

FINAL REPORT

SAMSI Summer Program on Challenges in Dynamic Treatment Regimes and Multistage Decision-Making June 18-29, 2007

Program Rationale

The management of chronic disorders, such as mental illness, substance dependence, cancer, and HIV infection, presents considerable challenges. In particular the heterogeneity in response, the potential for relapse, burdensome treatments, and problems with adherence demand that treatment of these disorders involve a series of clinical decisions made over time. Decisions need to be made about when to change treatment dose or type and regarding which treatment should be used next. Indeed, clinicians routinely and freely tailor treatment to the characteristics of the individual patient with a goal of maximizing favorable outcomes for that patient. To a large extent the tailoring of sequences of treatments is based on clinical judgment and instinct rather than a formal, evidence-based process.

These realities have led to great interest in the development of so-called “dynamic treatment regimes” or “adaptive treatment strategies.” A dynamic treatment regime is an explicit, operationalized series of decision rules specifying how treatment level and type should vary over time. The rule at each stage uses time-varying measurements of response, adherence, and other patient characteristics up to that point to determine the next treatment level and type to be administered, thereby tailoring treatment decisions to the patient. The objective in developing such multistage decision making strategies is to improve patient outcomes over time.

Methodology for designing dynamic treatment regimes is an emerging area that presents challenges in two areas. First, experimental designs for collecting suitable data that can be used efficiently to develop dynamic regimes are required. Second, techniques for using these and other data to deduce the decision-making rules involved in a dynamic regime must be developed. In both areas, input from researchers in a variety of disciplines and collaborations among them is critical.

Trials in which patients are randomized to different treatment options at each decision point have been proposed; however, little is known about when such trials should be conducted in lieu of the current approach of melding clinical judgment and expert opinion to formulate decision rules and using the standard two-group paradigm. An alternative approach is to conduct a series of randomized trials, as in agriculture and engineering; again, there is little guidance on how to implement this approach when the goal is to develop a dynamic regime.

Methods to make use of data in developing dynamic regimes involve complex considerations. The construction of optimized decision rules requires incorporating the effects of future decisions when evaluating present decisions, as is well-known to scientists working on improving multistage decision-making. Treatment given at any time may set a patient up for improved response to subsequent treatments or have delayed effects that either enhance or reduce effectiveness of subsequent treatments. The development of a dynamic regime hinges on how one operationalizes the relative importance of patient

outcomes over time. Researchers who work on multistage decision problems in other contexts (robotics, artificial intelligence, control theory) readily recognize these types of issues. A key challenge is to determine how to collect sufficient information to ascertain the “state” of an individual insofar as making treatment decisions goes. Computer scientists working the field of reinforcement learning and statisticians working in medical decision-making have quantified the properties that the state must possess; in addition, practical considerations associated with feasibility of collection, cost, and patient burden must be taken into account. Typically, a great deal of information is available at each decision point, and methods for best summarizing this information yet maintaining the summary’s usefulness for deciding on treatment alteration or type are required. Methods for feature extraction developed by statisticians and computer scientists are well-suited to this problem, but the focus on multistage decision-making rather than prediction requires evaluation of these methods from a different perspective.

Computational and inferential challenges arise in all of these endeavors; e.g., complexities of optimizing dynamic regimes can invalidate standard statistical inferential techniques, scientific considerations entail thinking beyond the standard loss functions familiar to statisticians, and the abundance of information at each decision point quickly leads to a “small n, large p” problem and the attendant computational issues. For some disorders, e.g., HIV infection, knowledge of the underlying within-subject biological has led to development of sophisticated mechanistic models for the processes governing disease progression and effect of treatment, which offer a scientific basis (via closed loop control methods) for designing dynamic regimes; however, this approach has not been widely explored or tested in samples of patients in this context.

Program Objectives

The critical need for development of methodology for dynamic treatment regimes, the considerable theoretical and practical challenges, and the relevance of work across a range of disciplines motivated the SAMSI Summer 2007 Program on Challenges in Dynamic Treatment Regimes and Multistage Decision-Making. The overarching goal of the Program was to stimulate and delineate the needed statistical, mathematical, and algorithmic research in this area. Specific objectives were to bring this area and the research opportunities it presents to the attention of statistical and applied mathematical scientists; to jump-start the necessary methodological development; and to nurture the necessary interdisciplinary collaboration between statistical and applied mathematical scientists, computer scientists, and health and behavioral science researchers. A further, key objective was to intrigue and engage junior researchers with relevant expertise and encourage their involvement in this area.

Program Leadership

The program was organized by Susan Murphy (University of Michigan) and the Local Scientific Coordinators Marie Davidian and Butch Tsiatis (North Carolina State University), with assistance from Daniel Scharfstein (Johns Hopkins Bloomberg School of Public Health), Joelle Pineau (McGill University).

Program Attendance

The program attracted XX participants in the first week, which featured tutorials and a workshop (see below). A total of 30 participants either continued from the first week or joined the program in the second week; the second week involved working groups and a final summary session (see below).

Program Activities

The program took place during the two weeks June 18-29, 2007 and involved four phases of activities:

Tutorials. The first three days of the program, June 18-20, were devoted to tutorial presentations by a diverse group of experts. The purpose of the tutorials was both to provide participants unfamiliar with the foundation necessary to assimilate more advanced material and to highlight the differences and similarities in perspectives, terminology, and methodological approaches across disciplines, including statistics, applied mathematics, computer science, and engineering. The tutorials were as follows:

- Monday, June 18: *Introduction to Causal Inference* (Miguel Hernan, Harvard School of Public Health, epidemiologist/statistician) and *Introduction to Dynamic Treatment Regimes* (Butch Tsiatis, North Carolina State University, statistician)
- Tuesday, June 19: *Introduction to Reinforcement Learning and Decision Theory and Markov Decision Processes* (Ron Parr, Duke University, computer scientist) and *Computational Challenges with High Dimensional Data* (Joelle Pineau, McGill University, computer scientist)
- Wednesday, June 20: *Introduction to Mechanistic Models and Control Theory* (Daniel Rivera, Arizona State University, control systems engineer) and *Introduction to Nonstandard Statistical Inference* (Susan Murphy, University of Michigan, statistician)

The tutorials inspired considerable interaction and discussion among participants during the breaks. Much discussion was devoted to beginning to understand the common themes and differences among competing approaches to the development of dynamic treatment regimes.

A poster session/reception was held on the evening of Wednesday, June 20, and featured posters presented by a diverse group of researchers, including statisticians, mathematicians, engineers, and computer scientists.

Workshop. On Thursday and Friday, June 21-22, a workshop was held, featuring talks presenting both the “big picture” of the methodological challenges as well as more advanced and targeted talks on research relevant to development of dynamic treatment regimes that

built on the foundation provided by the tutorials. After sets of related talks, a discussion session, with a defined question to frame the interactions among participants, took place.

The titles of the presentations and the speakers were as follows:

Thursday, June 21:

Introduction and Overview, Susan Murphy, University of Michigan

Sample Complexity of Policy Search with Known Dynamics, Ambuj Tewari, University of California, Berkeley

Discussion: What is the promise of work in sample complexity and upper bounds on generalization error?

Clinical data based optimal STI strategies for HIV: a reinforcement learning approach, Damien Ernst, now at University of Liege, Belgium

From Population to Individual Drug Dosing in Chronic Illness - Intelligent Control for Management of Renal Anemia, Adam Gaweda, University of Louisville

Discussion: What are the issues with model based work?

Estimation of the effect of dynamic treatment regimes under flexible dynamic visit regimes, Andrea Rotnitzky, Di Tella University and Harvard University

Asymptotic Bias Correction for Estimates of Optimal Dynamic Treatment Regimes, Erica Moodie, McGill University

Discussion: What are the issues that concern statisticians?

Adaptive stimulation design for the treatment of epilepsy, Joelle Pineau, McGill University

Bias and Variance in Value Function Estimates, Peng Sun, Duke University

Discussion: What are the key computational issues?

Friday, June 22:

Inference for Dynamic Regimes, Jamie Robins, Harvard University (this was a full morning presentation; the first part was tutorial in nature, and the second part involved current research of the speaker).

Working Groups. In the afternoon of Friday, June 22, the workshop participants planning to stay for the second week of the program divided into discussion groups centered around three key areas that would form the basis for the formal working groups that met during the

second week of the program. These areas were “Difficulties in Statistical Inference for Dynamic Treatment Regimes,” “Practical Challenges and Applications,” and “Translation Among Different Approaches.” Each participant chose a “primary” and a “secondary” group and spent time with both groups. During this afternoon session, group members “brainstormed” regarding the important challenges and questions centered around the group themes, produced lists of the most pressing problems, and outlined plans on how to address them for the following week. Originally, four working groups were planned; however, following the discussions, the consensus was to combine the planned “Bayesian Approaches” group with “Practical Challenges and Applications.”

A total of 30 participants either stayed on for the second week of the program or joined the program in the second week. On Monday, June 25, the three working groups convened for four days (June 25-28) of discussions. Participants met with both their primary and secondary groups, which held morning and afternoon sessions to discuss and prioritize challenges. During this period, the groups met as a whole for wide-ranging discussions and in subgroups to attack specific problems. Some additional, focused activities evolved:

- Participants created a wiki page for sharing information and posting references related to the program.
- In the program spirit of undertaking special measures to engage junior researchers, a group of junior participants (graduate students, postdocs, and assistant professors within a few years of receipt of their degrees), dubbed the “Under 40 Committee,” formed, and took the lead in planning further activities on Wednesday and Thursday.

The “Under 40 Committee” met in addition to the working groups to formulate activities and discussion points of interest to junior researchers. This group spearheaded the following activities:

- On Wednesday, a lunch group was convened during which several topics were discussed: (i) A list of journals relevant for different types of research articles on dynamic treatment regimes was developed; (ii) funding agencies that would support research on dynamic treatment regimes were identified, and strategies for writing grant applications in this area for these agencies were discussed; and (iii) academic departments, research centers, and other settings at which work on or interests in development of dynamic treatment regimes exists were identified as potential job opportunities for junior participants.
- On Thursday, a “brainstorming” session was convened in the afternoon, with the goal of creating a list of specific research questions to be solved. Following this, the “Under 40 Committee” members gave mini-presentations (5-10) minutes to the entire group of participants, with requests for feedback.

Overall, the working group and other discussions were very fruitful, particularly in bridging the gaps in terminology and perspectives across disciplinary boundaries.

Summary and Transitional Workshop. On Friday, June 29, the working groups met together. Each group presented their results, findings, and recommendations for the future to all participants. Discussion followed each presentation.

Program Outcomes

Central goals of the program were to expose participants to the area of dynamic treatment regimes and thereby spearhead new research on methodology for dynamic treatment regimes and the collaborations across disciplines critical to advancing the area. Follow-up with the participants demonstrates that the program had significant impact toward achieving these goals, as reflected in these representative testimonial comments:

“This workshop has altered my thinking about how I interrogate the clinician and structure things during the process of designing a clinical trial. I am also much more worried about selection bias.” (P. Thall, Biostatistics, M.D. Anderson Cancer Center).

[Referring to a collaboration developed during the program] “We almost certainly would not to cement a working relationship without the opportunity of meeting face to face that the SAMSI meeting afforded.” (E. Moodie, Epidemiology and Biostatistics, McGill University).

“Participation in the program directed my interests towards multi-armed bandit problems.” The work on multi-armed bandit, I believe, will be applicable in the dynamic treatment domain. I will soon look at specific applications. (A. Rakhlin, Computer Science, University of California, Berkeley)

“I am working with two of my students (Jinhui Ko and Sachiko Miyahara) on topics that I learned at the meeting. As a result of our participation in the workshop, we were able to enrich our Adaptive Treatment Strategies Reading Group Activities in our department (www.pitt.edu/~wahed/ATSRG/main.htm), discussing three talks from the workshop in details. Faculty and students in the Department (a group of about 10-12 researchers) became very much interested in this area.” (A. Wahed, Biostatistics, University of Pittsburgh).

“I am working on pain management decision-making in two different directions. One is treating pain via a two-stage program at the Eugene McDermott Center for Pain Management at UT Southwestern Medical Center (joint with Robert Gatchel in Psychology at UT Arlington). The other is treating pain via neurostimulation, which my colleagues are currently studying in rats (joint with Seoung Bum Kim and Jay Rosenberger in my department, J.-C. Chiao in Electrical Engr, and Yuan-Bo Peng in Psychology). I have students working on both directions. We have submitted proposals to NSF.” (V. Chen, Industrial Engineering, University of Texas at Arlington)

“For a young researcher, [the] summer program was excellent in connecting me with other researchers in this field... I learned a lot from the discussion at the program focused on young researchers, and it is great to have the wiki repository as a reference I look back at from time to time.” (M. Rosenblum, postdoc in Biostatistics, University of California, Berkeley).

Participants report that they have developed collaborations focused on dynamic treatment regimes research; some of these are as follows:

- J. Pineau (Computer Science, McGill University) and P. Sun (Fuqua School of Business, Duke University) have begun a collaboration on estimating components of the variance of an estimated policy (regime).
- E. Moodie (Epidemiology and Biostatistics, McGill University) and M. Rosenblum (postdoc in Biostatistics, University of California at Berkeley) are working on methods for estimating dynamic treatment regimes.
- E. Moodie (Epidemiology and Biostatistics, McGill University), P. Thall (Biostatistics, M.D. Anderson), and A. Wahed (Biostatistics, University of Pittsburgh) are collaborating on likelihood inference for adaptive treatment strategies (both frequentist and Bayesian).
- B. Carlin (Biostatistics, University of Minnesota) and P. Thall (Biostatistics, M. D. Anderson) have started the process of writing a book on adaptive clinical trials.
- A. Gaweda (Medicine, University of Louisville) and D. Rivera (Chemical Engineering, Arizona State University) are collaborating on a project entitled “Model Identifiably from Treatment Data.”
- M. Davidian (Statistics, North Carolina State University) and M. Kosorok (Biostatistics, University of North Carolina at Chapel Hill) have initiated discussions on a multidisciplinary project to develop dynamic treatment regimes for cystic fibrosis.

The program also inspired interdisciplinary exchange after the program, with a number of participants invited to give presentations or take sabbatic leave at the institutions of fellow participants, for example:

- E. Moodie (Epidemiology and Biostatistics, McGill University) traveled to M.D. Anderson Cancer Center (visiting P. Thall) to give a seminar on dynamic treatment regimes.
- J. Pineau (Computer Science, McGill University) traveled to the Department of Biostatistics at University of North Carolina at Chapel Hill (visiting M. Kosorok) to give a seminar on epilepsy and dynamic treatment regimes.
- A. Wahed (Biostatistics, University of Pittsburgh) traveled to McGill University (visiting E. Moodie) to give a talk on sample size estimation for dynamic treatment regimes.

- B. Carlin (Biostatistics, University of Minnesota) will spend Fall 2008 at M. D. Anderson Cancer Center to collaborate with P. Thall on a book on Bayesian adaptive clinical trials.
- S. Murphy will give a plenary talk at a conference (organized by Damien Ernst, Electrical Engineering, University of Liege, Belgium) on Machine Learning and Knowledge Discovery.

In addition to these interactions among program participants, participants in the program have or will disseminate results inspired by the program in invited talks, for example:

- M. Davidian (Statistics, North Carolina State University) and A. Tsiatis (Statistics, North Carolina State University) gave an invited presentation entitled “Novel Study Designs for Treatment Strategies that Reflect Actual Clinical Practice: Issues They Raise Regarding Implementation and Analysis” at the Research Conference for staff (including clinicians, statisticians, health care economists) at Duke Clinical Research Institute, Duke University.
- M. Davidian (Statistics, North Carolina State University) gave a talk entitled “Mathematical-Statistical Modeling to Inform the Design of HIV Treatment Strategies and Clinical Trials” at the 2008 ENAR Spring Meeting and at the 2007 FDA/Industry Statistics Workshop.
- Several program participants gave talks in a session entitled “Dynamic Treatment Regimes: Practice and Theory,” organized by P. Thall (Biostatistics, M. D. Anderson) at the 2008 ENAR Spring Meeting: S. Murphy (Statistics, University of Michigan), “Screening Experiments for Dynamic Treatment Regimes;” P. Thall, “Two-Stage Treatment Strategies Based on Sequential Failure Times;” and A. Rotnitzky (Statistics, Di Tella University, Argentina, and Biostatistics, Harvard), “Optimal Treatment and Testing Strategies with Possibly Non-ignorable Observation Processes.”
- Several program participants gave talks in a session entitled “On the Utility of Deterministic Models for Causal Effects” at the 2008 ENAR Spring Meeting: J. Lok (Biostatistics, Harvard), “Mimicking Counterfactual Outcomes to Avoid Deterministic Effects;” and J. Robins (Biostatistics, Harvard), “Causal Models for the Effects of Weight Gain on Mortality.”
- M. Davidian (Statistics, North Carolina State University) gave a talk entitled “An Introduction to Dynamic Treatment Regimes” at the School of Public Health at University of Alabama, Birmingham.
- A. Gaweda (Medicine, University of Louisville) gave a talk to the NIH Biomedical Computing Interest Group entitled “Application of Control Theory and Machine Learning to Drug Dose Determination” at NIH.

- S. Murphy (Statistics, University of Michigan) gave a talk on “Experiments and Dynamic Treatment Regimes” at the NIAID Biostatistics Research Branch.
- S. Murphy (Statistics, University of Michigan) gave a talk on “Novel Clinical Trial Designs” at the 98th Annual Meeting American Psychopathological Association, New York.
- S. Murphy (Statistics, University of Michigan) will present a talk on developing dynamic treatment regimes at the INFORMS meeting (organized by J. Ivy, Industrial Engineering, North Carolina State University).
- M. Davidian (Statistics, North Carolina State University) will be a discussant in a session on mathematical modeling and simulation to design treatment strategies at the Health Care Engineering Symposium (organized by S. Roberts, Industrial Engineering, North Carolina State University)
- M. Qian (Graduate Student in Statistics, University of Michigan) was selected to give a talk on “Model Selection for Individualized Treatment Rules with Lasso Penalty” at the Michigan Student Symposium for Interdisciplinary Statistical Sciences.
- A. Tsiatis (Statistics, North Carolina State University) has been invited to give a presentation to the Directors of the Institutes of NIH on novel methodological challenges in biomedical research; his talk will promote the need to study treatment as multi-stage decision making.

Also following from the program, several participants received grants to carry out methodological research on dynamic treatment regimes, including S. Murphy (Statistics, University of Michigan) and J. Pineau (Computer Science, McGill University), who, with two psychiatrist collaborators, has received a grant from NIMH; and D. Rivera (Chemical Engineering, Arizona State University), who received a NIH K25 award and a R21 grant to pursue research and training in dynamic treatment regimes. Several graduate student dissertations have been inspired or impacted by the program (Bibhas Chakraborty, Min Qian, Eric Laber, Lacey Gunter, Statistics, University of Michigan; Min Zhang, Statistics, North Carolina State University).

A number of publications (published, in press, submitted) and technical reports authored by participants have resulted from the program; a list is given IN THE LIST IN SECTION XX.

In summary, the program showed significant success meeting its objectives. Fruitful, synergistic collaborations among participants from complementary disciplines, which otherwise would not have materialized and which are critical to advances in the area of dynamic treatment regimes, have been established. Participants whose knowledge of the area was not extensive before the program have been inspired to undertake learning and research on dynamic treatment regimes. Junior researchers benefited significantly from

exposure to the ideas of senior researchers from diverse disciplines and from interaction with each other.

Specific Information Requested:

2. Governmental and Industrial Participation: None.

3. Publications:

Published or in press:

B.P. Hobbs and B.P. Carlin. Practical Bayesian Design and Analysis for Drug and Device Clinical Trials. *J. Biopharmaceutical Statistics*, 18, 54-80, 2008.

A.I. Oetting, J.A. Levy, R.D. Weiss, and S.A. Murphy. Statistical Methodology for a SMART Design in the Development of Adaptive Treatment Strategies. To appear in *Causality and Psychopathology: Finding the Determinants of Disorders and their Cures* (P.E. Shrout, Ed.) Arlington VA: American Psychiatric Publishing, Inc., 2008.

J.M. Robins, L. Orellana, M.A. Hernan, and A. Rotnitzky. Estimation and Extrapolation of Optimal Treatment and Testing Strategies. To appear in *Statistics in Medicine*, 2008.

S. Ross, B. Chaib-draa, and J. Pineau. Bayes-Adaptive POMDPs. To appear in *Neural Information Processing Systems (NIPS)*. 2008.

S. Ross, J. Pineau and B. Chaib-draa. Theoretical Analysis of Heuristic Search Methods for Online POMDPs". To appear in *Neural Information Processing Systems (NIPS)*. 2008.

Yang, Z., V. C. P. Chen, M. E. Chang, M. L. Sattler, and A. Wen. A Decision-Making Framework for Ozone Pollution Control. To appear in *Operations Research*, 2009. (COSMOS Technical Report 05-04).

Submitted or tech reports to be submitted:

J. Abernethy, E. Hazan, and A. Rakhlin, An Efficient Algorithm for Bandit Linear Optimization.

P. Bartlett, V. Dani, T. Hayes, S. Kakade, A. Rakhlin, and A. Tewari. High Probability Regret Bounds for Online Optimization.

A. Buu, R.A. Zucker, and S.A. Murphy. Statistical Methods for Estimating Latent Scores: A Methodological Comparison.

M.M.Fard, J. Pineau, and P. Sun. A Variance Analysis for POMDP Policy Evaluation.

E. Laber and S.A. Murphy. Small Sample Inference for Generalization Error in Classification Using the CUD Bound.

H. McGowan, R.L. Nix, S.A. Murphy, K.L. Bierman, and CPPRG, Investigating the Effects of Selection Bias in Dose-Response Analyses of Preventive Interventions.

S.A. Murphy and D. Bingham. Screening Experiments for Developing Dynamic Treatment Regimes.

S. Ross, J. Pineau, S. Paquet, and B. Chaib-draa. Online Algorithms for POMDPs: A Survey.

M. Zhang, A.A. Tsiatis, and M. Davidian. Inference on Treatment Effects from a Randomized Clinical Trial in the Presence of Premature Treatment Discontinuation.

4. Diversity: The group included a diverse set of participants, including individuals in computer science, chemical engineering, industrial engineering, statistics, biostatistics, applied mathematics, business and psychologists. We also included individuals at a variety of career stages from graduate student, postdoc, assistant professor, associate professor and full professor.

5. External Support: The Program Leaders received secretarial support during preparations for the program from NIH R21 DA019800, Methodology for Adaptive Treatment Strategies (S. A. Murphy, PI)

6. Summary of Exciting Research and Research Nuggets:

The management of chronic disorders, such as mental illness, substance dependence, cancer, and HIV infection, involves a series of clinical decisions on treatment made over time. In practice, these decisions are based on characteristics of and outcomes observed up until the current decision point for the individual patients, and are generally made based on clinical judgment rather than a formal, evidence-based process. Dynamic treatment regimes are formal sets of sequential decision rules that take as input time-varying measurements on the patient and output the next step of treatment, thus operationalizing clinical practice. Methodology for developing dynamic treatment regimes and for designing experiments to gather data to be used in such efforts is an emerging and challenging area. The SAMSI Summer Program on Challenges in Dynamic Treatment Regimes and Multistage Decision-Making brought together researchers from diverse disciplines, including statistics, applied mathematics, computer science, and engineering, to foster collaborations that can exploit the differing perspectives to accelerate advances in this area.

A *critical research nugget* is “the realization that scientists from a wide variety of disciplines can come together to collaborate and improve multi-stage decision making across a variety of fields from health-related fields (e.g. HIV, mental health, cancer) to autonomous helicopter flight, robotics, and artificial intelligence. This has had several impacts; in particular participants have formed teams:

- To develop a “mapping” relating the fields of reinforcement learning in computer science, control theory in engineering, and inferential techniques such as g-estimation in statistics, highlighting differences and similarities.
- To provide confidence measures for data-derived dynamic treatment regimes. This group involves computer scientists, business school members and statisticians.
- To consider both the combination of adaptive clinical trial design with the investigation of dynamic treatment regimes. This group involves both statisticians and biostatisticians.
- To further investigate causal issues in constructing dynamic treatment regimes from data. This group involves individuals from informatics and engineering.
- To investigate discussions on a multidisciplinary project to develop dynamic treatment regimes for cystic fibrosis. This group involves statisticians, applied mathematicians and biostatisticians.