Variable Selection for Optimal Decision Making

Lacey Gunter[1,2], Ji Zhu[1], and Susan Murphy[1,2]

[1] Department of Statistics, University of Michigan, Ann Arbor, MI 48109, USA
[2] Institute for Social Research, University of Michigan, Ann Arbor, MI 48109, USA

Abstract. This paper discusses variable selection for medical decision making; in particular decisions regarding when to provide treatment and which treatment to provide. Current variable selection techniques were developed for use in a supervised learning setting where the goal is optimal prediction of treatment response. These techniques often leave behind small but important interaction variables that are critical when the ultimate goal is optimal decision making rather than optimal prediction. While prediction of treatment response represents a first step in finding optimal decisions, this paper points out some key differences between prediction and decision making applications. The paper presents two new techniques designed specifically to find variables that aid in decision making and demonstrates the utility of these techniques on both simulated data and on real data from a randomized controlled trial for the treatment of depression.

1 Variable Selection for Decision making

We consider variable selection in the simplest decision making setting in which one must decide between two actions (usually treatments). Prior to taking an action, we obtain observations about a subject $X = (X_1, X_2, ..., X_p)$, and using this information we choose an action A. We then receive a response, R, an unknown function (maybe random) based on the action taken and the observations and patient outcomes subsequent to the action. The response, R, gives us some indication of the desirability of the chosen action and is often referred to as a reward. A policy, π, is a decision rule mapping the space of observations, X, into the space of the actions, A. The goal is to find a policy π^*, which maximizes the response. Alternate policies are compared via the expected mean response denoted by

$$V_\pi = E_\pi [R]$$

V_π is called the (average) Value for the policy π [1]. The optimal policy, π^*, is defined as

$$\pi^* = \arg\max_\pi V_\pi = \arg\max_\pi E_\pi [R] .$$

A simple example is a clinical trial to test two alternative drug treatments. In this case, the observation vector would consist of baseline variables, such as the patient's background, medical history and current symptoms. The action would be the treatment assigned to the patient and the response could be the patient's

condition and symptoms after receiving treatment. The goal is to discover which treatment is optimal for any given patient, using the data obtained in the trial.

There are multiple reasons why variable selection might be necessary in a decision making application. First, many learning algorithms either cannot support an excessive number of variables or their performance degrades as the number of variables grows. Thus, careful variable selection could lead to better models for decision making. Second, due to limited resources, it may be crucial that policies use only a small number of variables. Researchers are often unsure which variables would be most cost-effective to collect. Variable selection techniques could help identify these variables.

Currently, variable selection for decision making in artificial intelligence is predominantly guided by expert opinion. Some predictive variable selection techniques, such as Lasso [2], are also being used [3]. In medical decision making applications such as clinical trials, the combination of predictive variable selection techniques and statistical testing of a small number of interaction variables suggested by expert opinion are most commonly used [4, 5]. Little research has been carried out to evaluate these techniques in decision making or suggest how they might be improved.

When selecting variables for decisions making, a distinction should be made between variables that will be included merely to aid in the process of estimating the policy from data as opposed to variables that will be included in the decision rules. *Predictive* variables are variables used to reduce the variability and increase the accuracy of the estimation. Variables that prescribe which decisions are optimal for a given group of patients are *prescriptive* variables [7]. For optimal estimation results, it is best to select both types of variables. However, in the implementation of decision rules only prescriptive variables must be collected.

For a variable to be *prescriptive*, it must have a qualitative interaction with the action [6]. A variable X_i is said to qualitatively interact with the action, A, if there exists at least two distinct, non-empty, sets within the space of X_i for which the optimal action is different. In other words, there exists disjoint, non empty sets $S_1, S_2 \subset space(X_i)$ for which

$$\arg\max_a E[R|X_i = x_{i1}, A = a] \neq \arg\max_a E[R|X_i = x_{i2}, A = a],$$

for all $x_{i1} \in S_1$, and $x_{i2} \in S_2$. These variables are useful because they help decipher which action is optimal for each individual patient.

To see why a qualitative interaction is necessary, it is helpful to look at a few illustrative plots. The first plot, Figure 1(a), shows a variable, X_1, that does not interact with the action. In this plot, the optimal action is $A = 1$. Knowing the value of X_1 is useful for predicting the response for a given action, but will have no effect on which action should be chosen, thus it does not play a part in the actual decision making.

Figure 1(b) shows a variable, X_2, that interacts with the action, A, but does not qualitatively interact with the action. The optimal action is still $A = 1$. As with the variable X_1, knowing the value of X_2 is useful for predicting the response for a given action, but will not effect which action should be chosen.

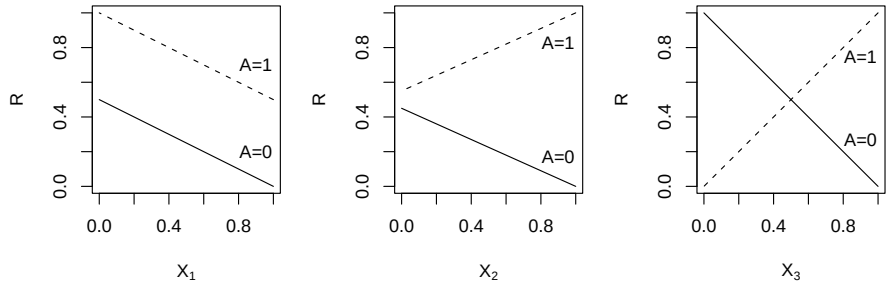

(a) No interaction (b) Non-qualitative interaction (c) Qualitative Interaction

Fig. 1. Plots demonstrating qualitative and non-qualitative interactions

The last plot, Figure 1(c), shows a variable, X_3, which qualitatively interacts with the action. The optimal action in this plot is $A = 0$, when $X_3 \leq .5$ and $A = 1$ when $X_3 > .5$. Knowing X_3 will impact the choice of action and likewise the response, thus it is useful for decision making.

The degree to which a *prescriptive* variable is useful depends on two factors:

1. *Interaction*: the magnitude of the interaction between the variable and the action. In other words, the degree to which the variable influences an action's effect on the response. For an action with two possible values, $A \in \{0, 1\}$, this is the degree to which the following quantity varies as x varies

$$E\left[R|X = x, A = 1\right] - E\left[R|X = x, A = 0\right] \tag{1}$$

2. *Proportion*: the proportion of patients whose optimal choice of action changes given a knowledge of the variable. That is, the proportion of patients for which the following holds:

$$\arg\max_a E\left[R|X = x, A = a\right] \neq a^* \tag{2}$$

where $a^* = \arg\max_a E\left[R|A = a\right]$.

Consider the plots in Figure 2. Figure 2(a) shows the relationship between R, A, and a variable X_4, with an underlying plot giving the distribution of the variable X_4. Figures 2(b) and 2(c) are similar to Figure 2(a), only they are for variables X_5 and X_6. Notice that both X_4 and X_5 have the same relationship with R and A. However, since the distribution of X_4 is centered at the point of intersection, half of the patients would do better choosing $A = 0$ rather than $A = 1$. Whereas, with X_5, a much smaller proportion of patients would benefit from choosing $A = 0$ rather than $A = 1$. For this reason, X_4 would be more useful in a decision making process than X_5. Now notice that the distribution of both X_4 and X_6 are centered at the point of intersection. However, the interaction between X_4 and A is much stronger than the interaction between X_6 and A.

Therefore, the effect of choosing the optimal decision is much greater given the variable X_4 than it is given the variable X_6. Thus X_4 is also more useful for decision making than X_6.

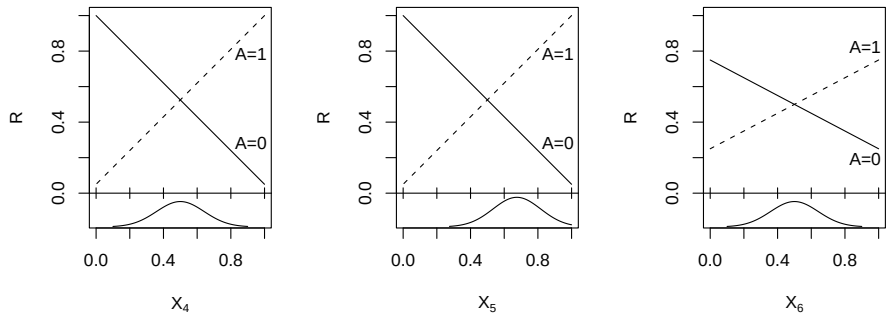

(a) Large interaction, large portion of change in action

(b) Large interaction, small portion of change in action

(c) Small interaction, large portion of change in action

Fig. 2. Plots demonstrating usefulness factors of qualitative interactions

Clearly *prescriptive* variables are important in prediction. However, in real world applications, we have found that individual variables, rather than interactions between variables, tend to explain most of the variability in the response and thus are most important for good prediction. This individual effect a variable has on the response is called the 'main effect' and most variable selection techniques are good at finding main effects. Furthermore, while treatment-covariate interactions do occur quite frequently in applications, most of these interactions are non-qualitative [6]. It is also not uncommon for some non-qualitative treatment-covariate interactions to have a moderately large effect on the overall variation of the outcome in real world settings, occasionally making these interactions important for good prediction. On the other hand, real world qualitative interactions tend to have a much smaller effect on the overall variation of the outcome. Thus, qualitative interactions can often be overlooked by variable selection techniques designed for prediction applications due to their small predictive ability and rare occurrence.

In the next section we present two new variable ranking methods specifically designed to find *prescriptive* variables. These methods rank the variables in X based on their potential for a qualitative interaction with the action.

2 Variable Ranking for Qualitative Interactions

The first ranking method is based upon the *interaction* and *proportion* factors for a qualitative variable discussed in the previous section, (see (1) and (2)). Assume

we have a data set of $j = 1, ..., n$ patients, with $i = 1, ..., p$ baseline observations for each patient, making up the $n \times p$ observation matrix X. Also assume we have an action, $A \in \{0, 1\}$, and a response, R, for each patient. Given an estimator of $E[R|X_i = x_{ij}, A = a]$ say $\hat{E}[R|X_i = x_{ij}, A = a]$, define the following quantities for $i = 1, ..., p$:

$$D_i = \max_{1 \le j \le n} \left(\hat{E}[R|X_i = x_{ij}, A = 1] - \hat{E}[R|X_i = x_{ij}, A = 0] \right) - \min_{1 \le j \le n} \left(\hat{E}[R|X_i = x_{ij}, A = 1] - \hat{E}[R|X_i = x_{ij}, A = 0] \right)$$

and

$$P_i = p_{i1}(1 - p_{i1}), \text{ where } p_{i1} = \frac{1}{n} \sum_{j=1}^{n} 1\{\arg\max_a \hat{E}[R|X_i = x_{ij}, A = a] = 1\}$$

where $1\{.\}$ is 1 one if '.' is true and 0 otherwise.

D_i is a measure of the magnitude of the interaction. P_i is a measure of the proportion of patients affected by a change in the optimal choice of action due to the inclusion of X_i. These two quantities can be combined to make a score, U_i, for ranking the variables. Define the score U_i by the following:

$$U_i = \left(\frac{D_i - \min_{1 \le k \le p} D_k}{\max_{1 \le k \le p} D_k - \min_{1 \le k \le p} D_k} \right) \left(\frac{P_i - \min_{1 \le k \le p} P_k}{\max_{1 \le k \le p} P_k - \min_{1 \le k \le p} P_k} \right) \quad (3)$$

So, the first variable ranking procedure will rank variables in terms of the U_i.

The second ranking procedure looks at the expected increase in the estimated optimal Value due to the knowledge of the variable X_i. It is similar in reasoning to the ideas discussed by Parmigiani in [8], dealing with the value of information. Define the score, S_i as

$$S_i = \hat{E}\left[\max_a \hat{E}[R|X = x_i, A = a] - \max_a \hat{E}[R|A = a] \right] \quad (4)$$

The scores, U_i and S_i can be used to rank the variables and have been defined generally to allow for different models of the relationship between the R, X, and A. In the numerical section that follows, we use a linear model to estimate the conditional expectation and obtain $\hat{E}$.

Since variable ranking is meant to be used as a first pass method and not a final variable selection method we suggest the following algorithm for variable selection.

1. Split the data set into K random sets to be used for cross-validation
2. Run a Lasso [2] on the main effects of X and A, but do not constrain the coefficient of A in the L_1 penalty function in (5). The standard Lasso minimization criterion is

$$\hat{\beta}(\lambda) = \arg\min_{\beta} \left\{ \sum_{j=1}^{n} (R_j - Z_j\beta)^2 + \lambda ||\beta||_1 \right\} \quad (5)$$

where the Z_j are the predictive vectors and λ is the Lasso penalty parameter.

3. Use K-fold cross-validation on the prediction error with the K random sets to select the value of the Lasso penalty parameter
4. Rank the variables in X using either U_i or S_i. Use the main effect variables suggested by Lasso with the selected penalty parameter from step 3 to decrease the variance in estimating $\hat{E}$
5. For $h = 1, ..., H$
   (a) Run a Lasso on the main effect variables chosen in step 3, A, and the interactions between A and the top h ranked variables from step 4. Again, do not constrain the coefficient of A in the L_1 penalty function in (5)
   (b) Select the value of the penalty parameter using K-fold cross-validation on the estimated Value with the K random sets.
      i. To estimate the Value first fit the variables suggested by the Lasso model for a given value of the penalty parameter into the chosen estimator of $\hat{E}$
      ii. Estimate the optimal policy, $\hat{\pi}^*(x) = \arg\max_a \hat{E}[R|X = x, A = a]$
      iii. Estimate the Value of $\hat{\pi}^*$ by:

$$\hat{V}_{\hat{\pi}^*} = \sum_{j=1}^{n} \frac{R_j 1\{a_j = \hat{\pi}^*(x_j)\}}{P(A = a_j|X = x_j)}$$

      where $P(A = a_j|X = x_j)$ is determined by the policy used to generate the actions in the data
6. Select the set of variables suggest by the Lasso model in step 6 that resulted in the highest estimated Value

In the section that follows, we reference this algorithm as either New Method U or New method S depending on the scoring function, U_i or S_i, used in step 4.

3 Numerical Results

To test the performance of the new techniques, a variety of simulations were conducted. The generative models contained at most one qualitative interaction and sometimes contained one or more non-qualitative interactions. Our methods were compared to the use of Lasso [2] on the main effects of X, A and all interactions between X and A. We did not constrain the coefficient of A in the L_1 penalty function and the penalty parameter was chosen by crossvalidation on the prediction error of the Lasso model. We reference this method as the 'Standard Lasso'. Due to limited space we briefly summarize our findings. For detailed simulation design and results please see the extended version of this paper [9].

In the simulations, we found that the two new methods selected the qualitative interaction (if one existed) more than any other interaction. The standard Lasso selected interactions with the highest predictive ability. While it was not uncommon for the standard Lasso to select the qualitative interaction a high number of times, many non-qualitative interactions were selected with equivalent and sometimes higher rates. The new methods also tended to give equivalent if not slightly higher average Values on test sets than the standard Lasso.

3.1 Nefazodone CBASP Trial

For most medical conditions there are multiple treatment options available. It is not always clear to the prescribing clinician which treatment will work best for a particular patient. In fact, some treatments might not work at all for certain patients. The methods presented in this paper can be used to aide in discovering patient characteristics that help determine the best treatment for the patient. To demonstrate this, we used the methods on data from a depression study to give insight on what variables might determine the optimal depression treatment for each patient.

The Nefazodone CBASP trial was conducted to compare the efficacy of three alternate treatments for chronic depression. The study randomized 681 patients with non-psychotic chronic major depressive disorder (MDD) to either Nefazodone, cognitive behavioral-analysis system of psychotherapy (CBASP) or the combination of the two treatments. Patients were given outpatient treatment for 12 weeks with assessments taken throughout the study. For detailed study design and primary analysis see [10]. A large number of baseline variables were collected on each patient. We consider the $p = 64$ variables listed in Table 1; these variables make up the X matrix. The outcome, R, was derived from the last observed 24-item Hamilton Rating Scale for Depression (HRSD) questionnaire [11], post treatment. Data from the 657 patients for which this response was available were used. For simplicity, we compared each individual treatment with the combination treatment; the first time with the action varying between the combination treatment and Nefazodone alone, and the second time with the action varying between the combination treatment and CBASP alone.

Since the techniques presented in this paper rely on cross-validation, it is conceivable that slightly different results could be obtained with different cross-validation splits of the data. Thus it is desirable to have a measure of this variability. To account for this, we took 100 bootstrap samples [2] of the data, ran the two new methods presented above and the standard Lasso on each bootstrap sample and recorded the interaction variables that were selected from each bootstrap sample and the direction of the slope of the interaction. This allowed us to obtain a percentage of time each interaction variable was selected by each technique. The percentages in Table 1 are the absolute value of the number of times an interaction was selected with a positive sign of the coefficient of the slope minus the number of times an interaction was selected with a negative sign of the coefficient. This adjustment helps reduce the number of spurious interactions.

The results of our analysis are listed in Table 1 and shown in Figures 3 and 4. Figure 3 shows plots of the adjusted percentage of time each variable was seleted by the standard Lasso and the two new methods U and S when the action varied between the combination treatment and Nefazodone alone. The x-axis is labeled with the variable numbers in Table 1. Table 1 also lists the adjusted percentage of time each variable was selected for the 3 methods. In Figure 3 we can see that with the standard Lasso, a large number of the variables were selected at least 20% of the time. It is somewhat difficult to decide how many variables to include

Table 1. Variable selection results for Nefazodone CBASP data. In the columns below, L stands for the standard Lasso, U stands for the new method U, and S stand for the new method S.

Variable	**Comb or Drug** % chosen			**Comb or CBASP** % chosen		
	L	U	S	L	U	S
1 Gender	21	0	0	16	0	0
2 Racial category	11	1	1	7	6	7
3-4 Marital status	2,13	0,1	0,0	10,3	0,0	0,0
5 Body mass index	2	1	1	5	0	0
6 Age in years at screening	9	0	0	5	0	0
7 Treated current depression	2	0	0	5	0	0
8 Medication current depression	23	0	0	9	0	0
9 Psychotherapy current depression	16	0	0	12	2	2
10 Treated past depression	16	1	1	5	0	0
11 Medication past depression	10	0	0	17	0	0
12 Psychotherapy past depression	38	0	0	29	0	0
13 Age of MDD onset	5	0	0	2	0	0
14-16 Number of depressive episodes	16,22,14	0,0,0	0,0,0	11,8,8	0,0,0	0,0,0
17 Length current episode	14	2	2	9	5	5
18-19 MDD type of current episode	19,18	1,1	1,0	13,16	0,2	0,1
20-21 MDD current severity	9,3	0,1	0,1	4,13	0,0	0,0
22-23 MDD chronic status	15,20	0,0	0,0	5,3	1,1	0,0
24 MDD threshold frequency	8	0	0	0	0	0
25 Dysthymic disorder current	4	0	0	1	0	0
26 Dysthymia initial onset	23	0	0	3	1	2
27 Length current dysthymia episode	1	0	0	5	4	4
28-29 Alcohol	28,46	12,17	10,15	23,19	5,4	3,1
30 Drug	1	0	1	21	0	0
31-32 Social phobia	11,28	2,3	1,2	9,2	0,0	0,0
33-34 Specific phobia	3,32	0,6	0,5	4,4	6,2	6,3
35 Obsessive compulsive	51	0	0	11	2	3
36-37 Post traumatic stress	9,2	0,0	0,0	16,5	10,1	10,0
38-39 Generalized anxiety	28,15	0,0	0,0	7,17	3,3	3,2
40 Anxiety disorder NOS	26	0	0	10	1	1
41-42 Panic disorder	26,27	0,0	0,0	31,11	6,5	4,4
43 Body dysmorphic current	8	0	1	1	3	3
44 Anorexia or Bulimia nervosa	22	3	2	8	1	1
45 Global assessment of function	5	0	0	2	0	0
46-47 Main study diagnosis	5,8	0,0	0,0	6,10	0,1	0,2
48 Severity of illness	1	0	0	2	0	0
49 Chronic or double depression	6	0	0	9	0	0
50 Total HAMA score	3	2	1	2	0	0
51 HAMA Sleep disturbance factor	3	0	0	1	1	0
52 HAMA Psychotic Anxiety Score	3	1	1	3	0	0
53 HAMA Somatic Anxiety Score	34	14	11	6	0	0
54 Total HAMD-24 score	4	1	1	5	0	0
55 Total HAMD-17 score	4	0	0	3	1	1
56 HAMD Cognitive Disturbance	10	0	0	7	1	0
57 HAMD Retardation Score	2	0	0	4	1	1
58 HAMD Anxiety/Somatic symptom	1	0	0	1	0	0
59 IDSSR Total Score	2	0	0	0	0	0
60 IDSSR Anxious depression type	5	0	0	11	0	0
61 IDSSR General/Mood Cognition	5	4	3	3	0	0
62 IDSSR Anxiety/Arousal Score	0	3	3	0	1	1
63 IDSSR Sleep score 1	10	0	0	0	2	1
64 IDSSR Sleep score 2	3	0	0	3	0	0

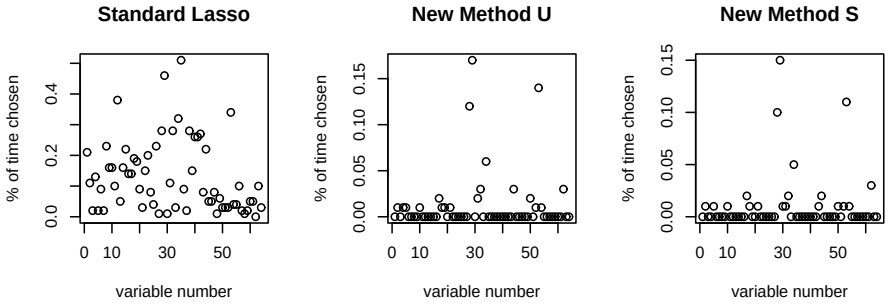

Fig. 3. Variable selection results for Nefazodone CBASP data where A varied between combination treatment and Nefazodone. In each plot, the x-axis is the variable number from Table 1, the y-axis is the adjusted percent of times variables were selected.

since there appears to be no obvious separation between the variables. The two new methods, selected 3 variables the most, the first two being indicators for *Alcohol* and the third being *HAMA Somatic Anxiety Score*. To contrast, the standard Lasso selected *Obsessive Compulsive Disorder* most frequently.

Figure 4 shows plots of the adjusted percentage of time each variable was selected by the three techniques when the action varied between the combination treatment and CBASP alone. Again, with the standard Lasso many variables were selected multiple times. However *Psychotherapy for past depression* and one of the indicators for *Panic Disorder* were selected most frequently. Neither of the new methods selected *Psychotherapy for past depression*. Both of the new methods, selected *Post Traumatic Stress Disorder* most often. With the standard Lasso, seven other variables where selected more often than *Post Traumatic Stress Disorder*.

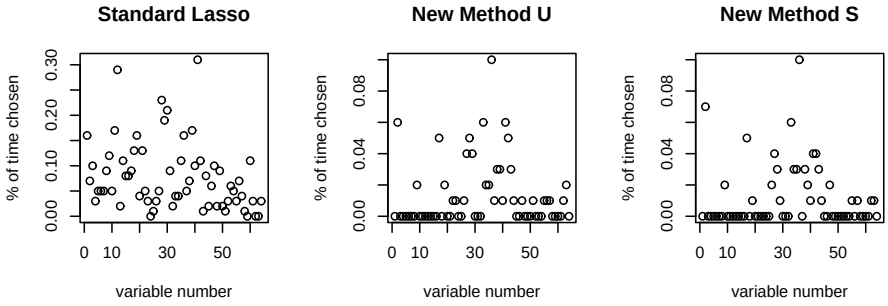

Fig. 4. Variable selection results for Nefazodone CBASP data, where A varied between combination treatment and CBASP. In each plot, the x-axis is the variable number from Table 1, the y-axis is the adjusted percent of times variables were selected.

4 Conclusion

In this paper, we discussed when a variable is important in decision making and why variable selection techniques designed for prediction may not perform well in a decision making setting. We presented two new techniques explicitly designed to select variables for decision making. We tested these methods in simulations and demonstrated them on a depression data set and the results look promising.

Acknowledgments We wish to thank Martin Keller and the investigators who conducted the trial 'A Comparison of Nefazodone, the Cognitive Behavioral-analysis System of Psychotherapy, and Their Combination for Treatment of Chronic Depression', for use of their data, and gratefully acknowledge Bristol-Myers Squibb for helping fund the study. We also acknowledge financial support from NIH grants R21 DA019800, K02 DA15674, P50 DA10075, and technical support from the Betty Jo Hay Chair in Mental Health, A. John Rush, MD at the University of Texas Southwestern Medical Center, Dallas.

References

1. Sutton R.S., Barto A.G.: Reinforcement Learning: An Introduction. MIT Press, Cambridge, MA (1998)
2. Hastie, T., Tibshirani, R., Friedman, J.: Elements of Statistical Learning. Springer-Verlag, New York (2001)
3. Loth, M., Davy, M., Preux, P.: Sparse Temporal Difference Learning using LASSO. In IEEE International Symposium on Approximate Dynamic Programming and Reinforcement Learning, Hawaii, USA (2006)
4. Krystal, J.H., Cramer, J.A., Krol, W.F.et al.: Naltrexone in the Treatment of Alcohol Dependence. New England Journal of Medicine **345** (2001) 1734-39
5. Reynolds, C.F., Dew, M.A., Pollock, B.G.et al.: Maintenance Treatment of Major Depression in Old Age. New England Journal of Medicine **354** (2006) 1130-38
6. Gail, M., Simon, R.,: Testing for Qualitative Interactions Between Treatment Effects and Patiwnt Subsets. Biometrics **41** (1985) 361372
7. Hollon, S.D., Beck, A.T.: Cognitive and cognitive behavioral therapies. In: Lambert, M.J. (ed): Garfield and Bergin's Handbook of Psychotherapy and Behavior Change: An Empirical Analysis, 5th edn. Wiley, New York (2004) 447492
8. Parmigiani, G.: Modeling in Medical Decision Making: a Baysian Approach. Wiley, West Sussex, England (2002)
9. Gunter, L., Murphy, S.A., Zhu, J.:Variable Selection for Optimal Decision Making. Unpublished manuscript. (http://www.stat.lsa.umich.edu/∼lgunter/aimelong.pdf)
10. Keller, M.B., McCullough, J.P., Klein, D.N.et al.: A Comparison of Nefazodone, the Cognitive Behavioral-analysis System of Psychotherapy, and Their Combination for Treatment of Chronic Depression. New England Journal of Medicine **342** (2000) 331-366
11. Hamilton, M.: Development of a rating scale for primary depressive illness. British Journal of Social and Clinical Psychology **6** (1967) 278-96