Published in final edited form as:
Annu Rev Med. 2008 ; 59: 443–453. doi:10.1146/annurev.med.59.062606.122232.

Adaptive Treatment Strategies in Chronic Disease

Philip W. Lavori and
Stanford University, lavori@stanford.edu

Ree Dawson
Frontier Science Technology and Research Foundation, dawson@fstrf.dfci.harvard.edu

Abstract

An adaptive treatment strategy (ATS) is a rule for adapting a treatment plan to a patient's history of previous treatments and the response to those treatments. The ongoing management of chronic disease defines an ATS, which may be implicit and hidden, or explicit and well-specified. The ATS is characterized by the use of intermediate, early markers of response to dynamically alter treatment decisions, in order to achieve a favorable ultimate outcome. We illustrate the concept of ATS with some examples and describe the way that the effect of initial treatment decisions depends on the performance of subsequent decisions at later stages. We show how to compare two or more ATS, or to determine an optimal ATS, using a sequential multiple assignment randomized (SMAR) trial. We observe that clinical trials designers might find the ATS concept useful in improving the efficiency and ecological relevance of clinical trials.

Keywords

treatment policy; dynamic treatment

DEFINITION OF 'ADAPTIVE TREATMENT STRATEGIES'

An adaptive treatment strategy (ATS) is a rule for adapting a treatment plan to the changing state of an individual patient, taking into account both the history of previous treatments and the response to those treatments (1,2,3,4). For example, the clinical management of HIV infection may begin with a particular combination of anti-viral medications, and then as the patient's viral load and CD4 count change over time, the combination may be changed, or other treatments may be instituted. A clinician treating a patient with any chronic disorder for an extended period of time may be following an ATS, which may be more or less explicit. By contrast, a single treatment choice (such as whether to start a selective serotonin reuptake inhibitor (SSRI) in a first episode of major depressive disorder) may be viewed as an isolated decision when taken out of the context of the future treatment plans that depend on how the patient's symptoms respond or on the occurrence of side effects.

WHY CONSIDER A STRATEGY AS A WHOLE?

The reason for considering an ATS as a whole instead of focusing on the individual treatment decisions that make it up, follows from the observation that the overall consequences of a current treatment choice may depend on the performance of future treatments that are applied after the results of the current treatment are manifested. The optimal choice of current treatment may need to take those future treatments into account. For example, a highly toxic cancer

Corresponding Author: Philip W. Lavori, Department of Health Research and Policy, Stanford University School of Medicine, Stanford, CA 94305-5405, Tel: 650-725-6109, Fax: 650-725-6951, lavori@stanford.edu.

chemotherapy A may produce higher rates of response than a less toxic version B, but many of the failures that occur on B might yet be reversed by A after B. If so, then the ATS 'first B, and then if failure, A' may dominate the ATS 'A first, then if failure, B', perhaps after taking account of differential treatment mortality and morbidity. Such tradeoffs pervade clinical decision-making, including the choice of frontline antibiotic, initial anti-depressant therapy, sequencing of cancer chemotherapy, and many others. One very general trend in medical care is the gradual substitution of perhaps somewhat less effective but gentler treatments early in the course of clinical management. This substitution is based on the knowledge or hope that with better downstream follow-on or 'rescue' treatments one can afford to give up some immediate benefit in order to avoid side effects, or other costs, without sacrificing ultimate benefits.

An important feature of any ATS is that the individual decisions at each stage must be based on knowledge available at that stage. It would be ideal to know the results of a blood glucose management decision in terms of the clinical outcomes of blindness, cardiovascular events, and quality adjusted life-years, but it is necessary to depend on a measurable clinical indicator, such as the area under the glycosylated hemoglobin curve over recent time, or some other summary of the current state of glycemic control. Thus, the ultimate outcome 'figure-of-merit' might not be purely a function of the states used to drive treatments. For example, at one point in an ATS for treating lymphoma, the decision for hematopoetic stem cell transplant might depend on risk factors for relapse following complete remission, but the outcome is overall or progression-free survival.

Despite the ubiquity of two- or three-stage ATS of the kind described above, the evidence base for comparing different versions among themselves seldom comes from direct studies of entire ATS. We show how to compare two or more ATS, or more generally, to determine an optimal ATS, given the ability to do all the necessary experiments. Then we consider methods for doing the right experiments (clinical trials) to compare real ATS and make statistical inferences about their effects. We conclude with an observation that suggests that clinical trials designers might take an adaptive view of study treatment conditions, in order to improve the efficiency and ecological relevance of trials.

COMPARING AND OPTIMIZING ADAPTIVE TREATMENT STRATEGIES

To compare ATS we need a criterion for individual outcome (Y). For example, in managing major depressive illness, Y might indicate the achievement of a full remission, or might be a summary of quality of life over an extended period. In cancer chemotherapy, five-year survival might be the right measure, or if side effects burden should be considered, a measure of survival time without symptoms or toxicity. In the management of chronic pain, Y might represent total days without pain or excessive sedation. We also need a description of the patient states that will in sequence define the next treatment called for by a particular ATS, when considered along with the record of treatments assigned in the past. The standard notation (4) for an ATS is presented in Table 1.

We illustrate the optimization process for the two-stage adaptive scenario described in Table 2a, since that contains all the essential features and avoids excessive notation. For 60+ year old patients with untreated diffuse large B-cell lymphoma, the question arises (5) as to the use of rituximab (R) in combination with the standard induction regimen (CHOP: cyclophosphamide, doxorubicin, vincristine, and predisone), and also in its subsequent use in the maintenance of patients who have a response. For purposes of this discussion, we think of the first state as a constant, and for patients who do not respond to the induction treatment, we lump all subsequent treatments into a catch-all 'standard of care' (SOC). In this simplified scenario, there are four possible ATS, depending on whether CHOP or R-CHOP is used for

induction, and whether maintenance rituximab (MR) or observation (OBS) is used in patients who respond to the induction. More elaborate and complete ATS might include specifications of what to do with younger patients (defining the first state based on age, for example), and those who do not respond (by choosing a specific set of rules for the SOC). To fix ideas, Y might be taken as vital status at 2 years.

For any particular ATS, we can imagine the average Y that would result if all patients were managed with that ATS, and we can define the optimal ATS as the one with the best average Y, among all ATS (6,7,8). The construction of an optimal 2-stage ATS by 'backward induction' is a useful way to see how the future performance of the strategy defines the best choice at each stage. The key point is that optimization starts at the end and works backward. The first step is to stratify the patient population by the pre-treatment state, the possible first treatment choices, and the possible results of that choice. For each of those strata, we find the best choice of second treatment. In the simple example, there are only 2 strata to consider, namely R-CHOP responders and CHOP responders (see Table 2b). In R-CHOP responders, whichever is larger, of x_1 or y_1, defines the best maintenance treatment in that stratum. We pick the best treatment in this stratum (call it B); the corresponding best survival, m_1, is known as the 'optimal benefit to go' in the literature (7). In CHOP responders, whichever is larger, of u_1 or v_1, defines the best maintenance strategy (B′) in that stratum (with survival m_1), and so forth. Note that the best choice in this stratum of CHOP responders does not have to be the same as the best choice in the R-CHOP responders. Furthermore, even though all strategies call for the same treatments in non-responders to R-CHOP and CHOP, the respective rates of survival, z_1, w_1, do not have to be the same, since the non-responders to R-CHOP are not necessarily the same patients as the non-responders to CHOP.

To back up one step, we have to compare the results of the first treatment option *in terms of the optimal choices for the second step.* Thus, we plug in the previous results to find that the contrast at the first stage depends on the various probabilities of survival. The choice of initial treatment comes down to a comparison of r and r', and completes the definition of the optimal ATS, since the subsequent decision has been identified already. The reader will be able to imagine the results of more stages and more scope of variation in treatment choices.

An alternative way of defining a 'good' ATS would be to start at the beginning and find the initial treatment for each initial state that produced the best distribution of second stages, and then for each pair of states resulting from that choice, find the best follow-on treatment in terms of Y. In our example, this would amount to choosing the induction strategy that produces the highest rate of response, and then choosing the maintenance strategy in responders that produces the best 2-year survival. It should be obvious from the algebra that this need not produce the best overall ATS. The first stage of this method looks ahead only to the next state, so we can call this the best 'myopic' strategy. The myopic method for finding good strategies may get trapped away from the optimal ATS at an early stage, since the myopic method cannot factor in the way that the subsequent treatments succeed or fail given the states that result from applying the first treatment. However, it seems clear that the kinds of studies that would be required to put an evidence base under the best myopic ATS are much faster and simpler than the ones required for the optimal ATS.

METHODS FOR COMPARING ATS

The practical utility of ATS will depend on the ability to estimate the quantities in Table 2b, and thus to compare and optimize ATS. There are broadly speaking three modes of inquiry into the relative merits of different ATS (4,6,7). They are (a) observational studies of the uncontrolled variation in sequential treatment choices that occur in practice ('data mining') (b) one-time randomized trials of whole ATS, in which patients are experimentally assigned to

entire ATS at the outset of the study, and the average Ys compared across those assignments (by ITT), or (c) sequential multiple assignment randomized (SMAR) trials, in which at each decision point the patient is randomized among the appropriate options for treatments, within strata defined by the history of prior treatments and responses.

DATA MINING

One of the benefits promised by the 'electronic medical record' (EMR) is that the patient's history of treatments and responses will be readily available to the clinician, thus making it possible to implement an ATS. The record left by such extended histories of sequential decision-making reveals the extent of variation in ATS apparently followed by different clinicians. In principle it seems possible to match each observed pattern of treatments to an explicit ATS, and use the outcome data (Y) to compare different ATS, and even to find the optimal ATS (among all those that are apparently used in the observed practice of clinicians). For example, it would seem to be possible to find the optimal medication sequencing strategy for managing hypertension, from the millions of records of such patients maintained by large HMOs or the Department of Veterans Affairs health system. Setting aside the shortcomings of current EMRs (especially their varying quality and lack of pertinent outcome data, Y) there is a more fundamental problem. The data mining or 'epidemiologic' approach has the same weakness as all non-experimental methods of measuring treatment effects, namely, the reliance on an untestable assumption that there are no unmeasured confounders of treatment assignment (NUC) (9,10,11). The central weakness of observational studies is exacerbated by the multi-stage sequential nature of the ATS, which requires a sequential version of NUC at each stage. Furthermore, the analytic methods needed to produce unbiased estimates from such sequential data (even assuming NUC) are different from those used by data miners in the context of single-stage treatments, mainly traditional classification and regression adjustment methods. Finally, a combinatorial explosion results from the many variations in apparent strategy that occur, most of which have no real claim to optimality.

Nevertheless, despite the weaknesses identified above, there is clearly a role for such data mining efforts. They provide crucial information on the strategies in current use, and the range of outcomes (Y) that occur. In designing a true experiment, one cannot ignore current practice, since an experimental intervention that takes clinicians and patients far off the usual road may fail because of non-adherence. The analysis of naturalistic strategies can also reveal the extent of naturally occurring non-adherence to the apparent ATS, and suggest additions to the ATS that accommodate that reality. In the management of hypertension, one of the 'states' that occur frequently for some patients is the state in which they do not take the medications that have been prescribed, or follow the diet recommendations that have been made, and may not even show up for regular disease-management appointments. Such a patient state may also occur frequently in the management of depression, type II diabetes, substance abuse disorders, and many others. Success or failure at dealing with the 'non-adherence' state in all its forms may be one of the most important determinants of overall success or failure of a treatment policy. When this is so, it is necessary to build efforts at managing the non-adherent patient into the ATS; otherwise the strategy is incomplete. That is, a complete ATS needs to specify what will be done if the patient is in any one of the possibly several different 'non-adherent' states. There is valuable information in the naturalistic data on rates of non-adherence and perhaps even clues for how to augment ATS to provide options for dealing with it.

Conversely, successful data mining for comparing treatments in the chronic disease context requires attention to the adaptive nature of decision making. Standard methods for the statistical adjustment for observed confounders of fixed treatments fail if applied to the time-varying treatment situation.

ONE-TIME RANDOMIZATION AMONG A LIST OF WHOLE ATS

The one-time (sometimes called 'baseline') randomized ATS trial looks much like a standard trial, and the randomization guarantees NUC, so inference is both familiar and well-grounded in the randomization. In a sense there is nothing new to studying ATS this way. For example, the VANQWISH trial (12) compared an 'invasive' management strategy (immediate angiography followed by revascularization if indicated) to 'conservative' medical management (followed by revascularization, should the patient exhibit ischemic changes) in patients with non-Q-wave myocardial infarction (NQMI), which we can again take as the (constant) initial state. Note that this is really a special case of a 2-stage strategy: initial treatment is either immediate angiography followed if indicated by revascularization, or medical management with stress-testing, the second state measures 'ischemic changes' for 'conservatively' managed patients (for 'invasive' management, a second state is not needed to drive treatment), and the second treatment is the (delayed) angiography/revascularization procedure for patients who develop 'ischemic changes' during non-invasive testing. The ultimate outcome is overall survival.

The VA Cooperative Study #006 (13) compared inpatient and outpatient geriatric evaluation and management units (GEM) on overall survival, randomizing patients to GEM/non-GEM inpatient care and to GEM/non-GEM outpatient care. The 4-way randomization was done at the outset, but of course the second stage randomization only had an effect in survivors of the inpatient stay. Since only vital status was used to drive treatment, this is an example of the states and Y measuring the same outcome.

The four DLBCL ATS described above could have been tested in a 4-way randomized trial in which patients were assigned before induction to the sequence they would follow. One drawback to this approach is that it creates an opportunity for patients to fail to adhere to their assigned ATS after induction but before maintenance. For example, it might be that patients who do not tolerate R-CHOP well but do have a response will be more likely to opt for OBS as a maintenance treatment, even if they are assigned to RM at the original randomization. If this happens often, there may be no satisfactory randomization-based inference from the study. On efficiency grounds, it also seems attractive to randomize in two steps, since the outcomes of the first treatment are 'pre-randomization' with respect to the second treatment assignment, and can be used in 'forced balance' schemes to improve the balance of patient characteristics across the second randomization. As it happens, the trial (5) of the 4 ATS for DLBCL described above did use a 2-stage randomization. The generalization of this 'second randomization' leads to the SMAR trial.

THE SEQUENTIAL MULTIPLE ASSIGNMENT RANDOMIZED TRIAL

The four ATS for DLBCL in our example were compared in a SMAR trial, with 2 year failure-free survival (FFS) as the primary outcome. For our purposes here we will ignore certain issues with the data, notably due to post-randomization exclusions caused by delayed pathologic findings. After exclusions, there were 267 R-CHOP and 279 CHOP assignments, and a second randomization in responders, 174 to MR and 178 to OBS. If we use the algebra of the example, but for FFS instead of OS, we can simplify the example considerably, because the FFS in the patients who do not achieve response after induction can be taken to be zero, by definition. The first backward induction step reveals that MR provides a higher FFS than OBS in both the CHOP responders (0.74 vs 0.45) and the R-CHOP responders (0.79 vs 0.77), so the optimal maintenance treatment is always MR, and the 'optimal benefit to go' is 0.79 for R-CHOP responders and 0.74 for CHOP responders. Then we can take a step back to the first stage and observe that since $z_1 = w_1 = 0$ by definition of FFS, we have $r = (0.77)(0.79) = 0.608$ and $r' = (0.76)(0.74) = 0.562$ so R-CHOP induction has the best combination of response and optimal

benefit to go. Thus, the evidence favors the use of rituximab in both induction and maintenance. The authors of the paper analyze the results in several different ways, but come to the same general conclusion. It is possible to provide statistical uncertainties (standard errors, confidence intervals, P-values) for the contrast of these results among all 4 ATS; these methods take account of the 'data sharing' in the estimates; all the patients who are randomized to R-CHOP and respond contribute data to the estimated results of both ATS that start with R-CHOP (and similarly for CHOP) (6,7,15,16).

The analytic methods for SMAR trials are formally the same as those needed for correct data mining of non-experimental data; but NUC is guaranteed for SMAR trial data, and the 'curse of dimensionality' can be avoided by deliberate judicious choice of limited options for the ATS tested in a trial this is not generally possible for the data miner, who needs to take her ATS as she finds them.

ATS AND THE INTENTION TO TREAT PRINCIPLE

The concept of ATS clarifies some of the issues surrounding the intent-to-treat (ITT) principle in clinical trials. Under this principle, the outcomes of patients in a trial are to be summarized and contrasted as the original randomization dictates, not (for example) according to treatments actually received or by restricting to patients who adhered to the assignment. The principle has been defended as providing inferences that are grounded in the randomization, and that refer to the 'pragmatic' requirement to compare the 'policies' of 'try to give treatment A' versus 'try to give treatment B'. The criticisms of ITT focus on the unreality of the comparison of such treatment policies when a large proportion of patients drop out of assigned treatments or into treatments to which they were not assigned by the randomization.

In ATS terms, the initial randomization assigns the first treatment. From that point on, each patient in the trial is involved with an unspecified ATS that determines subsequent treatments on the basis of changes in the patient's and clinician's state of equipoise. The ITT comparison averages over these ATS, and to the extent that a 'consumer' of the trial results has a different mix of ATS in mind, there will be dissatisfaction with the ITT comparison. For example, if all the non-adherence to protocol comes from the incidence of an unacceptable and unavoidable side effect of the experimental treatment, in a fraction of the population (not identifiable in advance) and the default second line treatment is clear to all, then the ITT principle is a sure guide to the true value of the treatment. This is because the underlying ATS are clear, and there are only two – the one that begins with the experimental treatment and then defaults to a standard if the side effect occurs, and the other begins and stays with the standard (control) treatment. As a trial departs from this neat scenario, the underlying ATS become less clear and more various, so averaging over them makes less and less definite sense to consumers.

For example, the ITT principle in the VANQWISH trial averages over the way that clinicians interpreted the results of stress testing in the conservative arm (what level of ischemic change led to angiography) and in both arms, over the way that angiographic results led to revascularization decisions (whether to intervene, what vessels to open, and what method to use). The optimal choice of conservative versus invasive strategy depends on how well the stress testing and delayed intervention 'rescues' patients who turn out to need interventional treatment. In the VANQWISH trial, the contrast between the two strategies (which was randomized) turned on how well the implicit ATS were implemented by the team at the study sites, and this was not under experimental control. In the event, the 'rescue' worked exceedingly well, so the 'conservative' strategy was overall the best one (since it kept about 2/3 of the patients from being catheterized during the index hospitalization, preventing complications from the diagnostic procedure and subsequent intervention), but one could imagine a situation in which the second stage was not so successful.

Thus, any clinical trial that uses the standard ITT method to compare treatments on long term definitive outcomes, effectively averages over the success or failure of the intervening (perhaps uncontrolled) treatments that are applied as the patient's course dictates during the trial. Thus, even though the comparison among randomized initial treatments (A_1) is not myopic (it uses Y as the criterion for success or failure), its interpretation depends on the specific mix of subsequent uncontrolled choices of A_2, and so forth (one might call these ITT comparisons 'hyperopic' since they can't see the intervening treatment variation). Consideration of the nature of the implicit ATS that will be averaged over can help define the value of the ITT analysis, and may lead to making them more explicit, perhaps by standardizing those parts of the ATS that have good 'defaults'. The next step involves taking the parts of the ATS that cannot be standardized (because there is disagreement) and randomizing over them in a SMAR trial. By doing so one brings subsequent treatment variation under experimental control, allowing surer inference about the primary treatments and also creating the possibility of inference about optimal choices for the subsequent options.

ATS AND EARLY INDICATORS OF RESPONSE

Thall (16) describes the difficulties of current designs early phase trials that use an 'early response indicator', showing that it is highly risky to use even a good early response indicator to select treatments to take forward, or to make the 'go/no-go' decision for a definitive trial. He points out the connection between more efficient phase II–III trials and the idea of ATS. This is a promising line of research that may help to improve the performance of the phase II – phase III transition. One inference from Thall's work is that the definitive treatment outcome (here, Y) and the early indicator (which corresponds to a state) need to be studied together. Current practice (which uses the early indicator to select treatments for definitive testing and then separately uses the definitive outcome to confirm efficacy in a large trial), completely avoids the issue of how the two relate to each other statistically, leading to failed phase III trials.

SUMMARY

The idea of 'personalized medicine' has seized the imagination of many medical researchers. Instead of 'one pill cures all', personalized medicine offers a strategy, a list of treatment recommendations that depend to the genetic makeup of the individual patient. Of course, physicians have been 'speaking prose all their lives'; the genotype is just the latest version of the usual adaptation of treatment to diagnosis. But each individual's germ-line genome is fixed over her lifetime, while everything else changes, including her state of health or illness, and response to treatments. Other than time-limited, acute interventions that cure the patient, the typical treatment decision occurs in a sequence of similar decisions, in the context of long-term management of a chronic condition. As medicine improves its ability to stave off mortality, the result is a growing list of previously acute conditions that become chronic (HIV infection, diabetes, and many others). The result is a growing need for adaptive treatment strategies: rules that change the treatment as the patient's history of illness and response to previous treatment evolves. Other terms for this kind of strategy include 'dynamic treatment', 'treatment policy', and 'empirical treatment'.

An ATS has three components: a list of states that capture the important current features of the patient's illness, a list of treatments from which to choose the one to apply at the current time, and a rule for choosing a new treatment on the basis of the history of current and past states and past treatments. The iterative application of the rule creates the ATS out of the available treatments, but the state descriptions are also fundamental to the ATS since they capture the information that is used to drive the decisions. The states and rules work together much the

way diagnostic technology and treatment guidelines do, the former defining the categories of patients and their diseases as a basis for the actions prescribed by the latter.

To evaluate an ATS one needs at least one metric of success (or failure), that captures the important features of the overall outcome of the patient. This can be as simple as the span of the patient's life from the start of the strategy, or as complex as a patient-centered utility based on quality of life integrated over that span. An ATS may be based on early markers of treatment effect, but the intent is to influence an overall ultimate outcome.

One of the benefits promised by the 'electronic medical record' (EMR) is that the patient's history of treatments and responses will be readily available to the treating clinician, making it possible to implement a coherent ATS. Presumably, mining the data that is embedded in the EMR can shed light on the natural variation of ATSs for a given condition, and perhaps suggest candidate strategies for testing in experiments. But randomization is still the key to reliable inference about treatment effects, especially in the context of adaptive treatment strategies. Testing the effectiveness of different ATS and finding optimal ones can challenge the current methods of evidence-based medicine. Methods for evaluating and comparing ATS are evolving, including experimental and observational paradigms for obtaining evidence, statistical methods for analyzing data relating to ATS, and issues relating to research ethics. The ATS perspective may also help resolve the difficulties presented by the ITT principle, and make it possible to do studies that have greater ecological relevance. This article is an early review of the state of the art of ATS, from the perspective of a methodologist interested in the specification, assessment, comparison, and optimization of ATS.

Acknowledgments

This work was supported by a grant from the National Institutes of Health (R01 MH01481) to Stanford University.

Terms/Definitions list Acronyms list

ATS
: Adaptive Treatment Strategies

SMAR
: Sequential Multiple Assignment Randomized

ITT
: Intention-to-treat

NUC
: No Unobserved Confounders

FFS
: Failure-Free Survival

EMR
: Electronic Medical Record

LITERATURE CITED

1. Lavori PW, Dawson R, Rush AJ. Flexible treatment strategies in chronic disease: Clinical and research implications. Biol Psychiatry 2000;48:605–614. [PubMed: 11018231]
2. Lavori PW, Dawson R. Developing and comparing treatment strategies: An annotated portfolio of designs. Psychopharm Bull 1998;34(1):13–18.

3. Lavori PW, Dawson R. A design for testing clinical strategies: biased individually tailored within-subject randomization. J Royal Statist Soc Series A 2000;163:29–38.

4. Murphy SA. An experimental design for the development of adaptive treatment *strategies*. Stat Med 2005;24:1455–1481. [PubMed: 15586395]

5. Habermann TM, Weller EA, Morrison VA, Gascoyne RD, et al. Rituximab-CHOP versus CHOP alone or with maintenance rituximab in older patients with diffuse large B-cell lymphoma. J Clin Oncol 2006;24:3121–3127. [PubMed: 16754935]

6. Lavori PW, Dawson R. Dynamic treatment regimes: practical design considerations. Clin Trials 2003;1:9–20. [PubMed: 16281458]

7. Murphy SA. Optimal dynamic treatment regimes. J Royal Statist Soc Series B 2003;65:331–366.with discussion

8. Moodie, EM.; Richardson, TS.; Stephens, DA. Biometrics. 2006. Demystifying optimal dynamic treatment regimes. (OnlineEarly Articles)

9. Robins JM. The control of confounding by intermediate variables. Stat Med 1989;8:679–701. [PubMed: 2749074]

10. Murphy SA, van der Laan MJ, Robins JM. Marginal mean models for dynamic treatment regimes. J Am Statist Assoc 2001;96:1410–1423.

11. Robins, JM. Optimal structural nested models for optimal sequential decisions. In: Lin, DY.; Heagerty, P., editors. Proceedings of the second Seattle symposium on biostatistics. New York: Springer; 2004.

12. Boden WE, ORourke RA, Crawford MH, et al. Outcomes in patients with acute non-Q-wave myocardial infarction randomly assigned to an invasive as compared with a conservative strategy. N Engl J Med 1998;338:1785–92. [PubMed: 9632444]

13. Cohen HJ, Feussner JR, Weinberger M, et al. Effectiveness of geriatric evaluation and management units and clinics. A VA Cooperative Trial. N Engl J Med 2002;346:905–912. [PubMed: 11907291]

14. Dawson R, Lavori PW. Sequential causal inference: Application to randomized trials of adaptive treatment strategies. Stat Med. 2007To appear

15. Lavori PW, Dawson R. Variance of the non-parametric G-estimator: Application to randomized trials of adaptive treatment strategies. Clinical Trials. 2007To appear

16. Thall PF. A review of Phase 2–3 clinical trial designs. Lifetime data analysis. 2007To appear

Table 1

Notation for an Adaptive Treatment Strategy

$S_1,S_2\ldots,S_K$	States used as inputs to the adaptive treatment decisions.
$A_1,A_2\ldots,A_K$	Treatment options at each stage.
$S_1,A_1,S_2,A_2,\ldots,S_K,A_K,Y$	The order of observation of states and treatments; Y is the overall outcome.
$d_1,\ldots,d_K$	The decision rules at each stage.
$d_1(S_1)$	The initial treatment, given the first state.
$d_2(S_1,A_1,S_2)$	The second treatment, given the first state, first treatment, and second state.
$d_k(S_1,A_1,A_{k-1},S_k)$	The treatment at stage k after a series of k states and k-1 treatments.
$E[Y(d)]$	The average outcome (for a population of patients), if all of them were treated with the ATS d.
$E[Y(d)]-E[Y(d')]$	The difference in average outcome of two ATS, d and d′.

Table 2

Table 2a: A 2-stage ATS for untreated diffuse large B-cell lymphoma in patients >60 years old

S_1	Constant over patients
A_1	Choice of induction therapy (CHOP or R-CHOP).
S_2	State of disease after induction therapy (response or not)
A_2	Maintenance therapy in responders (MR or OBS) or 'standard of care' in non-responders
Y	Vital status at 2 years after induction therapy

Table 2b: The backward induction for the DLBCL ATS

x_1	2-year survival under MR in patients who respond to R-CHOP
y_1	2-year survival under OBS in patients who respond to R-CHOP
$m_1 = \max(x_1, y_1)$	2-year survival under best maintenance option, B, in R-CHOP responders
u_1	2-year survival under MR in patients who respond to CHOP
v_1	2-year survival under OBS in patients who respond to CHOP
$m_1' = \max(u_1, v_1)$	2-year survival under best maintenance option, B′, in CHOP responders
f_1	probability of response to R-CHOP induction
h_1	probability of response to CHOP induction
z_1	2-year survival under SOC in patients who do not respond to R- CHOP
w_1	2-year survival under SOC in patients who do not respond to CHOP
$r = f_1 m_1 + (1 - f_1) z_1$	2-year survival in strategy beginning with R-CHOP and continuing with best second stage strategy in responders to R- CHOP
$r' = h_1 m_1' + (1 - h_1) w_1$	2-year survival in strategy beginning with CHOP and continuing with best second stage strategy in responders to CHOP
Myopic method	Choose best induction treatment (compare f_1, h_1), then choose maintenance strategy that has the best survival in the responders to that induction treatment