Medical Physics and Statistical Science
Exploring Interfaces and Building Collaborations

Fields Institute, April 4 -5 2017

ABSTRACTS

Day I

9:30 – 10:30: Medical Physics: Recent Developments

Parminder S. Basran, PhD, FCCPM
BC Cancer Agency, Victoria BC

Medical Physics and Statistical Sciences: Exploring Interfaces within Radiation Oncology

Two in five Canadians will develop cancer in their lifetime. While this is a worrisome statistic, the Canadian Cancer Society (2016) reports a steady decline in cancer mortality over the last three decades. Much of this positive news has been attributed to our better understanding of the causes, development, and treatments of cancer, fueled through advances in technology and the basic sciences. About half of all cancer patients receive therapeutic radiation for the purposes of curing the cancer, amplifying tumor response, reducing the risk of recurrence, cancer migration or side-effects, and the palliation of symptoms. Radiation oncology is essentially a geometrically and physically constrained problem of depositing therapeutic levels of radiation to a “clinically” defined volume, while simultaneously minimizing the radiation received to healthy functioning tissues. The increase in efficacy of radiation oncology observed within the last few decades has been largely due to advances in technology, computing power, coupled with improvements and integration of health information systems in cancer care. Consequently, more than any other cancer treatment modality, radiation oncology is inherently a ‘data-driven’ and ‘data-rich’ science.

Medical physicists play a vital role in the practice of radiation oncology through implementing and executing processes which ensure the safe delivery of radiation to cancer patients. They are also often involved in research and education in academic and clinical environments. Despite the fact that radiation oncology is ‘data-rich’, it can also be ‘information-poor”: the types of numerical tools and processes used by medical physicists remain quite basic, and the metrics used in assuring the quality of cancer care may not always be entirely relevant or useful.

The purpose of this presentation is to briefly introduce the cancer “problem”, present some of the challenges faced in radiation oncology from a medical physicist’s perspective, and explore the types of data problems faced in radiation oncology.

Frank Prato, PhD, FCCPM, ABMP, FCOMP
Lawson Health Research Institute, University of Western Ontario, and St. Joseph’s Health Care and London Health Sciences Centre

Implementation Hurdles of New Medical Imaging Technologies in Research, Clinical Trials and Patient Care

I will draw from my experience of introducing new imaging modalities into Canada which have included Bone-Mineral Density (BMD; 1981), Magnetic Resonance Imaging (MRI; 1982), Positron Emission Tomography with X-ray CT (PET/CT; 2002) and hybrid Positron Emission Tomography with MRI (hybrid PET/MRI; 2012).

Introduction into Basic Research: This includes challenges in maintaining state-of-the-art in an imaging modality that is experiencing fast growing technical development. For example in the first decade after the introduction of MRI there was still no consensus on field strength (0.15T to 1.5T) and there was fast paced major development in other MRI instrumentation. This was further complicated by the issue of safety of exposure to a new imaging procedure needed to define risk. Research associated with the impact on disease understanding associated with visualization of tissue characteristics never seen before was a major area of research which often took years to achieve consensus.

Introduction into Multi-Centre Clinical Trials: Major clinical trials often require hundreds of sites. For example the Bayer GadaCAD 2 study which evaluated the use of a MRI contrast agent in heart disease required 25 active sites with many more approached. The initial ADNI (Alzheimer Disease Neuro Initiative) trial included some 60 sites. Even though these studies were dependent on primary outcomes associated with MRI and PET imaging biomarkers, rather than just clinical outcomes, they have still been extraordinarily expensive. Significant effort is needed to investigate ways of reducing costs through effective statistical analysis. For example moving such studies from separate PET/CT and MRI studies to a single hybrid PET/MRI imaging session would reduce costs by a) increasing initial recruitment rates and reducing subject dropout by reducing the number of subject’s visits and improving subject safety and b) reduce the needed number of subjects by improving the accuracy and precision of the results. However big pharma studies are still restricted to PET/CT plus MRI rather than hybrid PET/MRI due to the slow adoption of a new modality. This has been somewhat related to the similar number of hybrid PET/MRI sites but more so to the concern regarding the number and technical misunderstandings that could be mitigated by strong statistical analysis quantitating benefits.

Introduction into Patient Care: Even when a new imaging modality is clearly superior in sensitivity and specificity it is often delayed and/or denied provincial reimbursement due to: a) lack of appropriate analysis of cost which should include, not only direct costs, but indirect costs such as costs of the wrong treatment for the

wrong disease, b) the failure to appreciate the impact on caregiver of the correct diagnosis even when effective treatment is not available such as discrimination between dementia Lewy Body from Parkinson's Disease and c) issues of medical specialty "ownership" when converging technologies such as hybrid PET/MRI are evolving.

Summary: These examples are a few of the many hurdles of implementation that could benefit from an effective collaboration between Medical Physicists and Statistical Scientists in the rapidly evolving field of Medical Imaging.

10:45 – 11:45: Dose Response; Simulation of Stochastic Models

Tim Ramsay, Ph.D.
Ottawa Hospital Research Institute
School of Epidemiology, Public Health and Preventative Medicine, University of Ottawa

A Phase I Time-to-Event Continual Reassessment Methodology (TITE-CRM) Dose Escalation Trial of CyberKnife® Stereotactic Body Radiotherapy (SBRT) Boost for Patients with Pancreatic Cancer

This goal of this talk will be to present an example of the use of relatively new design for a dose escalation study in which dose-limiting toxicity may occur many months after the experimental treatment is administered. In this example, toxicity is expected to occur between 3 and 9 months after the treatment. A normal continual reassessment design would require waiting at least 9 months between patients. Even if the target sample size were as small as 9 patents, this would mean that it would take at least 81 months, or almost seven years, to complete a Phase I trial. With the TITE-CRM design, one can recruit a new patient at any time, while still using all of the information available from previous patients who may not have completed the full nine-month follow-up period.

Patrick Brown, Ph.D.
Department of Statistical Sciences, University of Toronto

Modelling and Simulating Spatial and Spatio-Temporal processes

The 'Curse of Dimensionality' is a name used to describe the phenomenon that two-dimensional spatial models are an order of magnitude more complex to work with than one dimensional time series models, and three dimensional spatio-temporal processes are yet another further level of

difficulty. A 20 by 20 spatial image evaluated at 20 time points involves 8000 individual values and an 8000 by 8000 variance matrix (with 64 million entries). The talk will describe some basic spatial and spatio-temporal random processes and methods for overcoming the curse of dimensionality, including Markov Random field approximations and Spectral representations.

1:30-2:30 Medical Physics: Data Analytics in Radiation Therapy

Stephen Breen, Ph.D., MCCPM
Radiation Medicine Program, Princess Margaret Cancer Centre, Toronto
Department of Radiation Oncology, University of Toronto, Toronto

Applications of Statistical Process Control in Radiation Medicine

The practice of clinical medical physics has a long history of quality control. Measurements of operating parameters of cobalt units, medical linear accelerators, and brachytherapy devices have been made daily by medical physicists since the clinical application of these technologies. Historically, many of these measurements have been treated as isolated values, and the opportunity for meaningful quality improvement was lost. To maintain a consistent level of quality, and to develop a deeper understanding of the performance of clinical processes, we have started to use statistical process control, and have explored different types of control charts for varied applications. Control charts demonstrate the mean of a process and the upper and lower limits of expected variation. Variation beyond these limits is said to be “out-of-control”. Control charts are used to differentiate common cause (i.e. random) variation from variations due to assignable causes.

Patient-specific dosimetry is a common practice in the delivery of intensity-modulated radiotherapy and rotational treatments to verify the fidelity of planned and delivered doses. We have tracked our IMRT measurements on control charts to identify opportunities for improved beam models in our planning system.

Many features of contemporary radiotherapy - the multiplicity of tasks, the number of information hand-offs, the opportunities for work-arounds, and the multiplicity of designs of human-computer interfaces - produce numerous opportunities for patient safety incidents. For the last five years we have tracked the quarterly frequency of (1) patient safety incidents, (2) harmful and no-harm incidents, and (3) near misses. We have attempted to understand the causes of variation using statistical process control by using control charts to plot the number of events, the monthly event rate, and the intervals between consecutive events.

Lastly, we have retrospectively used yet another type of control chart – a multivariate chart – to predict the quality of LDR implants. By measuring several parameters of the implant, we may be able to predict the likelihood of edema after

implant. Monitoring several metrics within our program with SPC has allowed us to identify out-of-control events within our treatment delivery, brachytherapy, and incident learning programs, and provides a basis for improved radiotherapy and improved patient care.

Rowan Thomson, Ph.D.
Physics Department, Carleton University

Application of advanced patient-specific Monte Carlo dose calculations for brachytherapy

Calculations of radiation dose are critical in radiotherapy for treatment planning, as well as evaluating and understanding treatment outcomes. In recent years, there has been much progress in the development of advanced model-based dose calculations such as those employing Monte Carlo (MC) techniques. This presentation will focus on applications of fast MC codes for brachytherapy, which are radiotherapy treatments characterized by source implantation within or directly adjacent to the tumour. Dose-volume metric and radiobiological index analyses will be discussed for different treatment sites and patient cohorts. Ongoing challenges in evaluating patient treatments and future directions will be discussed.

3:00 – 4:00 Statistical Modelling

C. Devon Lin, Ph.D.
Department of Mathematics and Statistics, Queen's University

What are computer experiments, and how do we collect and analyze their data?

Many real-world phenomena like climate change and the spread of infectious disease can be simulated using numerical methods such as finite element analysis. With the help of numerical methods and physical models, computer experiments refer to those experiments that are performed using computers. They are efficient and cost-effective surrogates of physical experiments. In this talk, basic components, some common goals, and different types of computer experiments will be described. The general framework of design and modeling of computer experiments will be reviewed. An illustrative example will also be given.

Peter Kim, Ph.D.
Department of Mathematics and Statistics, University of Guelph
(Joint work with **Christine Lee**, Royal Jubilee Hospital (Victoria) and University of British Columbia)

The Gut-Brain Axis and Clostridium difficile Infection
The gut-brain axis describes the physiological connection between the gastrointestinal tract and the brain. There is strong evidence that the every

day function of the brain, hence the central nervous system, can be altered by modulating the gut microbiota. Although this complex biochemical connection would be difficult to quantify, there has been some progress in at least collecting the data that may potentially provide some quantifiable connection. At the Faculty of Health Sciences, McMaster University, a clinical trial was recently completed involving the hyper-virulent Clostridium difficile infection (CDI). In this clinical trial, an experimental treatment called fecal microbiota transplantation (FMT) was used to treat patients, where FMT provided strong statistical evidence of its efficacy in terms of clinical resolution.
A clinical observation encountered, was that patients felt almost immediate relief. To shed some light, a healthy donors stool is procured, processed and passed onto the patient as an enema. The microbiological process is that the gut microbiome of a healthy donor is engrafted into the gastrointestinal tract of the diseased patient thus altering the patients microbiome. Although CDI is not thought to be a neurological disease some valuable quantifiable information concerning the gut-brain axis could be revealed in the data. This could provide a basis for diseases of the central nervous system such as autism, multiple schlerosis, Parkinson's disease, to name a few.

Gut: The clinical data consisting of CDI resolution following an FMT(s) has been recorded. In brief following the last FMT, if the patients CDI symptom does not recur following 7 days following their last FMT, that patient is deemed clinically cured. As part of the protocol every patient was monitored by the research staff. From this group a subset of patients stool samples were sequenced at four time points: pre-FMT; followed by day 10, week 5 and week 13 following their last FMT. A donor stool provided the material for the FMT and we paired the donors stool with the corresponding patients pre-FMT stool sample. The donor and pre-FMT CDI patient, along with the three followups, were simultaneously sequenced using the Illumina MiSeq platform. We note that not all FMTs were initially successful hence of particular interest is to try and understand the clinical outcome using metagenomic (bioinformatic) variables as covariates. Metagenomics bears a structural relationship between covariates. Bacteria exhibit a treelike relationship with each other. Thus their operational taxonomic unit (OTU) proxies, typically the 16S rRNA gene, also exhibit a treelike structure with patterns of cycles at deep taxon classification. Our resolution depends on the length of the 16S rRNA region, the degree of lateral gene transfer, and the reliability of the reads.In addressing any research question, particularly in terms of modelling, we require a means of selecting the OTUs having dominating roles in the microbial systems at hand. The relationship between OTUs is such that some of them may represent the same type of bacteria. Alternatively, some spurious OTUs are artifacts of the laboratory sequencing protocol. Thus, rather than selecting OTUs on their individual merits, we want to select them based on their group affiliations using machine learning algorithms which was developed. Some key genera were identified.

Brain: The main objective of the CDI trial was to study the efficacy of FMT as a treatment alternative. Data was recorded at pre-FMT and post-FMT followups. In particular patients provided stool samples at their prescribed clinical followups. Ideally an MRI scan at each time point would have been superb to capture neurological changes as a result of altering the patients gut microbiota. However, this was not done and would have been difficult to do due to the number of patients involved and the timeline between hospital visits. Nevertheless all patients filled a health survey at each hospital visit using a well established health survey (SF-36) with the attempt to capture neurological changes as patients progressed through this experimental FMT treatment.
The SF-36 is a well established 36-item questionaire which groups the responses into eight components: Physical Functioning; Role Physical; Bodily Pain; General Health; Energy/Vitality; Social Functioning; Role Emotional; and Mental Health. The findings were somewhat dramatic in that pre-FMT patient scores were uniformly lower than that of the general Canadian population across all eight components. This questionnaire was subsequently repeated in the post-FMT CDI patients at followups. Neurological research has shown that some components (Bodily Pain, Physical Functioning, Social Functioning and Mental Health) are directly linked to different regions of the brain. One of the eight survey components investigated was Bodily Pain which is associated with the amygdaloid complex (amygdala). It has also been implicated in clinical disorders such as depression and anxiety, both of which impact sense of well-being. Using this as a proxy to an MRI scan, we were able to find interesting connections to the identified genera from the gut.

Day II

9:00 – 10:00 Medical Physics: Big Data in Imaging and Diagnosis

Anne L. Martel, PhD
Department of Medical Biophysics, University of Toronto
Imaging Research, Sunnybrook Research Institute

Deep learning: a universal tool for medical image analysis?

The availability of very large databases of labeled images, new methods for training deep neural networks and the increased availability of GPUs, has allowed computer scientists to make very rapid advances in the field of machine

learning in the last decade. Convolutional neural networks (CNNs) have been shown to outperform more traditional computer vision approaches in a wide range of applications and this has led to great interest in the use of deep learning for medical image analysis. In this talk I will give a brief introduction to CNNs and discuss their strengths and weaknesses for both segmentation and classification in the medical domain where datasets are generally very small compared to those used in other domains.

Tom Purdie, PhD, MCCPM
Department of Radiation Oncology, University of Toronto
Princess Margaret Hospital/University Health Network

Automation in Radiotherapy

The delivery of radiation for the treatment of cancer is a complicated process that requires both clinical and technical expertise to ensure radiation treatments are safe and effective. Sub-optimal treatment plans have the potential to result in significant detriment to the patient and several studies have shown radiotherapy plans, which deviate from established clinical guidelines, result in worse patient outcomes. Therefore, the current radiotherapy process requires substantial multi-disciplinary resources to both generate and verify radiotherapy plans are of high-quality.

In this talk, a previously validated machine learning platform customized for radiotherapy will be presented. The method automatically learns based on data from thousands of previously treated patients which relationships and patterns in radiotherapy image and radiotherapy plan data are best for deciding where dose should be placed and how dose should be delivered in radiotherapy plans without requiring any manual intervention. Therefore, the method can be used to both generate personalized radiotherapy plans and to quantitatively score radiotherapy plans for quality and classify plans that have errors.

The automated platform can readily be integrated into current clinical process to improve efficiency in the treatment planning and radiotherapy plan review process and to better utilize the vast data we have to ensure we are providing patients with highly personalized radiotherapy treatments.

10:15 – 11:15 Statistical Imaging and Complex Data

Joern Diedrichsen, Ph. D.
Department of Statistical and Actuarial Sciences, Western University

Analyzing functional activity in the human brain: What is the correct level of description?

Functional Magnetic Resonance Imaging (fMRI) provides an unprecedented opportunity to observe activity in the entire human brain at high spatial resolutions. How to understand the complex activity pattern that can be observed on the level of individual subjects, and how to relate them to behavioral or clinical characteristics of that person, however, is a matter of ongoing debate. Traditionally, researchers have normalized individuals brain spatially into a group template, and then compared the local activation in different brain regions across participants. This approach, however, neglects the substantial functional variability with which different brains are organized – with the exact spatial layout on the cortical sheet likely reflecting (unimportant) biological variation. I will present some new approaches from our lab to understand the structure of cortical and cerebellar activity patterns independent of their detailed spatial organization.

Farouk Nathoo, Ph.D.
Department of Mathematics and Statistics, University of Victoria

A Potts Mixture Spatiotemporal Joint Model for Combined MEG and EEG Data

We consider the problem arising when MEG and/or EEG are used to measure electromagnetic brain activity over an array of sensors at the scalp and it is of interest to map these data back to the sources of neural activity within the brain. This inverse problem is ill-posed and requires some assumptions or the use of priors in formulating a solution. The physics governing the forward problem assumes the subject's head to be a volume conductor with a nonuniform anisotropic conductivity tensor, and the electromagnetic field in the media is described by the system of Maxwell equations. We review some of the existing approaches to solving this inverse problem and discuss the mesostate-space model (MSM) proposed by Daunizeau and Friston (Neuroimage, 2007). We then propose a new Bayesian model that builds on the MSM and incorporates three major extensions: (i) We combine EEG and MEG data together and formulate a joint model for source reconstruction; (ii) we incorporate the Potts model to represent the spatial dependence in an allocation process that partitions the brain into a small number of latent mesostates; (iii) we formulate the mesostate dynamics in a manner that can represent the effective connectivity between mesosources. We formulate the model, discuss practical computational implementation, and provide some examples of its use. Joint work with PhD student **Yin Song**.

1:30 – 2:00 Statistical Learning

Irene Vrbik, Ph.D.

Unit 5 – IKBSAS, University of British Columbia, Kelowna, BC

Topics in Model-based Clustering and Classification

The exploration and extraction of meaningful patterns in data is an important research goal shared across many disciplines. Model-based methods offer a mathematically sound framework for finding implicit structures within data in an automatic or semi-automatic way. This talk discusses flexible model-based clustering and classification via finite mixture models. We demonstrate the performance of these techniques in simulation studies and for real data.