Curr Epidemiol Rep (2016) 3:225–232
DOI 10.1007/s40471-016-0079-3

EPIDEMIOLOGIC METHODS (D WESTREICH, SECTION EDITOR)

SMART Thinking: a Review of Recent Developments in Sequential Multiple Assignment Randomized Trials

Michael P. Wallace[1] · **Erica E. M. Moodie**[1] · **David A. Stephens**[2]

Published online: 2 July 2016

Abstract With the increasing interest in personalized medicine, over the last decade, sequential multiple assignment randomized trials (SMARTs) have become a more common fixture of the clinical trial landscape. Primarily of use in the identification of dynamic treatment regimes, they have experienced a shift from the more complex designs of the past to the considerably streamlined versions seen today. In this review, we summarize their history, outline recent and ongoing examples, and discuss some of the important methodological developments for their design and implementation.

Keywords SMARTs · Dynamic treatment regimes · Adaptive treatment strategies · Longitudinal data · Causal inference

Introduction

Personalized medicine is a general term for patient-centric disease management. Its fundamental philosophy—"treat the patient, not the diagnosis"—is grounded in the idea that outcomes may be improved if treatments are tailored to individual patient characteristics, rather than the traditional approach of choosing a single treatment for all patients based on the best population-average outcome. Despite this simple principle, optimizing disease management at the patient level presents numerous challenges, both theoretical and practical.

Dynamic treatment regimes (DTRs), also known as adaptive treatment strategies, are an important component of personalized strategies for treatment. DTRs are decision rules taking patient information (such as age, disease severity, or even response to prior treatments) as input, which then output treatment recommendations unique to that individual. A DTR could be as simple as "prescribe intervention A (say a traditional consultation) if patient is over 65 years of age, otherwise prescribe intervention B (a consultation performed via a mobile phone app)," but they may be considerably more complex.

A single-treatment decision, such as the previous example, is often referred to as an individualized treatment rule. More generally, a DTR is a *sequence* of decision rules, recommending treatment decisions at fixed points over a period of time (such as a follow-up period with multiple clinic visits). What makes a DTR truly "dynamic," however, is a capacity to provide different treatment recommendations dependent on response to previous treatments. This is particularly useful as it allows more complex relationships between current treatment, past treatment(s), and patient characteristics to be taken into account.

As a simple example, we consider a hypothetical study of treatments for depression. Suppose that we were interested in comparing various sequences of treatments, where patients receive either a primary pharmaceutical therapy or a counseling-based therapy for the first 6 weeks, before switching to one of two secondary pharmaceutical therapies. Our goal is to identify the sequence of treatments that optimizes outcome, as per some pre-defined response variable (for example, depression score at 12 weeks). This process is complicated by the potential for delayed effects of treatments and the possibility of interactions (synergy or antagonism)

This article is part of the Topical Collection on *Epidemiologic Methods*

✉ Michael P. Wallace
michael.wallace@mcgill.ca

[1] Department of Epidemiology, Biostatistics and Occupational Health, McGill University, 1020 Pine Avenue West, Montreal, QC H3A 1A2, Canada

[2] Department of Mathematics and Statistics, McGill University, Montreal, QC H3A 2K6, Canada

between treatments given at different stages of the treatment sequence. An example of the latter might be that patients who initially receive counseling will respond better to the later treatments—perhaps due to increased compliance. Alternatively, the initial treatment may have a delayed effect but perhaps only in patients with certain characteristics.

One route to identification of such interactions and effects is to consider more sophisticated trial designs. Sequential multiple assignment randomized trials (SMARTs) were first introduced as "biased coin adaptive within-subject designs" by Lavori and Dawson [1, 2], with the general framework proposed by Murphy [3]. SMARTS are characterized by various treatment sequences, with Fig. 1a as a schematic of a possible SMART for our preceding example. Here, patients are randomized to one of the two initial treatments and then, at 6 weeks, randomized again to one of the two secondary treatments. If randomization was done within each of the initial treatment groups, this could guarantee that a cohort followed each of the four possible treatment regimes (A-C, A-D, B-E, and B-F) allowing a direct comparison of each. These are referred to as "embedded regimes."

While the term "sequential" in SMART would suggest that these trials require randomization to be performed in sequence, it is easy to see that SMARTs could equally well be operationalized by randomizing participants *only once* to a *sequence* of treatments that is embedded in the trial. In the above example, where randomization is independent of all covariates, the design may be further simplified by noting its equivalence to a four-arm randomization to one of the four embedded treatment sequences or even as a factorial trial with the second randomization merely delayed in time. Such an example is therefore amenable to more established (and straightforward) techniques. Trials of this nature, where the second-stage treatment does not depend on any intermediate responses, can be used to learn about DTRs even though the trials themselves do not incorporate a dynamic component. However, as the goal of such trials is typically to learn about tailoring, it is often preferable to incorporate some elements of

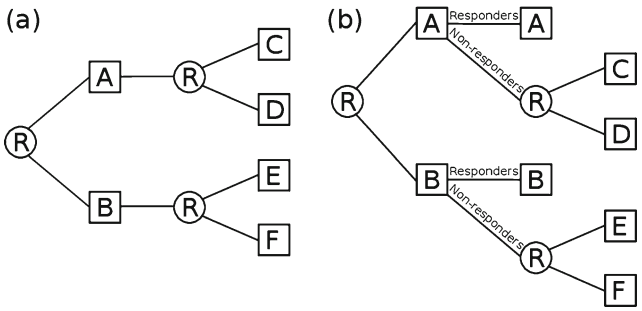

Fig. 1 Two common SMART designs, **a** all participants are re-randomized following initial treatment and **b** only participants who failed to respond to their initial treatment are re-randomized and the rest continue the same treatment. A *circled R* indicates randomization; a *boxed letter* indicates a particular treatment (with *different letters* not necessarily corresponding to different treatments)

that personalization directly into the trial design itself. This latter case is where SMARTs are of greatest utility. An example is depicted in Fig. 1b, where instead of full randomization at stage 2, only those patients who were deemed "non-responders" to the first treatment are re-randomized. We shall focus primarily on these more "dynamic" SMARTs, where a patient's treatment path cannot be known at study entry.

As we begin to think about these designs more carefully, it is critical to note the differences between a SMART and a DTR. While fundamentally distinct—a SMART is a study design, a DTR a treatment rule—it is important to additionally appreciate that a SMART is typically simple, with second-stage randomization often based on a very low dimensional intermediate measurement such as "early response." DTRs, in contrast, will typically be more complex, adapting treatment based on a larger number of intermediate measures. As we shall soon discuss, it is infeasible to conduct a SMART that involves tailoring on a large set of covariates, as the required sample size is rarely available. However, data from simpler SMARTs can nevertheless be used to estimate these more highly tailored DTRs.

In this article, we will review the properties, and consider examples, of SMARTs; summarize design considerations; and conclude with a discussion of current research and unmet needs in this vibrant and growing field.

Properties of SMARTs

Advantages

Two primary advantages of SMART trials are the ability to detect delayed effects of treatment and the possibility of discovering interactions between treatments given at different stages of the treatment sequence [4]. In addition, while SMARTs provide the ability to identify optimal treatment regimes, they also possess a number of other useful properties. Of particular use is that randomization may be based on individual patient characteristics which, in addition to improving patient outcomes, can improve patient experience during the course of the trial. For example, in a more complex study with a greater number of treatment options, a patient may be able to specify a class of treatments to which they are happy to be assigned. Similarly, randomization could be based on how well or poorly a patient responds to an initial treatment, avoiding potential ethical concerns such as switching a patient away from a treatment that appears to be effective. As randomization can still take place after the observation of such preferences or response, the benefits of randomization in terms of covariate balance are not lost.

While it may seem problematic to randomize based on characteristics such as these, as long as the randomization procedure is well documented, this may be taken into account

at the analysis stage. Suppose, for example, that a trial is conducted as in Fig. 1b, where "non-response" is defined as "desire to change treatment." The trial can then be viewed as randomization to one of four treatment options, each taking the form "initiate treatment with option A; if patient responds (does not wish to switch), remain on treatment A, otherwise change to treatment C." Thus, in Fig. 1b, there are a total of four embedded regimes that allow for tailoring, following the measurement of response to the first stage of treatment. As above, we can operationalize such a trial either by sequential randomization (first at baseline, then after observing response status) or by a single, up-front randomization to a particular regime that dictates how treatment will be given depending on response status. If response were multi-leveled, or a greater number of covariates used to define the set of allowable stage 2 treatments, the number of embedded regimes in—and correspondingly, the required sample size for—the trial would increase. For practical purposes, when randomization is based on intermediate outcomes, some low-dimensional summary measure such as patient preference or response status should be used [5].

The principle of randomization for non-responders features in two large, notable trials in mental health. The Sequential Treatment Alternatives to Relieve Depression (STAR*D) trial [6] was designed to assess the treatment options available for major depressive disorder. Recruiting over 4000 patients, the study was split into four levels (including one that was itself further split into two sub-levels), with patients proceeding to subsequent levels if they did not respond to their current treatment. While all patients were initially prescribed the same treatment (citalopram), non-responders were subsequently randomized up to three times. An interesting feature of this study was that non-responders could express a preference for their treatment options at the next randomization stage, a mechanism that helped avoid ethical concerns including randomization to an arm which maintained the use of a drug whose side effects were deemed intolerable by the participant.

The Clinical Antipsychotic Trials of Intervention Effectiveness (CATIE) for schizophrenia [7], meanwhile, recruited 1460 patients to study treatment sequences of antipsychotic drugs. Participants were initially randomized to one of five treatments and, if their symptoms worsened or the side effects became intolerable, could choose to be randomized to one of four secondary treatments within the first 18 months. Both CATIE and STAR*D have been heavily studied within the DTR literature [8–11]. A related CATIE study [12] considered treatments for Alzheimer's disease in a similar manner.

It has been purported that SMARTs may provide results that are more generalizable than traditional randomized control trials (RCTs) and that, by virtue of the greater variety of treatment options open to participants, trial dropout may be lower than in a traditional RCT. To date, there is no clear evidence of the former claim, although the STAR*D and CATIE trials were specifically designed to have few exclusion criteria so as to assess therapies in a broad clinical population. There is some evidence that the retention rate of a SMART is superior to that of traditional RCTs [13].

Challenges

An obvious limitation of SMARTs is one of power; every randomization step splits the cohort into smaller groups, and with multiple stages, we may find that the number of subjects following any given treatment regime to its conclusion is low. It is unsurprising that most SMART designs focus on just two stages of treatment; subjects are initially randomized to one of two (or more) treatments, followed up, then randomized to a secondary treatment at some later point. Despite this simplicity, however, there is still considerable flexibility in this approach. As previously noted, randomization itself may be dependent on patient characteristics, as may the point at which the second randomization takes place (for example, the second stage may be triggered by a clinical event such as disease severity passing a pre-determined threshold rather than simple calendar time). In addition, participants may not necessarily be randomized to a new treatment at the second stage if, for example, they respond well to the initial treatment.

A related challenge is that of the additional cost of running a SMART instead of a more traditional design. SMARTs are almost necessarily more complex than standard RCTs and will usually require larger sample sizes to power them sufficiently. In addition, the mere question of assessing power is itself non-trivial; as we will discuss in greater detail below, it is often advocated that one powers a study for simple two-arm comparisons of treatment "paths" through the SMART, leaving the estimation of an optimal DTR as a secondary goal, which is typically viewed as more exploratory than confirmatory.

Modern-Day SMARTs

STAR*D and CATIE, along with a number of other studies [14–18], were designed, run, and analyzed before the formal SMART design framework was developed and in particular before practical guidelines had been established. More recent SMARTs tend to be much more simplistic. It is typical for contemporary SMARTs to be limited to two stages, with at most two treatment options per participant per stage. For example, the most common SMART design is of the form illustrated in Fig. 1b.

An example of such a design was featured in a recent trial of neurobehavioral treatments for patients with metastatic malignant melanoma undergoing high-dose interferon-alpha therapy [19]. This study proposed randomizing 70 patients to one of two initial treatments (e.g., citalopram or methylphenidate) for 6 weeks, at which point symptoms would be evaluated. Response to treatment was defined as a Hamilton

Depression Scale score below 12, with those patients who respond continuing their initial treatment as they begin interferon treatment. Non-responders, meanwhile, were to be randomized to either augment their current treatment with the other initial treatment or switch to the initial treatment not yet tried.

This design gives rise to four different DTRs, which can be compared in terms of the primary outcome (in this case, maximizing adherence to interferon therapy). In addition, this design allows the investigation of secondary aims, including comparing the effects of the two initial treatments on questionnaire measures at 6 weeks and switching versus augmentation among non-responders in terms of the primary outcome.

The BestFIT study [20], which is currently recruiting participants, also follows the same general SMART design to compare weight loss treatments. Participants are initially randomized to either 3 or 7 weeks of a state-of-the-art therapy, with responders continuing with that treatment. Non-responders, meanwhile, are randomized to either augmentation of current treatment with portion controlled meals or switching to an "acceptance-based" therapy.

The authors identify the following two "critical questions" the study addresses: a primary aim to identify which secondary treatment is preferred and a secondary objective of when to identify sub-optimal response to the initial therapy. Despite the title of the paper mentioning the development of individualized weight loss treatments, this is listed as an "exploratory third aim," an issue that is again pertinent to sample size calculations.

This general SMART structure of maintaining "successful" treatments and randomizing all non-responders has been featured in numerous other studies, including another study of weight loss interventions [21], along with studies of attention hyperactivity disorder [22, 23] and multiple myeloma [24]. In contrast—but still adhering to a principle of low dimensionality—some studies re-randomize only those patients who fail to respond to one specific initial treatment (rather than all non-responders), such as the Adaptive Characterizing Cognition in Non-Verbal Individuals with Autism (CCNIA) Developmental and Augmented Intervention trial [25] for school-age, non-verbal children with autism spectrum disorder. In this study, patients were initially randomized to one of two treatments, one of which included a communication device. Researchers felt that it was inappropriate to re-randomize participants who had initially received the device to a different treatment (regardless of response). Consequently, those patients would continue with the same treatment, with those who had responded continuing at the same level, while non-responders would have their treatment intensified. Non-responders to the other initial treatment, meanwhile, would be randomized to one of two treatment options.

A final and more general example of this style of SMART design sees that *all* patients re-randomized to a second-stage treatment regardless of response (although response can determine which treatment options are available). Examples include the ExTENd study of alcohol dependence [26] and the study Adaptive Reinforcement-Based Treatment for Pregnant Drug Abusers [27]. This last design represents the more complex end of most modern SMARTs, wherein the two-stage case with two treatment choices at each stage of the SMART contains eight distinct embedded DTRs.

SMART Design Issues

When compared with earlier studies such as STAR*D and CATIE, the above-noted trials illustrate more recent trends in SMART design towards smaller, more manageable studies. This is largely due to the fact that a significant challenge in SMART design lies in the required power and sample size calculations. As noted above, SMARTs can easily incorporate a large variety of different treatment regimes, with the resulting sample size for any one of these embedded regimes necessarily very small. Furthermore, the complexity of the study question and design can lead to intractable sample size calculations since considerable knowledge is required (for example, knowledge of the stage 1 response rate may be necessary to power stage 2), and some participants may contribute information on the effect of more than one embedded regime in the trial. To this end, Murphy [3] recommends that the primary research question should focus on simple DTRs, with few intermediate tailoring variables on which randomization may depend and few treatment options at each stage. More generally, a SMART should be thought of as one study in a series, perhaps with the ultimate goal of a confirmatory randomized trial to compare a proposed DTR with standard care [3, 28], conducted as a traditional two-arm trial.

Design is further complicated by the different possibilities for the primary objective of a SMART. It is rare that a researcher is interested in DTRs that are as simple as those embedded in the trial, and yet, power calculations for a comparison of complex DTRs that tailor on possibly many more variables than those used for second-stage randomization do not exist. Moreover, there is typically insufficient information available in the literature or in a pilot study to adequately inform simulations. Thus, power is often based in primary objectives that do not focus on the dynamic nature of the treatments the analyst wishes to estimate. Instead, for example, common primary goals of SMARTs are identification of the best first-stage treatment or the best second-stage treatment for non-responders to a specific first-stage treatment (an approach taken by the trial for interferon therapy and BestFIT). In principle, such simple (non-dynamic) objectives can often be addressed by standard RCTs, but such standard trials do not permit secondary analyses that

incorporate tailoring. Thus, even with simplistic primary objectives that do not require sequential randomization, SMARTs are a necessity if our focus is specifically on DTRs, such as comparison of two specific regimes, or whether a specific DTR can be improved by further tailoring of treatment.

As in all trials, prior knowledge should be used to identify the treatments (set of DTRs) that should be considered. This could be based on pilot studies, scientific expertise, or through simulation-based approaches. A recent illustration of the latter is provided by Rich et al. [29] who, in the context of warfarin treatment, simulate SMARTs to identify which dosing strategies should be investigated. In contrast, other authors have conducted pilot studies to assess proposed treatment regimes before committing to a full-scale SMART. For example, Gunlicks-Stoessel et al. [30] conducted a 16-week pilot SMART with 32 participants to study DTRs for adolescent depression. The study aimed to address whether non-response to initial treatment should be determined at 4 or 8 weeks and whether non-responders should have their treatment intensified or augmented with a secondary treatment. Chronis-Tuscano et al. [31] describe a similar pilot SMART of ADHD treatments featuring 26 participants.

Having chosen a primary objective, power must be explored. For a simpler, non-dynamic primary aim, tools exist to compute power based on comparisons of main effects only [22] (such as comparing the main effect of first-stage treatment while controlling for second-stage treatment) or DTR comparison methods based on test statistics using marginal mean model variances (cf., e.g., http://sites.google.com/a/umich.edu/kidwell/home/tools-for-design-and-analysis). Alternatively, if the primary aim is to compare two embedded DTRs, the calculations of Dawson and Lavori [32] are more appropriate.

A further issue is one of optimal DTR estimation, whereby tailoring variables are used (post hoc) to identify personalized regimes not necessarily embedded in a SMART design. While numerous methods have been proposed for estimation of optimal DTRs [33, 34•], little work has considered them from a power perspective. Dawson and Lavori [32] discuss how their sample size methods impact optimal DTR estimation, but this remains a little-studied area.

The powering of SMART trials is not limited to the simplest, "everyday" settings already discussed. Another important, more specialized, class of SMART design focuses on clustered data, with interventions delivered at a group level. A recent example is the Adaptive Implementation of Effective Programs Trial (ADEPT) [35], which aimed to improve evidence-based practice use and patient outcomes in participants with mood disorders. In this setup, the treatment options—the use of external and/or internal facilitators—were administered to all patients within one of 80 outpatient clinics, with any subsequent re-randomization also administered in clusters. The design of this clustered SMART featured sample size calculations but only for the purpose of comparing initial treatments—i.e., another non-dynamic primary aim. Chakraborty et al. [34•, Sect. 5] derived sample size calculations for comparison of two or more embedded regimes within a clustered SMART design.

A recurring theme in the literature (and indeed, one we have highlighted here) is the problem of small sample sizes. Tamura et al. [36], for example, have recently discussed SMART designs in the context of rare diseases, describing a "small N" SMART (or snSMART). Cheung et al. [37] consider studies with staggered recruitment and introduce the concept of a SMART with Adaptive Randomization (SMART-AR), whereby intermediate analysis is used to inform randomization probabilities as patients enter and proceed through a study. This approach proposes that at first, when data are few, only "simple" designs should be considered, but as more participants are recruited, the overall study can become more complex. If patients are recruited quickly, this will reduce to a standard (non-adaptive) SMART, and so, SMART-AR is most advantageous when enrollment is spread over time.

Conclusion

As interest in personalized medicine grows, so does the prevalence of SMART or SMART-like studies. Their appeal is grounded in a number of attractive features, not least the ability to embed numerous dynamic treatment regimes for direct comparison. Unfortunately, as early studies such as CATIE and STAR*D demonstrate, it is easy for SMARTs to become so complex that large samples (and considerable expense) are required if meaningful results regarding treatment tailoring are to be found.

Consequently, this has led more contemporary studies to be much more limited in scope. Almost all of the SMARTs discussed in this review featured only two stages of randomization, with at most eight embedded DTRs. This necessarily restricts the extent to which truly individualized treatments can be pursued, while simultaneously requiring more complex (and expensive) trials than those following more established designs. Indeed, nearly all of the examples of modern-day SMARTs that we have discussed have a stated primary goal of main effect comparisons of first- and second-stage treatments, for which traditional RCTs can usually offer insight.

A major focus of recent methodological work has been sample size calculations. For this, the primary aim of any study becomes paramount, with older methods focusing on powering for the aforementioned main effect comparisons. Powering for more direct DTR comparisons has received some recent attention, but is still limited in scope, and little

work has been done towards what could be considered the ultimate goal of optimal DTR estimation.

A closely related area is the development and estimation of more tailored DTRs (see [33] for a comprehensive review). Numerous methods have been proposed including, among many others, g-estimation, Q-learning, and dynamic marginal structural models [38–40]. All may be used on data generated from a SMART to estimate treatment regimes that tailor on numerous variables, not limited to those used to determine treatment allocation in the SMART. These methods may also be applied to observational data, potentially giving access to much greater patient numbers at reduced cost and complexity relative to a SMART. Many such methods rely on knowing or modeling how treatment is allocated, and all rely on the assumption of no unmeasured confounding. Hence, a SMART would always be preferable, even if it is not always feasible. For an introduction to DTRs and their estimation, the interested reader is referred to Wallace and Moodie [41].

While more complex methods are required to estimate more deeply tailored regimes, simple regimes that tailor only in the variable used to determine second-stage randomizations can be conducted quite simply using intention-to-treat analysis, where the *intended* treatment is determined by the outcome of the randomization. As SMARTs have seen second-stage treatment triggered by patient dissatisfaction with current treatment, and the context of most SMARTs is such that alternative treatments are not easily accessible, non-compliance with assigned treatment is rare. Thus, "as-treated" analyses have not been popular in the SMART literature. In principle, however, all methods for DTR estimation could be applied to data with non-compliance without additional modifications.

Table 1 Key points

- A SMART is a particular design of randomized trial that can be used to examine sequences of treatments. SMARTs may incorporate tailoring of treatment allocations based on participant response but need not.
- A dynamic treatment regime (DTR), in contrast, is a "rule book" for treatment allocation that may be determined through experience (trial and error by individual physicians), through analysis of observational data, or by analysis of data arising from SMARTs.
- A SMART design can be operationalized either by randomization performed in sequence or by a single, up-front randomization to a sequence of treatments or a dynamic treatment regime.
- A feasible SMART will typically incorporate only one or two intermediate covariates in the tailoring of treatment (or treatment randomization).
- A good DTR may require tailoring on several covariates. Such regimes can be estimated from SMART study data but require specialized methods of analysis since SMARTs allocate randomized treatment based on few covariates.
- Powering SMARTs is a complex (and ongoing) challenge. Current work focuses on powering not for DTR comparison but simpler objectives such as main effect comparisons.

Although SMARTS do offer a valuable path towards DTR analysis, there is something of a disconnect between the high level of tailoring about which we wish to learn and what SMARTs can realistically provide (see Table 1). At present, despite their inherent complexity, most SMART designs do not involve the level of tailoring which we are ultimately interested in using in clinical practice. In many situations, it may well be preferable to focus on observational data to inform a traditional two-arm RCT comparing a proposed DTR and standard care. Nevertheless, SMART design continues to be of active interest to the research community, and as further work is done—both in theory and in practice—we hope to see more of its potential realized.

Acknowledgments This work is funded by Natural Sciences and Engineering Research Council (NSERC) of Canada Discovery Grants. Dr. Moodie is sponsored by a Chercheur-Boursier career award from the Fonds de Recherche du Québec-Santé (FRSQ), and Dr. Wallace is partially supported by a training scholarship from the Canadian Network for Advanced Interdisciplinary Methods for comparative effectiveness research, Drug Safety and Effectiveness Network (CAN-AIM DSEN).

Compliance with Ethical Standards

Conflict of Interest The authors declare that they have no conflict of interest.

Human and Animal Rights and Informed Consent This article does not contain any studies with human or animal subjects performed by any of the authors.

References

Papers of particular interest, published recently, have been highlighted as:
- Of importance

1. Lavori PW, Dawson R. A design for testing clinical strategies: biased adaptive within-subject randomization. Journal of the Royal Statistical Society, Series A. 2000;163:29–38.
2. Lavori P, Dawson R. Dynamic treatment regimes: practical design considerations. Clinical Trials. 2004;1:9–20.
3. Murphy SA. An experimental design for the development of adaptive treatment strategies. Statistics in Medicine. 2005;24(10):1455–81.
4. Thall PF, Wooten LH, Logothetis CJ, Millikan RE, Tannir NM. Bayesian and frequentist two-stage treatment strategies based on sequential failure times subject to interval censoring. Statistics in Medicine. 2007;26:4687–702.
5. Almirall D, Lizotte D, Murphy S. SMART design issues and the consideration of opposing outcomes, a discussion of evaluation of viable dynamic treatment regimes in a sequentially randomized trial of advanced prostate cancer by Wang, Rotnitzky, Lin, Millikan, and Thall. Journal of the American Statistical Association. 2012;107(498):509–12.
6. Rush AJ, Fava M, Wisniewski SR, Lavori PW, Trivedi MH, Sackeim HA, et al. Sequenced Treatment Alternatives to Relieve

Depression (STAR*D): rationale and design. Controlled Clinical Trials. 2004;25:119–42.
7. Stroup TS, McEvoy JP, Swartz MS, Byerly MJ, Glick ID, Canive JM, et al. The National Institute of Mental Health Clinical Antipsychotic Trials of Intervention Effectiveness (CATIE) project: schizophrenia trial design and protocol development. Schizophrenia Bulletin. 2003;29:15–31.
8. Pineau J, Bellemare MG, Rush JA, Ghizaru A, Murphy SA. Constructing evidence- based treatment strategies using methods from computer science. Drug and Alcohol Dependence. 2007;88(Supplement 2(0)):S52–60. **Customizing treatment to the patient: adaptive treatment strategies. Available from:** http://www.sciencedirect.com/science/article/pii/S0376871607000270.
9. Shortreed SM, Laber E, Lizotte DJ, Stroup TS, Pineau J, Murphy SA. Informing sequential clinical decision-making through reinforcement learning: an empirical study. Machine Learning. 2011;84(1-2):109–36. Available from: doi:10.1007/s10994-010-5229-0.
10. Shortreed SM, Moodie EEM. Estimating the optimal dynamic antipsychotic treatment regime: evidence from the sequential multiple-assignment randomized Clinical Antipsychotic Trials of Intervention and Effectiveness schizophrenia study. Journal of the Royal Statistical Society, Series C. 2012;61(4):577–59.
11. Zhao Y, Zeng D. Recent development on statistical methods for personalized medicine discovery. Frontiers of Medicine. 2013;7(1):102–10. **Available from: doi:**10.1007/s11684-013-0245-7.
12. Schneider LS, Tariot PN, Lyketsos CG, Dagerman KS, Davis KL, Davis S. National Institute of Mental Health Clinical Antipsychotic Trials of Intervention Effectiveness (CATIE): Alzheimer disease trial methodology. American Journal of Geriatric Psychiatry. 2001;9:346–60.
13. Moodie EEM, Karran J, Shortreed SM. A case study of the generalizability and retention rate of a SMART study. 2015; Submitted.
14. Habermann TM, Weller EA, Morrison VA, Gascoyne RD, Cassileth PA, Cohn JB, et al. Rituximab-CHOP verses CHOP alone or with maintenance rituximab in older patients with diffuse large B-cell lymphoma. Journal of Clinical Oncology. 2006;24:3121–7.
15. Matthay KK, Villablanca JG, Seeger RC, Stram DO, Harris RE, Ramsay NK, et al. Treatment of high-risk neuroblastoma with intensive chemotherapy, radiotherapy, autologous bone marrow transplantation, and 13-cis-retinoic acid. The New England Journal of Medicine. 1999;341:1165–73.
16. Stone RM, Berg DT, George SL, Dodge RK, Paciucci PA, Schulman PP, et al. Granulocyte-macrophage colony-stimulating factor after initial chemotherapy for elderly patients with primary acute myelogenous leukemia. New England Journal of Medicine. 1995;332:1671–7.
17. Stone RM, Berg DT, George SL, Dodge RK, Paciucci PA, Schulman PP, et al. Postremission therapy in older patients with de novo acute myeloid leukemia: a randomized trial comparing mitoxantrone and intermediate-dose cytarabine with standard-dose cytarabine. Blood. 2001;98:548–53.
18. Tummarello D, Mari D, Graziano F, Isidori P, Cetto G, Pasini F, et al. A randomized, controlled phase III study of cyclophosphamide, doxorubicin, and vincristine with etoposide (CAV-E) or teniposide (CAV-T), followed by recombinant interferon-alpha maintenance therapy or observation, in small cell lung carcinoma patients with complete responses. Cancer. 1997;80:2222–9.
19. Auyeung SF, Long Q, Royster EB, Murthy S, McNutt MD, Lawson D, et al. Sequential multiple-assignment randomized trial design of neurobehavioral treatment for patients with metastatic malignant melanoma undergoing high-dose interferon-alpha therapy. Clinical Trials. 2009;6:480–90.
20. Sherwood NA, Butryn ML, Forman EM, Almirall D, Seburg EM, Crain AL, et al. The BestFIT trial: a SMART approach to developing individualized weight loss treatments. Contemporary Clinical Trials. 2016;47:209–16.
21. Naar-King S, Ellis DA, Idalski AC, Templin T, Jacques-Tiura AJ, Brogan KH, et al. Journal of Clinical Child and Adolescent Psychology. 2015.
22. Nahum-Shani I, Qian M, Almirall D, Pelham WE, Gnagy B, Fabiano GA, et al. Experimental design and primary data analysis methods for comparing adaptive interventions. Psychological Methods. 2012;17:457–77.
23. Nahum-Shani I, Qian M, Almirall D, Pelham WE, Gnagy B, Fabiano GA, et al. Q-learning: a data analysis method for constructing adaptive interventions. Psychological methods. 2012;17(4):478–94. Available from: http://search.proquest.com.proxylib.umich.edu/docview/1314310841?accountid=14667.
24. Mateos MV, Orio A, Martínez-López J, Gutiérrez N, Teruel AI, de Paz R, et al. Bortezomib, melphalan, and prednisone versus bortezomib, thalidomide, and prednisone as induction therapy followed by maintenance treatment with bortezomib and thalidomide versus bortezomib and prednisone in elderly patients with untreated multiple myeloma: a randomised trial. The Lancet Oncology. 2010;11:934–41.
25. Kasari C. Developmental and augmented intervention for facilitating expressive language (CCNIA). Bethesda: National Institutes of Health; 2009. Available from: http://clinicaltrials.gov/ct2/show/NCT01013545.
26. Oslin D. Managing alcoholism in people who do not respond to naltrexone (EXTEND). Bethesda: National Institutes of Health; 2005. Available from: http://clinicaltrials.gov/ct2/show/NCT00115037.
27. Jones H. H. Jones. Reinforcement-based treatment for pregnant drug abusers (HOME II). National Institutes of Health, Bethesda, MD. 2010; Available from: http://clinicaltrials.gov/ct2/show/NCT01177982
28. Murphy S, Lynch K, Oslin D, et al. Developing adaptive treatment strategies in substance abuse research. Drug and Alcohol Dependence. 2007;88:s24–30.
29. Rich B, Moodie EEM, Stephens DA, Platt RW. Simulating sequential multiple assignment randomized trials to generate optimal personalized warfarin dosing strategies. Clinical Trials. 2014;11:435–44.
30. Gunlicks-Stoessel M, Mufson L, Westervelt A, Almirall D, Murphy S. A pilot SMART for developing an adaptive treatment strategy for adolescent depression. Journal of Clinical Child and Adolescent Psychology. 2015;18:1–15.
31. Chronis-Tuscano A, Wang CH, Strickland J, Almirall D, Stein M. Personalized treatment of mothers with ADHD and their young at-risk children: a SMART pilot. Journal of Clinical Child and Adolescent Psychology. 2016.
32. Dawson R, Lavori PW. Sample size calculations for evaluating treatment policies in multi-stage designs. Clinical Trials. 2010;7:643–52.
33. Chakraborty B, Moodie EEM. Statistical methods for dynamic treatment regimes. New York: Springer-Verlag; 2013.
34.• Kosorok MR, Moodie EEM. Adaptive treatment strategies in practice. 2016. **This book provides an overview of the personalized medicine literature, from planning trials to data analysis. It is written with a wide variety of readers in mind, ranging from the discussion of broad principles to specialist implementations**.
35. Kilbourne AM, Almirall D, Eisenberg D, Waxmonsky J, Goodrich DE, Fortney JC, et al. Protocol: Adaptive Implementation of Effective Programs Trial (ADEPT): cluster randomized SMART trial comparing a standard versus enhanced implementation strategy to improve outcomes of a mood disorders program. Implementation Science. 2014;9.
36. Tamura RN, Krischer JP, Pagnoux C, Micheletti R, Grayson PC, Chen Y, et al. A small n sequential multiple assignment randomized trial design for use in rare disease research. Contemporary Clinical Trials. 2016;46:48–51.

37. Cheung YK, Chakraborty B, Davidson KW. Sequential multiple assignment randomized trial (SMART) with adaptive randomization for quality improvement in depression treatment program. Biometrics. 2016;71(2):450–9.
38. Robins JM. Optimal structural nested models for optimal sequential decisions. In: Lin D, Heagerty P, editors. Proceedings of the second Seattle symposium on biostatistics. New York: Springer-Verlag; 2004. p. 189–326.
39. Watkins CJCH. Learning from delayed rewards, PhD dissertation. Cambridge, England: Cambridge University Press; 1989.
40. Sutton RS, Andrew GB. Reinforcement learning: an introduction. Cambridge: MIT press; 1998.
41. Wallace MP, Moodie EEM. Personalizing medicine: a review of adaptive treatment strategies. Pharmacoepidemiology and Drug Safety. 2014;23(6):580–5.