The International Journal of Biostatistics

Volume 6, Issue 1

2010

Article 6

A Comparison of Variable Selection Approaches for Dynamic Treatment Regimes

Peter Biernot, *McGill University*

Erica E. M. Moodie, *McGill University*

Recommended Citation:

Biernot, Peter and Moodie, Erica E. M. (2010) "A Comparison of Variable Selection Approaches for Dynamic Treatment Regimes," *The International Journal of Biostatistics*: Vol. 6: Iss. 1, Article 6.

DOI: 10.2202/1557-4679.1178

A Comparison of Variable Selection Approaches for Dynamic Treatment Regimes

Peter Biernot and Erica E. M. Moodie

Abstract

In estimating optimal adaptive treatment strategies, the tailor treatment variables used for patient profiles are typically hand-picked by experts. However these variables may not yield an estimated optimal dynamic regime that is close to the optimal regime which uses all variables. The question of selecting tailoring variables has not yet been answered satisfactorily, though promising new approaches have been proposed. We compare the use of reducts—a variable selection tool from computer sciences—to the S-score criterion proposed by Gunter and colleagues in 2007 for suggesting collections of useful variables for treatment regime tailoring. Although the reducts-based approach promised several advantages such as the ability to account for correlation among tailoring variables, it proved to have several undesirable properties. The S-score performed better, though it too exhibited some disappointing qualities.

KEYWORDS: adaptive treatment strategies, dynamic treatment regimes, variable selection, categorical variables, binary outcomes, reducts, STAR*D, depression

Author Notes: Peter Biernot's studies were supported by a PGS-M Scholarship from the Natural Sciences and Engineering Research Council of Canada. Erica Moodie is supported by a Discovery Grant from the Natural Sciences and Engineering Research Council of Canada and an Operating Grant from the Canadian Institutes for Health Research. We would like to thank Dr. Joelle Pineau and Mr. Mahdi Milani Fard for their valuable insights. We thank the STAR*D team for supplying us with the data to carry out the investigation, and Dr. A. John Rush for comments on the manuscript. STAR*D was supported by funds from the National Institute of Mental Health.

1 Introduction

To effectively manage chronic or relapsing illnesses, physicians often adjust treatment to patient characteristics. For example, a physician must select the optimal first-line treatment, decide when to switch to another treatment, and choose the next treatment to try if the initial treatment failed (Murphy, 2005). A dynamic treatment regime, also known as an adaptive treatment strategy, is a function that takes information about a patient (baseline characteristics, treatments history, and intermediate outcomes) as inputs and returns an individually tailored treatment suggestion as the output (Murphy, 2003; Robins, 2004). Dynamic treatment regimes may be investigated via observational data assuming no unmeasured confounding or in sequential multiple assignment randomized trials (SMART) (Murphy, 2005; Collins et al., 2007). Unlike conventional trials where the treatment assignment is fixed at the beginning of the trial, in a SMART, treatments are re-randomized throughout the course of the trial and the randomization probabilities may depend on a participant's history. Rather than evaluating effectiveness of each treatment separately and potentially out of context, sequences of treatments are evaluated with the objective of optimizing some ultimate outcome (Lavori and Dawson, 2008). It may, for example, be the case that a treatment that is slow to act but leads to a greater number of remissions of disease is preferred over a treatment that provides relief to a smaller proportion of individuals but does so in a shorter time-frame.

There are two primary modelling tools used in statistics to estimate dynamic treatment rules: structural nested mean models and marginal structural mean models. Structural nested models (Robins, 1986; Murphy, 2005; Robins, 2004) do not define the class of regimens to be estimated – in particular, the variables used to tailor treatment need not be selected a priori – but rather seek to find the set of optimal rules among all rules using all (or a potentially large subset of all) variables, given the mean model. Marginal structural models may be used to estimate either statically-optimal treatment regimes (van der Laan et al., 2005) or optimal dynamic regimes (van der Laan and Petersen, 2007; Robins et al., 2008). In the marginal structural modeling approach to estimating dynamic regimes, the rules to be estimated must be supplied by the analyst.

Studies often collect information on a large set of easy-to-measure covariates (e.g., multiple surveys of mental health status and functioning) from which a smaller subset of variables must be selected for any practical implementation of treatment tailoring. While much of the statistical research on dynamic treatment regimes assumes a predefined class of regimes indexed by a small

number of pre-selected variables, interest here lies in defining the variables which index the class of regimes. That is, we do not wish to rely on the a priori hand-picking of tailoring variables. Rather, we seek to use the data to select a subset of tailoring variables that renders the estimated optimal dynamic treatment rule as close as possible to the optimal rule which uses all variables. We restrict our attention to the randomized setting.

In this paper, we compare two methods of selecting variables with which to adapt treatment for optimal dynamic treatment strategies. The first method considered is a variable ranking strategy based on the S-score, developed by Gunter et al. (2007). The S-score approach has many attractive features, but may not select the best combination of variables. We also considered a variable selection method based on reducts, an approach taken from the computer science literature, to see whether this might lead to an improved variable selection algorithm. The term “reduct” comes from rough set theory, a branch of computer science concerned with knowledge reduction. Rough set theory has been used in variety of applications: prediction of business failure, financial investment, bioinformatics and medicine, fault diagnosis, spatial and meteorological pattern classification, music and acoustics and feature selection (Shen and Jensen, 2007). Given the versatility of rough sets as evidenced by the above applications, we also investigated whether combination of S-scores and reducts would lead to a better variable selection algorithm.

The paper is structured as follows: Section 2 introduces notation and the two methods under consideration, S-scores and reducts. Section 3 applies the methods to the Sequenced Treatment Alternatives to Relieve Depression (STAR*D) clinical trial, a multistage randomized study of patients with nonpsychotic major depressive disorder. Sections 4 and 5 explore characteristics of the two variable selection approaches, and Section 6 concludes.

2 Background

2.1 Setting

An important step in the formulation of a dynamic treatment regime is the choice of variables to be used to decide when to switch treatments and to select the optimal treatments. Tailoring variables are variables that are used to adapt treatment to the individual. They often include baseline characteristics, intermediate outcomes, measures of adherence, and the presence of side effects (Murphy et al., 2007). Currently, much of the research on variable selection has focused on the supervised learning setting, where interest is in optimal

prediction of the response; less work has been done in the medical decision-making setting, where interest lies in choosing the best treatment for a given patient (Gunter et al., 2007). For the purpose of variable selection in a randomized setting where confounding is not present, Gunter et al. (2007) distinguish between predictive variables (used to increase precision of estimates) and prescriptive variables (used to adapt treatment strategies to patients, i.e. tailoring variables). Prescriptive variables must qualitatively interact with the treatment: that is, the choice of optimal treatment varies for different values of such variables. For example, consider a situation in which (i) for low levels of a patient’s hormone, the optimal action is to prescribe the drug A, (ii) for high levels of the same hormone the optimal action is to prescribe the drug B, and (iii) intermediate levels of the hormone require no treatment at all. The hormone level is a prescriptive variable, since different values of the variable result in different recommendations of the optimal treatment. Gunter et al. (2007) note that the usefulness of a prescriptive variable can be characterized by two quantities: the magnitude of the interaction and the proportion of the population for whom the optimal treatment changes given the knowledge of the variable.

Consider the following setting and notation, adapted from Murphy (2003). It is of interest to choose between *actions* (or, more commonly in the medical and epidemiological literature, between *treatments* or exposures): $A = (A_1, A_2, \dots, A_K)$ in K intervals, with domain (treatment options) $A_j \in \mathcal{A}_j$. Covariates measured on a subject at interval j , prior to action A_j , are denoted $X_j = (X_{j1}, X_{j2}, \dots, X_{jm})$, $j = 1, \dots, K$, with domain $\mathcal{X}_j$. Subsequent to the K actions, a response, R , is observed. The response is an unknown function of the action taken, the covariates, and possibly other unmeasured subject characteristics. A *dynamic treatment regime*, π , is defined to be a set of decision rules – one rule for each interval $j = 1, \dots, K$ – that maps the space of covariates X_j onto the space of actions A_j . More informally, a dynamic treatment regime provides a list of decision rules of the form “if an individual possesses certain characteristics (x_j) at the j th exposure interval, assign that individual to the specific action (treatment) a_j .”

Note that in data from a randomized study, assignment to the action space is independent of subjects’ covariates; however, using the data, we wish to estimate dynamic treatment regimes to assign actions to future subjects given their covariates. Thus, the objective is to find the dynamic treatment regime that will maximize the expected response obtained by following that regime.

2.2 S-scores

As proposed by Gunter et al. (2007), the *S-score* for a (univariate) variable X is defined as:

$$S_X = \hat{E} \left\{ \max_{a \in \mathcal{A}} \hat{E} [R|A = a, X] - \max_{a \in \mathcal{A}} \hat{E} [R|A = a] \right\}. \quad (1)$$

The S-score of a variable captures the expected increase in the response that is observed by adapting treatment based on the value of that variable. S-scores combine two characteristics of a useful prescriptive variable: the interaction of the covariate with the treatment and the proportion of the population exhibiting variability in the covariate's value. A high S-score for a covariate is indicative of a strong qualitative interaction between the variable and the treatment, as well as a high proportion of patients for whom the optimal action would change if the value of the variable were taken into consideration.

Gunter et al. (2007) used this scoring method to rank potential prescriptive variables in a one-interval setting. The performance of the S-score method was found to be superior to the standard lasso in terms of consistent selection of a small number of variables from a large set of covariates of interest. However in the implementation of the S-score ranking, each variable is evaluated separately, without taking into account potential correlation between variables. Two variables that are highly correlated may have similar S-score but may not both be necessary for decision making. Of course, the S-score may be modified in a straight-forward fashion to examine the utility of sets of variables given the use of other variables. The S-score could be used sequentially, first as in equation (1) to select the variable, X , with the highest score, then select a second variable with the highest S-score given the use of X as a prescriptive variable, using the altered definition

$$S_{Z|X} = \hat{E} \left\{ \max_{a \in \mathcal{A}} \hat{E} [R|A = a, X, Z] - \max_{a \in \mathcal{A}} \hat{E} [R|A = a, X] \right\}$$

and so on. In this paper, we compare the single-variable ranking strategy approach as implemented by Gunter et al. (2007) to an approach specifically designed to automatically select sets of covariates.

2.3 Reducts

The concept of *reducts*, derived from rough set theory in computer science, is another approach to variable selection. The notation as outlined by Swiniarski and Skowron (2003) was modified to conform with the setting and notation

described above. In rough set theory, the data from a sample are represented by a table where rows represent objects (or individuals), $U = (u_1, u_2, \dots, u_n)$, and columns represent all measured covariates, $\mathbf{X}$, where $\mathbf{X}$ is a finite and non-empty set of variables. The variables can be further partitioned into *condition* and *decision* variables. In our application we will consider the response, R , to be a decision variable while the remaining variables, B , will be considered condition variables.

Define the indiscernibility relation between two objects with respect to some set of variables C as follows: two objects are indiscernible with respect to C if and only if they have the same value for every variable in set C . Clearly, the indiscernibility relation is an equivalence relation and so it partitions the dataset into equivalence classes. Denote the partition by U/C and let $C(u)$ denote the equivalence class to which object u belongs. Further, define *C-lower approximation* of V ($V \subseteq U$) as follows:

$$C_*(V) = \{u \in U : C(u) \subseteq V\}$$

Thus, *C-lower approximation* of V is the set of all objects u that can be classified as certainly belonging to V based on the variables in C . Now, define a *positive region* of the partition U/R with respect to set of *condition* variables B as follows:

$$POS_B(R) = \bigcup_{V \in U/R} B_*(V)$$

The *positive region* is the set of all objects u that can be uniquely classified into one of the equivalence classes of partition U/R based on the variables in B . Given a dataset, with condition variables, B , and decision variable, R , a *decision-relative reduct* is a set of variables $B' \subseteq B$ such that $POS_{B'}(R) = POS_B(R)$ and for any $B'' \subset B'$ $POS_{B''}(R) \neq POS_B(R)$. That is, a reduct is a minimal subset of the condition variables, $B' \subseteq B$, that classifies objects (or individuals) into unique decision equivalence classes, U/R , as well as does the complete set of condition variables (B).

To further clarify these concepts, consider the simple covariate data in Table 1, with R the decision variable. Then:

$$U/R = \{\{u_1, u_3\}, \{u_2, u_4, u_5\}\}$$

$$U/(B_1, B_2, B_3) = \{\{u_1, u_2\}, \{u_3\}, \{u_4\}, \{u_5\}\} \Rightarrow$$

$$POS_{B_1, B_2, B_3}(R) = \{u_3, u_4, u_5\}$$

$$U/(B_1, B_2) = \{\{u_1, u_2\}, \{u_3\}, \{u_4\}, \{u_5\}\} \Rightarrow POS_{B_1, B_2}(R) = \{u_3, u_4, u_5\}$$

$$U/(B_1, B_3) = \{\{u_1, u_2\}, \{u_3, u_4\}, \{u_5\}\} \Rightarrow POS_{B_1, B_3}(R) = \{u_5\}$$

$$U/(B_2, B_3) = \{\{u_1, u_2, u_4\}, \{u_3\}, \{u_5\}\} \Rightarrow POS_{B_2, B_3}(R) = \{u_3, u_5\}$$

$$U/(B_1) = \{\{u_1, u_2\}, \{u_3, u_4\}, \{u_5\}\} \Rightarrow POS_{B_1}(R) = \{u_5\}$$

$$U/(B_2) = \{\{u_1, u_2, u_4, u_5\}, \{u_3\}\} \Rightarrow POS_{B_2}(R) = \{u_3\}$$

$$U/(B_3) = \{\{u_1, u_2, u_3, u_4\}, \{u_5\}\} \Rightarrow POS_{B_3}(R) = \{u_5\}.$$

(B_1, B_2) is a decision-relative reduct of the data since it is a minimal set of variables such that $POS_{B_1, B_2}(R) = POS_{B_1, B_2, B_3}(R)$. Using this reduced set of condition variables we can classify individuals into decision classes as well as we could with the complete set of condition variables. Thus, reducts help to eliminate redundant attributes while preserving the information in all measured covariates regarding the similarity of individuals in a sample. Note that reducts can only be implemented for categorical condition and decision variables.

Table 1: Sample covariate data.

Subject	Experience (B_1)	SAS Programmer (B_2)	Bilingual (B_3)	Hired (R)
u_1	1-2 Years	Yes	No	No
u_2	1-2 Years	Yes	No	Yes
u_3	3-5 Years	No	No	No
u_4	3-5 Years	Yes	No	Yes
u_5	5-7 Years	Yes	Yes	Yes

Often in applications, there are many reducts. When this occurs, a method of selecting between the reducts is required. For the specific application of selecting prescriptive variables, we compare two approaches: (1) selecting the variables that most frequently occur in the reducts and (2) selecting amongst reducts by choosing the set of covariates with the highest average S-score. We believed that the second, hybrid approach would combine the strengths of two attractive variable selection techniques, by using reducts to identify minimal sets of variables that are not strongly correlated with one another, and then selecting amongst the reducts via the average S-score to obtain a set of covariates that provided the best prescriptive variables.

3 Motivating example: STAR*D

The methods described in the previous section are applied in the context of the Sequenced Treatment Alternatives to Relieve Depression (STAR*D)

clinical trial, a multistage randomized study of patients with nonpsychotic major depressive disorder. STAR*D is the largest and longest randomized clinical trial of depression, and one of only a very small number of studies in any field to consider more than two stages of treatment (Rush et al., 2003).

3.1 The STAR*D trial design

STAR*D aimed to compare different treatment alternatives following lack of response to citalopram (CIT), an antidepressant in the class of selective serotonin reuptake inhibitors, in patients with major depression. Initially, 4041 patients aged 18-75 were enrolled in the study. STAR*D aimed to study the effectiveness of depression treatments in individuals with major depressive disorder, and encompassed a wide array of treatment options.

There were several ‘levels’ of treatment in the trial, corresponding to different stages of treatment randomization or choice. In Level 1, all participants were assigned to citalopram (CIT). Those without satisfactory improvement entered Level 2, where they were randomized to one of four medication switch options: sertraline (SER), bupropion (BUP), venlafaxine (VEN) or cognitive therapy (CT); or to one of three citalopram augmentation options: bupropion (CIT+BUP), buspirone (CIT+BUS) or cognitive therapy (CIT+CT). If treatment failed at Level 2, patients were re-randomized in Level 3. However, before proceeding to Level 3, patients for whom Level 2 treatment involved was cognitive therapy (CT or CIT+CT) and for whom that treatment was not beneficial were eligible for two switch options (Level 2A): venlafaxine (VEN) or bupropion (BUP). Level 2A was used to ensure that all patients who progressed to Level 3 did not respond to two different antidepressant medications. Patients without satisfactory response at Level 2 or 2A were eligible to proceed to Level 3, where they were randomized to two switch options, mirtazapine (MRT) or nortriptyline (NTP), or two augment options, lithium (Li) or thyroid hormone (THY), each used to augment their medication in Level 2 or 2A. Patients without satisfactory improvement at Level 3 were eligible to proceed to Level 4, where they were randomized to one of two switch options: tranylcypromine (TCP) or a combination of mirtazapine and venlafaxine (MRT-VEN). Table 2 summarizes the treatment randomization alternatives at each level.

Key results from the trial can be found in Gaynes et al. (2008), examining decision rules of the form: “If Level j treatment fails, what should the next Level treatment be?” All analyses were performed on an interval-by-interval basis. Based on efficacy alone, none of the treatment alternatives were found to be superior to the others. The data were also analyzed by Pineau et al. (2007) in an analysis that explicitly took into consideration the sequential nature

Table 2: Treatments at different Levels of the STAR*D trial.

Level 1	Level 2	Level 3	Level 4
CIT	SER, BUP, VEN, CT, CIT+BUP, CIT+BUS, CIT+CT Level 2A VEN, BUP	MRT, NTP, CIT+Li, CIT+THY, BUP+Li, BUP+THY, SER+Li, SER+THY, VEN+Li, VEN+THY	TCP, MRT+VEN

BUP = bupropion, BUS = buspirone, CIT = citalopram, CT = cognitive therapy, Li = lithium, MRT = mirtazapine, NTP = nortriptyline, SER = sertraline, TCP = tranylcypromine, THY = thyroid hormone, VEN = venlafaxine

of the trial, using instance-based reinforcement learning to suggest treatment allocation decision rules. Two tailoring variables were used: start-of-Level Quick Inventory of Depressive Symptomatology (QIDS) score and the rate of change in QIDS score in the previous Level; these variables were handpicked. It is therefore of interest to consider other, data-driven approaches to prescriptive variable selection.

3.2 Variable selection in STAR*D

The response of interest, R , was chosen to be the binary indicator of whether a patient had a remission from depression (as defined by a QIDS score no greater than 5) at any stage. The potential tailoring variables were mostly qualitative, with only a handful of continuous variables. Some variable processing was performed prior to analysis. All continuous variables were discretized (recall that reducts cannot be considered for continuous-valued variables). Variables indicative of increase or decrease were combined into a single ‘change’ variable. For example, the variables ‘QIDS appetite decrease’ and ‘QIDS appetite increase’ were combined into QIDS appetite change. Additionally, some very sparse categories of variables were collapsed where it was possible to form substantively reasonable groupings. Finally, with 10 treatment options at Level 3 and only 376 individuals, S-scores were not estimable for several variables. These variables were simply removed from consideration as potential tailoring variables at this level. In total, there were between 75 and 98 potentially tailoring variables at each level.

Level 1 was used in the trial primarily as a screening tool to identify difficult-to-treat depression; thus, with only one treatment under consideration, reducts can be used to identify predictive variables but there was no possibility of identifying tailoring variables. For that reason, no S-scores were obtained at Level 1 and only variables in reducts obtained in Levels 2 through

4 were considered for S-score calculation. For the purpose of this analysis, Level 2A was not included as the data from that stage were very sparse.

Reducts were obtained using the Rough Sets Exploration System (RSES, www.roughsets.org). Since the response was binary (remission or lack thereof), logistic regression was used to obtain the S-scores. Variables selected by each of the three methods (S-score ranking, variables selected by percent frequency of occurrence in reducts, and the hybrid reducts/S-score approach) are presented in Table 3. At each Level, the size of the smallest reduct was used to determine the number of variables to select.

Relative reduct frequencies for the top-ranked variables in Levels 2, 3 and 4 ranged from 34.4-75.5%, 32.5-59.3% and 21.6-48.0% respectively, while the S-scores for top-ranked variables ranged from 0.0607-0.1204, 0.1734-0.2270, and 0.0724-0.0858, for Levels 2, 3, and 4, respectively. This suggests that even when there is relative stability in the selection of a variable via reducts (since, for example, some variables appeared in more than three-quarters of the reducts at Level 2), the prescriptive power of the variables may be relatively low (S-scores no greater than 0.1204 at Level 2). Note also that the three approaches selected different sets of variables, although some variables were common to all three approaches. The two reducts-based variable selection approaches tended to be more similar to one-another: nine of ten, six of eight, and three of five variables selected for Levels 2, 3, and 4, respectively, were common between the approaches.

3.3 Remarks on the STAR*D analysis

This application of variable selection was undertaken for comparative rather than practical purposes. In particular, although the variable ‘research center’ (RC) was selected by all three approaches, which would imply the need for decision rules which depend on research center (in addition to individual patient characteristics), one would typically not opt to use this as a tailoring variable since decision rules based on these data could not then easily be generalized to research centers not included in the study. These results nevertheless underline the importance of controlling for center effects in any analysis of the data.

STAR*D followed a SMART design, in which subjects were randomly allocated to a given treatment when the previous treatment failed, as defined by treatment interruption due to intolerable side effects or lack of efficacy of the treatment in controlling illness. It is possible that the best dynamic treatment regime is one with decision nodes at times other than treatment failure; for instance, better long-term outcomes might be observed if treatment decisions

Table 3: STAR*D variables selected by each of the three methods considered.

Level	Relative reduct frequency (%)	Hybrid approach	S-score
Level 2	RC	RC	RC
	Self-outlook	Resid	Marital status
	Psychomotor	Degree	Self-outlook
	Appetite	Age	Side effects (int)
	Degree	Thous	Resid
	Sleep	Appetite	Headache
	Age	Self-outlook	Mid-noct. insomnia
	Thous	Psychomotor	Anxiety
	Resid	CNS effects	Rtown
	Poor conc.	Sleep	Appetite
Level 3	Self-outlook	Marital status	Marital status
	Sexual effects	Mid-noct. insomnia	Setting
	Side effects (int)	Self-outlook	OTHER
	Mid-noct. insomnia	Psychomotor	Self-outlook
	Weight	Side effects (int)	Resid
	Psychomotor	CNS effects	White
	CNS effects	Sexual effects	Involved
	Headache	Energy	Side effects (frq)
Level 4	RC	RC	Anxiety
	CGI	Suicidal	Diff. sleep
	Suicidal	Skin	Sleep
	Thous	No GU effects	Self-outlook
	Headache	CGI	RC

Age = Age at enrollment, Degree = Highest degree received, CGI = Clinical Global Impressions score, Side effects (frq) = Side effect frequency, Side effects (int) = Side effect intensity, CNS effects = Central nervous system side effects, Poor conc. = Poor concentration, Energy = Decreased energy, OTHER = Other symptoms, Resid = Current residence, RC = research center, Appetite = QIDS Appetite changed, Setting = Clinical Setting, Sexual effects = Sex side effects, Involved = QIDS Involvement, Skin = Skin side effects, Diff. sleep = Difficulty sleeping, Sleep = Sleep side effects, Mid-noct. insomnia = QIDS Mid-nocturnal insomnia, Psychomotor = QIDS Psychomotor slowing, Suicidal = QIDS Suicidal ideation, Weight = QIDS Weight change (last 2 weeks), Self-outlook = QIDS Self-outlook, Thous = Total number of persons in household, No GU effects = No genital/urinary symptoms, White = White race

were made monthly with treatment switching occurring if improvement is not observed after three months on a treatment. See Zhao et al. (2009) for a reinforcement-learning approach to estimating both optimal therapy and optimal timing of second-line treatment in a non-small cell lung cancer setting. Such a regime cannot be identified from the STAR*D data, where treatment changes were made only at therapy failure. Rather, the analysis undertaken

here attempted to identify the optimal set of prescriptive variables to select a new therapy for an individual at the point of current treatment failure. Specifically, the goal was to estimate the optimal dynamic treatment rule that was closest to the true optimal rule which tailored treatment based on all of the variables available (75 or more per Level) while allowing the data to select a more manageable subset of tailoring variables.

4 STAR*D experiments

In this section, we describe the results of several experiments using the STAR*D data to examine the properties of S-scores and reducts (with the percentage frequency rather than hybrid approach). The experiments were designed to consider the sensitivity of each method to the presence of noise variables in a dataset, the sensitivity of reducts to the omission of an apparently important tailoring variable (RC) from the dataset, and the behavior of the methods when variables in a dataset are highly correlated. The focus in these experiments is on the Level 2 data, which provided the maximum sample size, using the remission at any point in the trial as the response of interest, and selecting from among 50 potential prescriptive variables.

4.1 Experiment: Sensitivity to noise variables

In this experiment, we studied the sensitivity of both variable selection approaches criteria (S-scores and relative reduct frequency) to irrelevant noise variables in the data. We investigated the effect of adding noise variables with between two and 16 categories. The noise variables were labeled $N2, N3, \dots, N16$, where the number indicates number of categories in that generated variable. For example, a variable labeled $N3$ was generated as a random sample with replacement of the values (1,2,3). The noise variables were then added to the potential prescriptive variables.

Six of the ten top-ranked variables selected by S-scores were noise variables containing seven or more categories. The remaining four variables selected by S-scores were RC, marital status, QIDS self-outlook, and side effect intensity, which consisted of 12, 3, 4, and 3 categories, respectively. These results suggest that a variable with many categories can have a high S-score, even if it is irrelevant. The relative reduct frequency approach also selected many noise variables: seven of the top ten variables were irrelevant, and each of these had at least five categories. Thus it appears that both S-scores and reducts may select variables with a large number of categories over variables that are truly

prescriptive. In particular, RC may have been highly ranked by both methods due its large number of categories.

4.2 Experiment: Correlation between variables

In this experiment, we investigated how well the two variable selection methods handle correlation between categorical variables by examining the effect of misclassifying variables. We misclassified one each of a two-category, three-category, and four-category variable (poor concentration, sleep side effects, and mid-nocturnal insomnia, respectively) to varying degrees, to allow for varying degrees of similarity or correlation between the original variable and the newly formed, misclassified version of the variable. We compared the S-scores of the misclassified variables to the original, correctly classified versions of those variables. We also extracted reducts from a list of potential tailoring variables consisting of 50 potential prescriptive variables as well as the misclassified variables, and examined the frequency with which a misclassified version of a variable appeared in the same reduct as the original variable.

S-scores of the misclassified variables were occasionally larger than S-scores of the original variables. However, the variables used in the experiment had very low S-scores and thus neither these variables nor the highly correlated, misclassified versions of these were highly ranked by the S-scores approach. The greater the rate of misclassification, the more likely that a variable would appear in the same reduct with the original, correctly classified version of the variable. In a few instances, variables with low misclassification rates (5%) appeared in the same reduct as the corresponding original variable, indicating that a reduct can contain variables that are at least 95% similar to one another. Thus, this experiment suggests that both S-scores and reducts may be susceptible to selecting sets of prescriptive variables that contain variables with similar information.

4.3 Experiment: Collapsing research center

We investigated the effect of clustering the RC variable on S-scores and relative reduct frequencies. Without clustering, RC was the top-ranked variable in Level 2 selected by both variable selection methods. In this experiment, research centers were clustered into a four-category variable in two ways: informatively (based on center-specific recovery rates) and randomly. Using S-scores, the variable maintained its a ranking when clustering was based on recovery. Random clustering had the effect of lowering both the S-scores and the relative reduct frequencies. Typically, a random four-category clustering of

RC caused the variable to lose its place in the ten highest-ranked variables selected via the S-score. These results may suggest that RC does in fact contain useful prescriptive information and was not selected in the analysis of Section 3 purely because of its large number of categories. RC remained highly-ranked using the reducts approach to variable selection under both methods used to collapse the variable into four categories.

Finally, we investigated the effect of excluding research center on the relative reduct frequencies of other variables and their rankings. When RC was omitted from the list of potential tailoring variables, changes in variable rank of 30 or greater were not uncommon; this may indicate that relative reduct frequency is not a good measure of a variable’s prescriptive utility. Upon excluding RC from the list of potential tailoring variables, there was typically an increase in the size of the reducts. This could indicate that RC included relevant predictive information that needed to be accounted for by several additional variables; however perhaps the more likely explanation is simply that a 12-category variable – regardless of its prescriptive power – is a likely candidate for inclusion in reducts.

4.4 Summary of STAR*D experiments

The experiments performed indicated that in finite samples, S-scores may not be able to distinguish well between variables that are prescriptive, predictive, or simply irrelevant. There was also some suggestion that S-scores may tend to select many-categorized variables over those with less variability. Relative reduct frequencies appeared to perform badly by several measures. The method tends to select irrelevant variables, and preferentially selects variables with more categories. However, one positive attribute of the reducts approach observed is that it was relatively rare that very similar variables appeared together in the same reduct.

Sensitivity analyses to the results presented in this section were performed in the form of additional experiments in which the response variable was taken to be Level-specific remission rather than the general recovery outcome, corresponding to recovery at any stage. The conclusions from the sensitivity analyses were similar to the primary experiments (results not shown).

5 A simulation study

The results of the experiments performed on the STAR*D data prompted us to further investigate the properties of S-scores and reducts. A simulation

was designed to study the effect of adding irrelevant variables in addition to predictive and prescriptive variables to the list of potential tailoring variables.

We considered one four-category predictive variable (x_1), two prescriptive variables (x_2 and x_3 , taking on three and two values, respectively), and one binary treatment (Tx). The following samples sizes were considered: 100, 250, and 500; for each of those sample sizes, 100 datasets were generated. Additionally, a response variable, R , was created using the following linear predictor:

$$\begin{aligned} \text{logit}(\text{Pr}(R)) = & - 0.667 - 0.167 \times I[x_1 = 1] + 0.167 \times I[x_1 = 2] \\ & + 0.333 \times I[x_1 = 3] + 1 \times I[x_2 = 1] - 1.333 \times I[x_2 = 2] \\ & + I[x_3 = 1] + 0.667 \times I[Tx = 1] - 1.333 \times I[Tx = 1] \times I[x_2 = 1] \\ & + I[Tx = 1] \times I[x_2 = 2] - 1.333 \times I[Tx = 1] \times I[x_3 = 1]. \end{aligned}$$

In addition to the five main variables (x_1, x_2, x_3, Tx, R), irrelevant (noise) variables were added to each of the datasets. In total seven noise variables were added: three of which had four categories, two of which had three categories, and two of which had two categories.

As can be seen in Table 4, the greater the number of categories, the larger the average S-score for irrelevant variables; additionally, the average S-score of the predictive variable (x_1) was comparable to the average S-score for noise variables with the same number of categories. However, as sample size increased, the average S-score for the predictive and the noise variables decreased while the average S-score for the prescriptive variables (x_2 and x_3) remained approximately unchanged. The distribution of the S-scores of prescriptive variables appeared to follow a symmetric, unimodal distribution while the S-score distributions of predictive and noise variables had a long right tail with considerable mass around zero (results not shown).

Variables with a greater number of categories occurred more frequently in reducts, and both prescriptive and predictive variables appeared with similar frequency in reducts as noise variables with the corresponding number of categories. Adding the treatment variable to the reduct analysis resulted in only minor changes to the relative reduct frequencies of the predictive and the noise variables (Table 4). However, when interactions were also considered in the analysis, the relative reduct frequencies for the prescriptive variables and the treatment variable decreased markedly.

These simulations confirm many of the observations in the STAR*D experiments. The relative reduct frequency approach to variable selection appears to be biased towards selecting variables with greater numbers of categories, and the approach does not distinguish well between prescriptive, predictive,

Table 4: Simulation study results. Mean S-scores are based on 100 datasets; relative reduct frequencies are based on all the reducts.

No. values:	Predictive	Prescriptive		Noise								Tx	Interactions	
	4	3	2	4	3	2	4	3	2	4		2	3	2
Sample size	<i>Average S-score</i>													
100	0.036	0.071	0.041	0.038	0.028	0.015	0.043	0.032	0.014	0.041				
250	0.018	0.063	0.039	0.022	0.015	0.008	0.021	0.017	0.007	0.024				
500	0.014	0.073	0.045	0.014	0.008	0.006	0.014	0.009	0.006	0.016				
<i>Relative reduct frequencies (%) - Tx excluded</i>														
100	72.68	65.68	53.41	72.21	65.29	54.28	72.94	64.59	52.41	71.86				
250	88.60	82.15	67.93	88.76	82.31	66.12	86.28	81.16	64.96	83.97				
500	97.89	91.55	73.24	97.89	85.92	79.58	96.48	90.14	78.87	97.89				
<i>Relative reduct frequencies (%) - Tx included</i>														
100	67.80	62.98	51.35	67.95	60.96	52.60	69.34	60.92	49.61	67.69	53.28			
250	83.90	78.84	63.75	83.90	77.40	62.24	82.45	76.32	60.65	81.08	62.96			
500	93.94	91.18	73.28	93.39	83.20	73.00	93.66	87.05	74.10	90.91	77.96			
<i>Relative reduct frequencies (%) - Tx and interactions included</i>														
100	66.53	54.49	44.47	66.77	60.16	52.06	68.45	59.44	48.11	67.48	37.31	30.21	24.77	
250	84.02	69.17	54.56	83.57	75.82	62.45	82.51	76.10	60.49	81.72	42.15	33.85	28.26	
500	93.54	83.78	62.32	93.05	83.66	75.49	94.15	86.71	73.78	92.07	48.90	35.24	31.34	

and noise variables. The S-score approach, on the other hand, appears to perform better in that S-scores of prescriptive variables were higher than those of predictive and noise variables in large samples.

6 Discussion

We compared a recently proposed and highly promising variable selection technique, S-score ranking, with an approach that is novel in the statistical literature, reducts, for the purpose of selecting tailoring variables for optimal dynamic treatment regimes. The approaches were applied in the context of the STAR*D trial and further investigated via experiments with the STAR*D data and simulation. In this work, attention was aimed at the choice of prescriptive variables while the utility function was taken for granted. In the motivating example, the STAR*D data, recovery as a function of the QIDS score was identified by the trial physicians to be the outcome of interest. It is not unusual in the dynamic treatment regimes literature that the utility function is selected by substantive collaborators. However, the choice of utility may influence the method used to select prescriptive variables. Reducts, for example, are not appropriate for continuous outcomes.

The exposition here focused on the randomized trial setting. In an observational data setting, methods available typically require that all confounding variables are known and measured. In such cases, the confounding variables must be used by, for example, including these in regression models. This can be accomplished with S-scores by forcing the confounding variables into all models that assess other covariates for prescriptive power.

In our experiments and simulations, reducts did not perform satisfactorily: they selected irrelevant (noise) variables, were biased towards variables with more categories, and, contrary to our expectations, occasionally allowed for moderately correlated variables to appear in the same reduct. Most importantly, reducts did not select prescriptive variables over predictive or noise variables. The S-score approach performed better, in that prescriptive variables typically ranked higher than predictive and noise variables in large samples, although this method, too, exhibited limitations. For example, S-scores selected correlated and hence redundant variables. Variable selection for treatment tailoring remains an interesting area in the development of dynamic treatment regimes methodology, particularly as interest in multi-phase trials is growing.

References

- Collins, L. M., S. A. Murphy, and V. Strecher (2007). The multiphase optimization strategy (MOST) and the sequential multiple assignment randomized trial (SMART): New methods for more potent e-health interventions. *American Journal of Preventive Medicine* 32(5S), S112–118.
- Gaynes, B. N., A. J. Rush, M. H. Trivedi, S. R. Wisniewski, D. Spencer, and M. Fava (2008). The STAR*D study: Treating depression in the real world. *Cleve Clin J Med*. 75(1), 57–66.
- Gunter, L., J. Zhu, and S. A. Murphy (2007). Variable selection for optimal decision making. In *Proceedings of the 11th Conference on Artificial Intelligence in Medicine*.
- Lavori, P. W. and R. Dawson (2008). Adaptive treatment strategies in chronic disease. *Annual Review of Medicine* 59, 443–453.
- Murphy, S. A. (2003). Optimal dynamic treatment regimes. *Journal of the Royal Statistical Society, Series B* 65, 331–366.
- Murphy, S. A. (2005). An experimental design for the development of adaptive treatment strategies. *Statistics in Medicine* 30, 1455–1481.
- Murphy, S. A., D. W. Oslin, A. J. Rush, and J. Zhu (2007). Methodological challenges in constructing effective treatment sequences for chronic psychiatric disorders. *Neuropsychopharmacology* 32, 257–262.
- Pineau, J., M. G. Bellemare, A. J. Rush, A. Ghizaru, and S. A. Murphy (2007). Constructing evidence-based treatment strategies using methods from computer science. *Drug and Alcohol Dependence* 88(S2), S52–S60.
- Robins, J. M. (1986). A new approach to causal inference in mortality studies with sustained exposure periods - application to control of the healthy worker survivor effect. *Mathematical Modelling* 7, 1393–1512.
- Robins, J. M. (2004). Optimal structural nested models for optimal sequential decisions. In D. Y. Lin and P. Heagerty (Eds.), *Proceedings of the Second Seattle Symposium on Biostatistics*, New York, pp. 189–326. Springer.
- Robins, J. M., L. Orellana, and A. Rotnitzky (2008). Estimation and extrapolation of optimal treatment and testing strategies. *Statistics in medicine* 27, 4678–4721.

- Rush, A. J., M. Fava, S. R. Wisniewski, P. W. Lavori, M. H. Trivedi, and H. A. Sackeim (2003). Sequenced treatment alternatives to relieve depression (STAR*D): Rationale and design. *Controlled Clinical Trials* 25, 119–142.
- Shen, Q. and R. Jensen (2007). Rough sets, their extensions and applications. *International Journal of Automation and Computing* 4, 217–228.
- Swiniarski, R. W. and A. Skowron (2003). Rough set methods in feature selection and recognition. *Pattern Recognition Letters* 24, 833–849.
- van der Laan, M. J. and M. L. Petersen (2007). Causal effect models for realistic individualized treatment and intention to treat rules. *International Journal of Biostatistics* 3.
- van der Laan, M. J., M. L. Petersen, and M. M. Joffe (2005). History-adjusted marginal structural models and statically-optimal dynamic treatment regimes. *International Journal of Biostatistics* 1.
- Zhao, Y., M. R. Kosorok, D. Zeng, and M. A. Socinsky (2009). Reinforcement learning strategies for clinical trials in non-small cell lung cancer. Technical report, U.N.C. at Chapel Hill Department of Biostatistics Technical Report Series.