Ethics of AI -Applying ALETHEIA framework to oncology

Speaker: Marianne Aznar

The introduction of Artificial Intelligence in the clinic raises new challenges in terms of workflow integration and ethical applications. While much of the research and development has so far focused om performance, other issues such as ethics, fairness and workforce training/ replacement have not been discussed in the same depth. Healthcare is not this only domain where AI such challenges. In radiotherapy, the aviation industry has long held as an exemplar in terms of safety and quality assurance processes. In this presentation, we will discuss the ALETHEIA framework developed by Rolls Royce for the aerospace industry, and its extension to oncology.

Starting a clinical computing service in Radiotherapy Physics
- Jonny Lee, University Hospitals Dorset NHS Foundation Trust

Aims and/or Background:

To improve support for the various software in the Radiotherapy Department, a specific role of "Radiotherapy Systems Lead Physicist" was created. This was to mainly to work on long term projects: dealing with system obsolescence, data storage/archiving, Information Asset management and associated compliance, leading on the upgrade or replacement of our current Record and Verify system and ProKnow. The post was also to manage the existing system administrator and to build a relationship with Trust IT. It is hoped this will be the start of a small Radiotherapy Clinical Computing group at University Hospitals Dorset.

Methods:

The Job Description and Person Specification were written, aiming to achieve an 8B banding. This was to ensure sufficient pay to attract candidates with sufficient experience of radiotherapy physics and computing to be effective.

***Results*:**

We successfully had the post approved within the Trust and recruited to the post.

Discussion around results

In my view there is the necessity for a specialist in computing in all Radiotherapy Physics departments; at a minimum working to keep systems compliant and planning ahead to avoid obsolescence. The role is of a similar level of challenge, specialist knowledge, responsibility and authority to that of a Radiation Protection Advisor and we banded it as such.

***Conclusion*:**

We have successfully increased resources in the area of computing within our department.

Key Words: IT, Radiotherapy Physics

Clinical decision support in triaging patients for surgery

Stephen Lake, Pre-op Team and Digital Services, Liverpool University Hospitals NHS FT

Background.

The NHS recovery plan requires perioperative care pathways to start at the latest when listing for surgery (NHS England, 2022). The 2023/24 NHS Standard Contract (NHS England 2023a) requires that all providers must implement a system of early screening, risk assessment and health optimisation for all adults waiting for inpatient surgery. NHS England (2023b) has provided guidance that aims to improve outcomes from surgery, reduce short-notice cancellations, and reduce length of stay. If patient's that would fail a pre-operative assessment are identified early, then they can begin complementary alternatives instead. If co-morbidities are unmanaged, e.g., diabetes, then this can be addressed before pre-op while on a waiting list. A patient may also be supported to alter their lifestyle choices, while on the waiting list, to reduce the risk of the anaesthetic and improve outcomes. Guidelines are available on the routine preoperative tests for elective surgery (NICE, 2016) to inform screening during a pre-op assessment and NHS England (2023c) are developing a suite of decision support tools. The pre-op team in Liverpool have extended a decision support tool, OSCAR, to assist in triaging patients for surgery that has been clinically evaluated (Murthy et al, 2008).

Methods

A booking form has been developed in the Trust's order communications system for surgeons to list patients with and provide details of relevant co-morbidities. An application has been developed for a healthcare assistant to record a patient's observations in the surgeon's clinic. A second application has been developed for a patient to record a subset of relevant medical history in the surgeon's clinic. All three sets of information are then made available to the pre-op team the next day in a new triage tool in OSCAR, complemented by previous investigations, internal referrals, past pre-op assessment reports and clinical decision support (CDS). The CDS provides advice on managing the case as well as suggesting/recommending investigations based on the medical history provided. The pre-op team can then outcome the triage case, hold it for review or hold it for action. The outcomes are: 1) Routine pre-op; 2) Shorter pre-op – low risk case; 3) Early pre-op to optimise case; or 4) Reject case – back to GP/Surgeon.

Results.

Approximately 5% of surgery cases were cancelled at the pre-op stage in 2022. 716 of these cases were cancelled for clinical reasons. The top 6 clinical reasons for cancelling a case at pre-op can be seen in table 1. The triage tool has now been piloted for several weeks. Figure 1 shows part of the triage tool where the CDS displays predicted investigations and advice. The question mark button shows an explanation of how the CDS identified this investigation.

Table 1. Cancelled cases at pre-op in 2022. Top 6 clinical reasons.

Clinical reason	Number cancelled
Cardio & Exercise Tolerance	184
Diabetes	141
Respiratory	90
Hypertension	54
Blood tests, inc. anaemia	46
Frailty & Gerontology	40

Figure 1 CDS predicting tests

Difficulties have arisen with the availability of healthcare assistants to carry out the baseline observations and guide a patient in completing their medical history questionnaire while in the surgeon's clinic. There is also a shortage of staff in the pre-op team to triage the case.

Discussion

Diabetes, anaemia, hypertension and respiratory long-term conditions are recommended as the ones to investigate early in the NHS England document (2023b). Our data also showed that cardiac issues and exercise tolerance was the most common clinical reason for cancelling a case at the pre-op stage. An ECG could reveal such issues early on. There is no one trained to read an ECG in the surgeon's clinic and those ordering the investigation need to be the one following it up. The NHS England guidelines (2023b) also suggest that some patients may not require a pre-op assessment, if deemed fit early on. The pre-op team in Liverpool see this as contradictory since a patient can only be deemed fit once they have had a pre-op assessment, as can be seen in their choice of outcomes.

Conclusion.

Triaging at the decision to operate stage has already been shown to improve outcomes, reduce cancellations and provide a better patient experience at several places around the UK, e.g., Newcastle, Oxford, UCL and Southampton (NHS England 2023b). Finding the resources to obtain the information needed and carry out the triage is going to be a challenge. Automatically pulling in other past sources of information and complementing this with CDS should increase effectiveness in an efficient way.

Key references.

Murthy BV, Lake SP, Fisher AC. (2008) Evaluation of a decision support system to predict preoperative investigations*. Br J Anaesth.* 2008 Mar;100(3):315-21.

NHS England (2022) *Delivery plan for tackling the COVID-19 backlog of elective care*. Online at: https://www.england.nhs.uk/coronavirus/delivering-plan-for-tackling-the-covid-19-backlog-of-elective-care/
[Accessed: 8th July 2023]

NHS England (2023a) *2023/24 NHS Standard Contract*. Online at: https://www.england.nhs.uk/nhs-standard-contract/23-24/
[Accessed: 8th July 2023]

NHS England (2023b) *Earlier screening, risk assessment and health optimisation in perioperative pathways: guide for providers and integrated care boards*. Online at: https://www.england.nhs.uk/long-read/earlier-screening-risk-assessment-and-health-optimisation-in-perioperative-pathways/
[Accessed: 8th July 2023]

NHS England (2023c) *Decision Support Tools*. Online at: https://www.england.nhs.uk/personalisedcare/shared-decision-making/decision-support-tools/
[Accessed: 8th July 2023]

NICE (2016) *Routine preoperative tests for elective surgery NICE Guideline NG45*. Online at: https://www.nice.org.uk/guidance/ng45
[Accessed: 8th July 2023]

Construction of a Data Engineering Pipeline for Training a Clinical Risk Classification Model

[1]Luke Bird
[1]Medical Physics, The James Cook University Hospital, Middlesbrough, UK.
[1]Luke.Bird@nhs.net

Background. Despite the general drive towards digital transformation in the NHS, most data in healthcare remains stored in textual format. This represents a challenge for downstream analytics, particularly in machine learning (ML), where numerical inputs are required for training. Clinical texts must first be mined for the required clinical entities, engineered into a suitable format, and then aggregated before model training can begin. ML projects in the current healthcare data environment become particularly time-intensive at the pre-processing stage of development.

As part of an in-work apprenticeship, a project was designed to train and implement a classification model for pre-treatment risk stratification of patients who are prescribed pelvic radiotherapy. The first stage of this project was to design a data engineering pipeline for collating textual data from remote sources and then appropriately encoding and aggregating clinical entities into training, validation, and test sets.

Methods. Python 3.8 was used to build the data engineering pipeline. The natural language processing (NLP) libraries, *SpaCy* and *SciSpaCy*, formed the basis of the pipeline. Clinical context data were extracted from initial consultation reports (*Microsoft Word* 97-2003 documents), stored in a remote hospital data management system and initially converted to PDF format. This allowed for sections of text to be extracted using the inherent PDF coordinate system.

Pre-trained 'transformer' and 'named entity recognition' models were imported from the *SpaCy* and *HuggingFace* NLP libraries, and the *NegSpaCy* library was used to include sematic negation (e.g., absence of symptoms). These models were used alongside regular expressions to extract relevant clinical entities from each processed section of text. Specifically, these entities were: Patient age, ethnic origin, TNM staging codes, ICD-10 diagnosis codes, clinical comorbidities, past medical conditions, current medication, and frequency of alcohol consumption.

The *Metathesaurus* clinical vocabulary tool (developed by the US *National Institute for Health)* was used to map clinical entities to SNOMED-CT codes. This step standardised the training input to ML models and can be used to retain the semantic structure of the text.

Once collected, all entities were collated into a tabulated Data Frame object using the *Pandas* library, which can then be used to feed into the training of ML and deep learning models.

Results. This method was successful in extracting clinical entities from textual data and encoding them into standardised, ML-ready inputs. Due to the rigidity of regular expression rules, variations in textual formatting occasionally led to missed entity extraction. However, with the combination of NLP models, over 90% of clinical texts were successfully translated to complete model training inputs.

Discussion. Having developed a data engineering pipeline that can process textual data, the next step is to apply this to historical on-treatment comments written by treatment staff. Once extracted, it will form the basis of a project to generate and store patient pathways and experiences as knowledge graphs. These will enable further predictive models to be created in the future.

Conclusion. The establishment of an automatic data extraction and engineering pipeline allows current and historical clinical texts to be mined. It will allow ML-techniques to be implemented in future projects.

Case study: Challenges of providing a sustainable service in a radiotherapy department without a dedicated computing section

William Beasley, Jonny Lee

Background. In the radiotherapy department at University Hospitals Dorset (UHD) we have an array of different systems. IT is managed solely by Trust IT, who build, host and maintain our virtual servers. This reliance on Trust IT has some benefits but has indirectly resulted in a number of issues faced by the department in the past few years. This abstract seeks to highlight some of these issues and describe how we, as a medium-sized radiotherapy centre without a specialist scientific computing department, are working towards establishing a resilient service with respect to our IT systems.

Methods. Some of the problems encountered over the past few years are summarised below. 1) A Citrix environment was created for Mosaiq several years ago, which was subsequently used for Eclipse. Trust IT did not have anyone with Citrix expertise and the Citrix environment was not maintained for many years (no regular reboots, no hotfixes) to the point that Eclipse was becoming unstable. 2) We could not upgrade Eclipse to the latest version due to complex incompatibilities and interdependencies with operating system versions and other systems. 3) The amount of storage space on our Mosaiq servers required to store CBCT images is becoming unsustainable, with the nightly server snapshots soon starting to impact the clinical day. 4) Our surface-guided radiotherapy (SGRT) solution (AlignRT) depends on a physical server, which has been located in various positions around the hospital. Downtime on this server, caused as a direct result of its being a physical PC, has been experienced on several occasions; this has resulted in downtime of the radiotherapy service.

Results. 1) Fortuitously, Trust IT now have some in-house Citrix expertise, and have been able to help improve some of the performance problems we were experiencing. We were able to apply the latest Citrix hotfixes and establish weekly rebooting of the servers. Ultimately, we were running severely outdated servers (Windows Server 2008) and our problems were not solved until we upgraded Eclipse. 2) The upgrade of Eclipse addressed the performance issues, but it was a complex process that required upgrades of several servers and also a decommissioning of the Mosaiq Citrix environment, which was not compatible with the version of Citrix required for Eclipse. Interdependence of different systems is now something assessed and planned for proactively. 3) Approximately 30 GB of CBCT data are added to the Mosaiq server each week, and a request to IT to increase storage is required every month. The manufacturer has suggested that we migrate to the cloud, but this will not address the issue in the short/medium term. We are exploring options, including alternative storage architecture and in-house scripts for archiving data onto a separate storage location. 4) A physical server was originally purchased from the manufacturer, and IT were reluctant to create a virtual environment due to misunderstandings in the requirements. We are again working on virtualising the server, having had discussions with the IT department and allayed some of their worries; however, due to IT resourcing issues, it is likely to be many months before this will be performed.

Discussion. At UHD we have been working to resolve a number of issues related to historical IT-related setups. By engaging with IT and freeing up resources within the physics department to be able to dedicate time to these issues we have been able to work towards building a resilient and sustainable service. Establishing a good working relationship with Trust IT, as well as increasing knowledge of staff in the physics department, will help us achieve this aim. However, we are still vulnerable to resource availability within the Trust IT department itself.

Conclusion. When no dedicated computing department exists, sole reliance on Trust IT to provide the department's IT needs can result in unforeseen problems that can negatively impact the clinical service. In such a setup, having physicists with dedicated time and knowledge to actively work with IT to solve problems, helps provide a basic level of resilience that can enable a department to function efficiently.

perfusion2diagnosis: from research to clinical application, and back again

Neil O'Brien[1], James Leighs[1], Matt Guy[1,2], Livia Tossici-Bolt[2], Sofia Michopoulou[2], Chris Kipps[3]

[1] Scientific Computing (SciCom) | [2] Nuclear Medicine Physics (NMP) | [3] Neurology

UHS NHS Foundation Trust

Background. We routinely acquire 99mTc-HMPAO brain perfusion images from patients with cognitive issues [1], which assist in diagnosis and treatment planning. Nuclear Medicine Physics (NMP) was involved in a multi-centre research project, investigating methods to improve upon the standard visual-only interpretation of scan results, which gave rise to research software. SciCom was asked to help develop inhouse software for effectively and consistently applying the methods identified during this research, centred around comparing the patient's scan with a control database to identify areas of abnormality. No equivalent commercial software was identified.

Methods. SciCom developed both our in-house Software Quality Management System (SQMS) and perfusion2diagnosis (p2d) in parallel and each influenced the other. Particularly, the code that makes up p2d was mostly written by two developers, with every line having been peer-reviewed by at least one aside from the author, in dedicated git branches, before merging into the main branch. Significant testing was carried out, including automated unit testing within SciCom, then involving NMP, and finally clinical validation in collaboration with a Consultant Neurologist. Under the SQMS we documented requirements & risks, and produced a detailed user/developer guide, test plan & results logs. All documentation was reviewed and approved by relevant colleagues ahead of Top Management Review, after which the software was put into use, with surveillance and support in place via SQMS procedures. We developed an extensible framework for identifying relevant incoming DICOM images and invoking the desired processing software, before sending results to appropriate systems; NMP can now create clinical reports using our GUI to add annotations where required (including specifying an origin voxel in the scan), send the output to our system, and receive a detailed report within a few minutes.

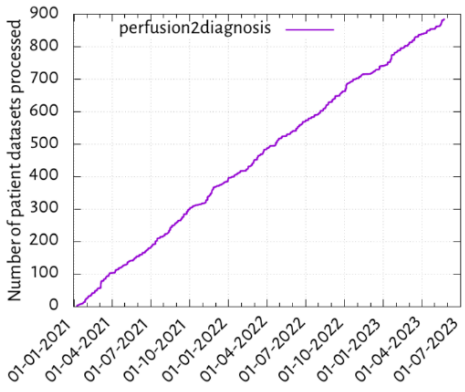

Results. From January 2021 to July 2023, p2d has produced around 900 individual clinical reports. Our support process has resulted in several minor improvements to the software, which have gone through significant in-house SQMS review prior to deployment, to improve legibility of report text, and add useful graphics to help NMP assess their inputs and the automated image registration outcome. Related research software is under development to investigate the applicability and usefulness of AI-based statistical analysis of segmented images, in support of NMP's research projects and with neurologist input, with a view to eventually integrate useful outputs into clinical reports.

Discussion. p2d was one of our larger medical device projects to date, due to the complexity of the analysis required, and was the first multi-developer project in SciCom. It saves time for our colleagues and provides consistency of analysis. Positive feedback has been received from NMP & clinicians using the system. We have had only a small number of brief failures to analyse, due to issues like power failures or lack of sufficient disk space on the system receiving the reports. Our close collaboration with NMP helps ensure our developments remain relevant and applicable. Our SQMS has provided confidence in our code and the results it generates, at the "cost" of requiring increased time and effort vs. writing code without these review and documentation requirements.

Conclusion. Close collaboration with colleagues across departments can help effectively apply development processes for software as a medical device, in this case providing clinical image analysis reports with minimal manual intervention, reducing required time per study and ensuring consistency of analysis.

Key references.[1] The impact of regional 99mTc-HMPAO single-photon-emission computed tomography (SPECT) imaging on clinician diagnostic confidence in a mixed cognitive impairment sample. A.M.J. Prosser, L. Tossici-Bolt, C.M. Kipps, https://doi.org/10.1016/j.crad.2020.04.016

Automated Clinical Audit of Planned vs Delivered Treatment Details

J.P.C. Baldwin, R. Wheller
Scientific Computing, Department of Medical Physics and Engineering,
Leeds Teaching Hospitals NHS Trust LS9 7TF

Background. In 2021, a previously unknown risk was identified in the radiotherapy pathway and investigated by Radiotherapy Physics [1]. When importing an RTPlan into the MOSAIQ Oncology Management System (OMS) it is possible to select a linac that is unable to deliver the plan. The error was identified when routine pre-treatment planning checks identified a flattering filter free (FFF) 6MV treatment being prepared as a flattened 6MV treatment [1]. Three root causes were identified [1]:

1) When there is a mismatch in between requested energy and what the linac can deliver, the user will be prompted to select a different energy/fluence type. This can lead to human error.

2) During checking of the parameters in the OMS, the mouse wheel is used to scroll down the page and this can change parameters when a drop down box is selected.

3) The third party commercial software used locally for plan checking does not detect errors in the FFF status.

The identification of this error triggered an audit into our previously delivered plans to identity if this had occurred previously.

Methods.

The overall project uses a combination of Python and SQL programming. An audit pipeline was set up to extract all RTPlans stored within the department's radiotherapy PACs archive via a DICOM SCP service. The Python script then extracts the required DICOM tags of each plan. This data is uploaded to an SQL server where it is compared to the data stored within MOSAIQ.

Initially the project was limited to a FFF and energy audit based on the identified risk discussed in the Background section. The outcome of the initially audit resulted in a request to expand the script to all data that is taken from the RTPLAN and stored in MOSAIQ. This will be used as an automated weekly, retrospective audit to ensure all checking processes are being carried out correctly.

Results. The initial audit was able to query over 74,000 RTPlans representing roughly 44,000 patients between 2012 and 2021. The audit identified a single patient who had received 6MV instead of 6MV FFF for a 40Gy in 15 fractions IMRT breast treatment. This resulted in a 25% increase in delivered dose to the target volume [1]. Due to the nature of the treatment the change in organ at risk dose was minimal [1]. The outcome of this audit triggered the need for an on-going, automated, audit tool.

Discussion. The initial clinical audit demonstrated the advantages of an automated audit pipeline. The patient's RTPlan represents only 0.0013% of all RTPlans, likely too small to be identified in a regular, manual, audit. Even though the harm for this incident was determined to be minimal, the patient required a follow up to ensure this minimal harm. This would not have been possible without being able to automatically process large data sets. We aim to leverage this automated audit tool for all planning parameters, especially where the current third-party commercial checking software is lacking.

Conclusion. Stage one of this project demonstrated the power of automated, big data, audits and we aim to extend this process to all planning parameters. We will present the methods and results of this extended process and highlight the issues encountered.

Key References

[1] UK Health Security Agency, "Safter Radiotherapy. Triannual RTE analysis and learning report Issue 36," UK Health Security Agency, London, 2021.

A Quality Management System for Radiotherapy TPS Scripts

Matt Ward (Medical Physics Technician), Laura Smith (Clinical Scientist), Virgiliu Craciun (Head of Scientific Computing), University Hospital Southampton NHS Foundation Trust

Key words: *radiotherapy, automation, quality management system, scripting*

Background: Automation of treatment planning workflows in radiotherapy is a major focus in our department to improve efficiency, consistency and ultimately to provide the best plan quality for our patients. Most radiotherapy treatment planning systems (TPSs) provide scripting functionality, allowing for automation of some repetitive, time-consuming aspects of the planning workflow. These scripts may exist as sections of code in a proprietary language such as Philips Pinnacle^3 or as files in an open-source language such as Python (Raysearch Laboratories, Raystation). On a basic level, users can "record" a script which executes a series of pre-recorded interactions with the user interface (UI). The scope of these recorded scripts is often quite limited, prompting those who need more advanced functionality, i.e., iterating over loops, to edit the code outside of the relative safety of the TPS. The aim of this project was to adopt a more robust script development life cycle and to implement a bespoke in-house quality management system (QMS), in line with new medical software guidelines [1] and the ISO 9001:2015 standard.

Methods: *Project Management:* A "Scripting Portal" was established in Redmine (our internal open-source project management system) where users are assigned specific roles: user, developer, tester, or manager. "Tracked workflows" were designed using the issues tracking system, which breaks the script development into stages to encourage an agile software development cycle and details the script's priority, category (e.g., quality assurance tool) and treatment site. QMS documentation was developed and tailored to these workflows – a script risk assessment and a testing and review form. Together the documentation and tracked workflow ticket provides an audit trail from the initial script request, through the risk assessment approval and testing and validation process.
Technical: Existing scripts from our TPS (Philips, Pinnacle^3) were copied onto a remote server which was initialised as a Git repository and a categorisation of scripts based on their function was performed. A secondary remote was established on GitHub, allowing users to clone the repository onto their local drive for development. Each new script/bug fix is developed on a feature git branch, and the associated documentation is saved in the repository alongside the script files. This allows the documentation to be visible by everyone by mirroring the repository directly to the Scripting Portal. Each developer has access to Visual Studio Code with ftp-sync extension; this automatically syncs any local development work to a user-specific directory on the TPS server for on-the-fly testing. This directory can only be accessed via specific credentials, ensuring that scripts in the development stage cannot be used on clinical patients. Once the tracked workflow process has been followed, the script manager is responsible for final approval and transfer of scripts into the clinical script directory. Following final verification testing on the clinical system, the script is deployed for clinical use, the feature branch is merged with the master branch, and the request in the Scripting Portal ticked status is set to Sorted.

Results/Discussion: Since the implementation of the script QMS in December 2022, we have successfully followed the workflow for the development and deployment of 21 script requests. Due to several clinical protocol updates and changes to naming conventions, a vast number of script updates were required in a short space of time. Without the scripting QMS in place, it would have been much more challenging to implement these changes in a safe and timely manner.
The Scripting Portal is also a central hub for education - a comprehensive Wiki page provides resources for new developers, such as a syntax help guide and details on why the QMS has been established to meet the guidelines. Future work includes creating interactive tutorial videos and script worked examples to enable a wider staff group to safely contribute to TPS scripting.

Conclusion: We have implemented an in-house QMS for TPS scripting which has facilitated the timely and safe development of in-house scripts for automation of treatment planning workflows.

Key references: [1] IPEM (2022) *Best-practice guidance for the in-house manufacture of medical devices and non-medical devices, including software in both cases, for use within the same health institution. Version 2.0* Available at: https://www.ipem.ac.uk/media/vp0ewy01/ipembe-1.pdf (Accessed: 10 July 2023).

FSynapse: A Software Toolbox to Support Development of Multimodal Clinical Decision Support Tools for Falls and Syncope Assessment

Ivan Lifanov, Laura Perez Denia, Ciara Rice, Gerard Boyle, Rose Anne Kenny, Ciarán Finucane

Introduction

Falls and syncope affect 1 in 3 older adults aged over 65 annually, with unexplained fallers five times more likely to suffer recurrent falls [1]. Active Standing (AS), Head-up Tilting (HUT) and Carotid Sinus Massage (CSM) with continuous beat-to-beat blood pressure (BP) monitoring are recommended as diagnostic tests for syncope and falls [2]. Near infrared spectroscopy (NIRS) and other complementary multimodal physiological measures are novel tools of emerging clinical importance in this field. [3, 4]. However, there is currently a lack of suitable analytical tools to support analysis of this multimodal data in clinical research and practice. Here we describe the design, verification and validation (V&V) of a novel analytical framework for joint analysis of BP and NIRS data building upon our previous work in this area [5].

Methods

A software toolbox, FSynapse, was designed to meet predefined user requirements including file linkage, and synchronization, signal visualisation and quality assessment, signal filtering and extraction of clinically relevant features. FSynapse was also designed to enable flexible data management to support data visualisation, bio statistical analysis and future AI based clinical decision support tools. The system was developed in MATLAB R2021b and designed to cater for thousands of patient records. After completing an agile based rapid prototyping development cycle, the system architecture was subject to a verification process and then validated on a large real-world clinical dataset.

Results and Discussion

FSynapse is organized into 3 functional steps, which are implemented as individual scripts referencing common toolbox functions, with the overall execution flow governed by a master command script. Firstly, the user launches the automated file extraction step. Secondly, files are semi-automatically linked and synchronized. Finally, the signal quality is assessed, signals are pre-processed, features and clinical definitions are extracted. The outputs generate includes signal data and feature sets for each test, which can facilitate further analysis and A.I. applications, summary and visualizations. After successfully passing V&V tests, FSynapse was used to process available test recordings. A large sample of multimodal AS, HUT and CSM files was processed successfully with reported features and clinical definitions generated. Analysis of this multimodal data is now being used to create machine learning based phenotypes to support the application of NIRS and clinical decision support tools in syncope and falls diagnostics.

Conclusion

FSynapse is a promising tool for multimodal data processing, which can facilitate clinical and research data curation, biomarkers discovery, and AI based clinical decision support.

Key references

1. Bhangu, J., et al., *The prevalence of unexplained falls and syncope in older adults presenting to an Irish urban emergency department.* Eur J Emerg Med, 2019. **26**(2): p. 100-104.
2. Brignole, M., et al., *2018 ESC Guidelines for the diagnosis and management of syncope.* European Heart Journal, 2018. **39**(21): p. 1883-1948.
3. Pérez-Denia, L., et al., *Increased multimorbidity is associated with impaired cerebral and peripheral hemodynamic stabilization during active standing.* J Am Geriatr Soc, 2022. **70**(7): p. 1973-1986.
4. Mol, A., et al., *Multimodal Monitoring of Cardiovascular Responses to Postural Changes.* Frontiers in Physiology, 2020. **11**.
5. Soraghan, C.J., et al. *TILDA Signal Processing Framework (SPF) for the analysis of BP responses to standing in epidemiological and clinical studies*. in *IEEE-EMBS International Conference on Biomedical and Health Informatics (BHI)*. 2014.

Developing a vendor neutral cardiac implant device checker for MR safety queries

J Moggridge[1], A Papadaki[1], J Thornton[1]
[1]Neuroradiology, National Hospital for Neurology and Neurosurgery, UK

Background. In the protocolling process for MR scans at the National Hospital for Neurology and Neurosurgery, patient Cardiac Active Devices (CADs) such as pacemakers and ICDs are checked to determine whether they are MR Conditional and what scan conditions need to be applied to ensure patient safety. Both device and leads need to be MR conditional with some combinations limited to 1.5T field strength or having an exclusion zone for scanning. To facilitate this, CAD manufacturers provide web or pdf-based resources with information on the MR conditions in which a device can be safely scanned. Navigating these resources can be complex and time consuming. Having a single tool that could be used to search for a CAD model from any manufacturer and return information required to scan safely could simplify the process and improve efficiency.

Methods. A combination of manual and automated techniques was used to build a database of CAD models, associated leads and the location of MR conformance resources provided by the manufacturer. Conformance information from pdf resources was entered into the database manually and for web-based resources, a web application using Flask[1] and Selenium[2] was developed to scrape equivalent data from the relevant web page.

The web application was further developed to provide a user interface for checking a device model number along with the model numbers of up to 3 associated leads. When a check is triggered, the application searches the internal database for the model(s) to return the respective resource needed to confirm MR compliance. If the resource is web-based, Selenium is used to access the site, complete any text or selection inputs using the provided details, trigger a check before scraping the MR scan condition results. These results are returned to the end user as a new page on the web application along with a screen shot of the web site as evidence that the check was successful. If the device compliance is described in a pdf, the application looks for the relevant information within the internal database and returns condition textual information along with a screen shot of the relevant part of the pdf.

Results. The application is being tested as the first check for implant safety when protocolling patients' that have CADs. Feedback from the MR radiographers has been positive but there has been one case where the application has been unable to find the device due to it being missing from the database as a US resource was used and another where the device was found and, although MR Conditional, was displayed as non-MR Conditional due to the application incorrectly reading the number of associated leads from the designated resource.

Discussion. This project presented multiple challenges: firstly the web resources provided by the manufacturer were difficult to navigate automatically due to a lack of API and the use of dynamic web page display techniques where data and web elements would only become accessible once inputs to the relevant part of the web form had been entered; secondly the contents and information provided in output statements relating to MR scan compliance varied between manufacturers making it hard to display a standardised check result. These challenges mirror those faced by an individual looking to run their own conformance check on the resources themselves and hence highlight a need for a single, standardised resource.

Conclusion. User feedback suggests that the tool does save time, but further work is required to ensure that the application's database in complete (including foreign device models) and that all data from a manufacturer check is interpreted and displayed correctly along with a mechanism to ensure that resources are regularly checked for updates. This project indicates a need for a single site on which all manufacturers of implant devices can upload their devices' MR conditions making it easier to check compliance and ultimately improve patient safety.

Keywords: application development, web scraping, cardiac implant device, MR safety

Key references.

1. Flask {https://flask.palletsprojects.com/en/2.3.x/} Jul 2023

2. Selenium {https://www.selenium.dev/} Jul 2023

A dictionary-driven pipeline for the calculation of clinical key performance indicators

[1]Man Him Law, [1]Neide Simões-Capela, [1]Daria Cosentino, [1]Timothy Cross
[1]*Clinical Informatics – Clinical Operations Group, HCA Healthcare UK*

Background. Clinical key performance indicators (KPIs) [1][2] are used to monitor performance and ensure clinical excellence in hospitals. A straightforward scripting approach for KPI calculations would be to hard-code their calculation criteria within each script. However, this leads to repetitive coding, suboptimal runtime, and difficulties with tracking KPI definition changes across scripts which may result in reporting conflicting information across products. A dictionary-driven approach that calculates the KPI cube [3] stored in a central place can tackle all these issues.

Methods. Our dictionary-driven pipeline includes 3 main pieces: 1. the raw data sources stored across multiple tables in a SQL database [4], containing the data feeds coming from a plethora of clinical systems; 2. our KPI dictionary, stored in the same database; and 3. an R [5] script that is triggered on a schedule to update the KPI cube (stored in another SQL table). The dictionary table includes all the information necessary for the calculation of the KPIs (i.e., name of the source table, filters to be applied, dimensions used for the required aggregation). The script loops across the KPIs in the dictionary to load the data source in the R working environment, apply the required transformations and perform the final aggregation using the definition and the aggregation dimensions specified in the dictionary itself (as shown in Fig.1). For example, to calculate the KPI "Patient Falls" defined as the number of patient falls per 1000 bed days, the script will: 1. pull incidents in the last 2 years from the back-end table of the relevant incidents reporting system; 2. use the KPI definitions from the dictionary (*IncType == 'PAT' & RepApproved != 'REJECT' & IncSubCategory == 'FALLS'*) to filter the data; 3. aggregate the data by the month of the incident reported and by facility (as indicated in the relevant dictionary columns); and finally 4. divide the count of incidents by the number of bed days x1000.

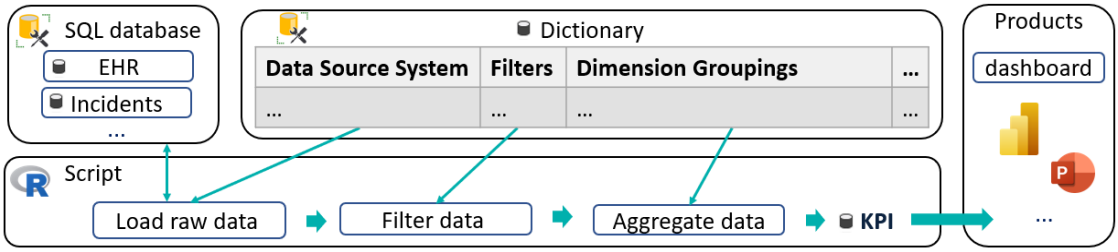

Figure 1. A simplified schema of the dictionary-driven pipeline.

Results. Having a central unique KPI dictionary and cube removes the chances of conflicting KPI values displayed across different products. The dictionary keeps up-to-date documentation of the KPI definitions that can be easily shared with our end-users without incurring the risk of showing outdated or conflicting definitions. The dictionary is proved to be easily scalable, allowing analysts to define new KPIs and thus to expand the cube with little overhead. The dictionary also allows for the calculation of the same KPI aggregated using different dimensions (e.g., at facility, ward, or consultant level) and/or at different timeframes (e.g., monthly, quarterly, or yearly).

Conclusion. The dictionary-driven KPI calculation pipeline allows us to increase the frequency of re-calculations of critical KPIs, to quickly add in new KPIs on demand (e.g., COVID-19 related KPIs), and to ensure the robustness of our KPI calculations across products.

Key Words: *Key Performance Indicators, R, SQL*

Key References.

1. Doucette, D. and Millin, B. *Should key performance indicators for clinical services be mandatory?* Can J Hosp Pharm. 2011 Jan;64(1):55-7. doi: 10.4212/cjhp.v64i1.989.
2. R. Latinovic, "Key performance indicators: What are they all about?," PHE Screening, https://phescreening.blog.gov.uk/2015/07/10/key-performance-indicators-what-are-they-all-about/ (accessed Jul. 7, 2023).
3. Gray, J., Chaudhuri, S., Bosworth, A. *et al.* Data Cube: A Relational Aggregation Operator Generalizing Group-By, Cross-Tab, and Sub-Totals. *Data Mining and Knowledge Discovery* 1, 29–53 (1997). https://doi.org/10.1023/A:1009726021843
4. Silva, Yasin & Almeida, Isadora & Queiroz, Michell. (2016). SQL: From Traditional Databases to Big Data. 413-418. 10.1145/2839509.2844560.
5. RStudio Team (2020). RStudio: Integrated Development for R. RStudio, PBC, Boston, MA URL http://www.rstudio.com/

In-House Software Development in Radiotherapy - Lessons Learnt

M Tyyger[1]

1. Leeds Cancer Centre, Leeds Teaching Hospitals NHS Trust, Leeds, UK

Background: Developing in-house software in Clinical Scientific Computing (CSC) requires collaboration with a variety of departments and specialisms, all of which can have differing challenges and nuances. However, common themes can be identified which support project success regardless of the specific specialism involved. This work discusses lessons learnt from four years of software development with Radiotherapy (RT) Physics.

Methods: Fifteen clinical in-house software development projects between 2019 and 2023 were reviewed. Thirteen projects were based inside a treatment planning system (TPS) using its application programming interface (API), and two projects were stand-alone software.

Results: Six projects were clinically released, with three enabling clinical pathways. Six projects were under development, mostly small projects automating repetitive tasks. Three projects were ultimately cancelled. Of these, two were ready for clinical deployment but were never used due to changes in clinical practice, and thus differing software requirements. The final cancelled project was halted midway through development due to rapidly increasing project complexity and scope.

Discussion: Three common themes were identified through post project review.

Successful projects tended to be small scale with very specific use cases and requirements, although one large clinical pathway defining project was deployed. However, typically large scale projects were cancelled, sometimes even after software development was finished. Poor project choice was ultimately the issue. The largest potential clinical impact was desired; however, this meant broad project scopes involving multiple complex treatment pathways and significant staff resources. These projects were susceptible to requirement changes and increased resource demand due to: clinical pathway nuances, unexpected API behaviour, and other unforeseen roadblocks. One lesson learnt is to dissuade groups from initiating large scale projects until there is certainty of team capacity, expertise, API knowledge, and software development lifecycles. Instead proposing the initial idea is deconstructed into smaller independent projects, including separate prototypes. Smaller projects can achieve clinical benefit sooner, and prototyping enables thorough evaluation of complexity, workload, and potential roadblocks.

Staff workload meant it was not always possible to assign experienced staff members to a project with large time commitments. Instead, multiple inexperienced staff members were provided, each with a small time commitment. Large inexperienced programming teams drastically increased project management overhead and caused significant delays to development. The lesson learnt is to not equate a wide inexperienced team to a narrow experienced one, even when the total staff time commitments are equal.

Long-term success was reliant on developing staff expertise, with CSC staff gaining clinical experience and clinical staff gaining understanding of software development. This helped staff to communicate in the same language and understand the challenges and risks each side were facing. However, giving inexperienced programmers suitable clinical projects whilst still achieving deadlines was challenging to manage. The local solution is a “toolkit”, developed by experienced staff, which manages boiler-plate programming such as logging and error management. This provides a shallower learning curve for safe software development, allows staff to focus on easier individual tasks, and reduces the overall training workload.

Conclusion: A range of collaborative development projects have been attempted locally with RT physics to varying degrees of success. Reviewing these projects has concluded that successful in-house software development requires: selecting the correct project by identifying the expertise and capacity of staff, starting small when uncertain, developing the skillset of multidisciplinary staff groups long-term, and not simply using large teams of inexperienced staff when experienced staff are unavailable.

Key references: None