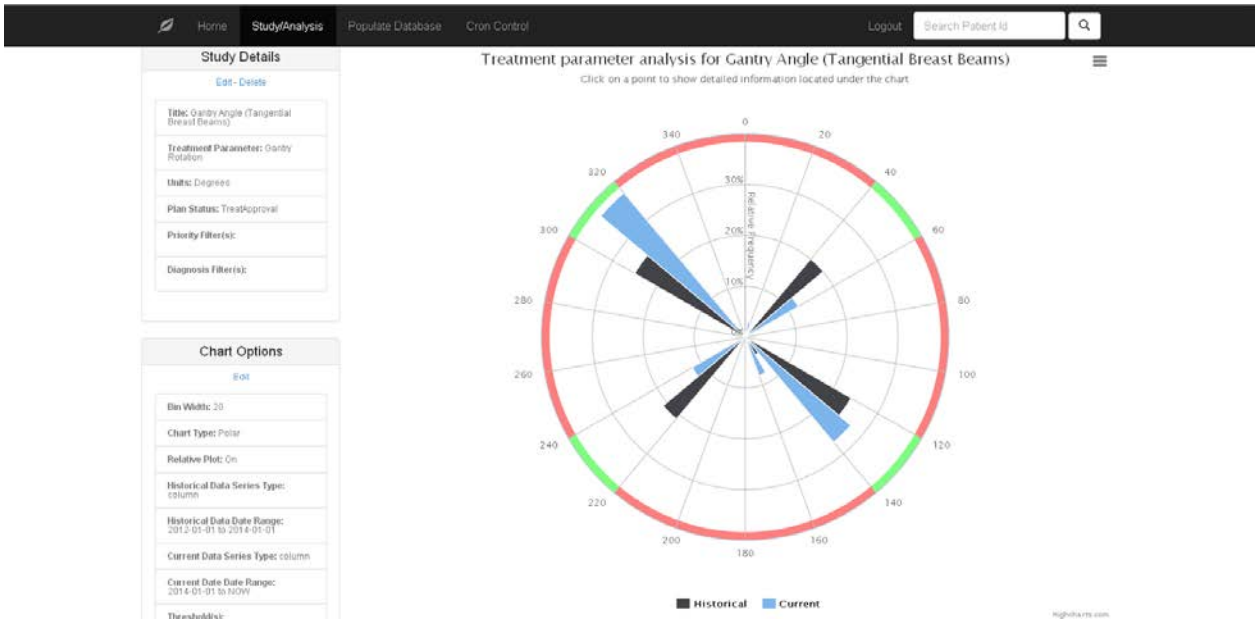

Figure 5: Screenshot from our software comparing current and historical gantry angles during tangential breast treatments. 95% of the historical data lie within the green region. Angles are shown in 20-degree bins, such that the historical and current delays appear side-by-side for the same bin.

Limitations

The treatment planning delay study was conducted using a diverse range of patient demographics and patient diagnoses. Smaller cohort samples with less spread in patient information such as, age group, diagnosis group, primary oncologist, would have produced more constructive results. Using patient cohorts for treatment planning delay analysis will be added to our software at a future date.

The DVH study was primarily limited by its inherently small sample size due to infrequently-utilized techniques. The sample size will hopefully soon be expanded by encouraging more clinicians to use our software. Increased participation will increase the statistical certainty of our data and thereby improve the efficacy of our software.

Recommendations

A number of outstanding features must be incorporated before evaluation of software can be fully established. The first issue is that the software needs a proper alerting system. The optimal method to determine clinical outliers and bring them to the attention of the clinical team is not obvious and several methods will have to be examined. An appropriate alerting system will be developed in consultation with clinical staff to provide proper EHR monitoring so to improve compliance and reduce errors in patient treatments. A system comprising an electronic dashboard on the clinic's internal website and automated emails to key team leaders is envisaged.

Student Learning

Development of this web-based software involved writing many lines of code in multiple programing languages. My first task was to familiarize myself with the ARIA database to coherently search for relevant data and communicate the information to our treatment registry. During this task I gained concrete knowledge about different types of queries and functions in both MySQL and PERL. I created structures of functions to help facilitate organization and data manipulation. This concept is known as object-oriented programming. I got the insight that was necessary to easily switch between databases and model information based on how they are

identified in the database. For example, I developed a custom-built PERL package that defines a data class (called an *object*) to control any meta-data related to appointments. It is a useful tool for setting and getting appointment information, comparing two appointment objects, and updating appointment objects. On one hand there was an investment of time required to construct such a framework of code but the other hand, this has allowed me to create new data types that are not already defined in the language itself which are useful building-blocks for writing standalone software.

During the development of our software, I had to deal with one unique fact: I needed to provide an instructional and conceptual graphical interface at the front-end while maintaining the complex communication between programming languages at the back-end. So, I had to design a user-friendly interface considering various human factors, all the while allowing the user to intuitively understand the purpose of the software. The web-based user interface was mainly focused on step-by-step processes and descriptions of each section of the software. Writing middle-end custom PHP scripts allowed our treatment registry to record user configurations and writing front-end custom JavaScript, HTML, and CSS code allowed the background processes to be explained and displayed to persons with no technical background.

I realized that I needed to think in advance about my audience, and about what and how to communicate. I also learned that I can use three-sectional frameworks (back-end, middle-end, and front-end) in a very early phase of the design process to provide a template for potential future projects. For example, after the kick-off of our software, we presented fully working pilot

studies to several members of our staff. Demonstrating the inter-connected and automated features ignited curiosity and enthusiasm for learning from data and a new machine learning research project involving a research team from the McGill School of Computer Science is currently under development.

Clinical teams were keen to evaluate specific and unique aspects of their clinical workflow. Writing structured and organized software allowed me to optimize and facilitate flexible domains incorporated into the software based on ideas inspired by the clinical team. For example, members of the clinical team suggested future projects built around provision of realistic waiting times for patients at our centre. Most of the technological requirements can be developed from our software. So by focussing on conceptual design and communication, this is something we can strive to develop using the models put forth by our software.

References

Non-applicable.

Appendices

The tool that was developed as a result of this project was sent on file via email.