Statistic. Springer-Verlag, 1991.

[10] G.M. Provan and J.R. Clarke. Dynamic network construction and updating techniques for the diagnosis of acute abdominal pain. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(3):299–307, 1993.

[11] M. Ramoni and P. Sebastiani. Robust learning with missing data. Technical report, Knowledge Media Institute, The Open University, 1996.

[12] D.J. Spiegelhalter, P.A. Dawid, S.L. Lauritzen, and R.G. Cowell. Bayesian analysis in expert systems. *Statistical Science*, 8:219–283, 1993.

[13] D.J. Spiegelhalter and S.L. Lauritzen. Sequential updating of conditional probabilities on directed graphical structures. *Networks*, 20:157–224, 1990.

[14] J.A. Tatman and R.D. Shachter. Dynamic programming and influence diagrams. *IEEE Transactions on Systems, Man, and Cybernetics*, 20(2):365–379, 1990.

[15] R. Winkler. The assessment of prior distributions in bayesian analysis. *American Statistical Association*, 62:776–800, 1967.

[16] R. Winkler. *Introduction to Bayesian Inference and Decision*. Holt, Rinehart and Winston, Inc., 1972.

els. Therefore, automated learning of probabilities for those models are demanded. On the other hand, we may not have sufficient data to learn probabilities from. Based on our experiment with DynaMoL, we think that a feasible way out of this dilemma is through Bayesian learning, data abstraction and model abstraction. As our case study and comments from our expert show, model abstraction may lose some of the aspects in a decision problem, but the abstract model can be instructive and informative at a high level.

We observe that in dynamic decision problem modeling, prior distribution elicitation is not much more complicated than in static decision problem modeling. This is because the decision stages provide the expert with a context to compare the exponents of Dirichlet or Beta distributions across decision stages.

In our learning system, all the influence views are built by the modeler in collaboration with the domain expert. One of our future tasks is to learn influence-view structures from data or refine partial influence-view structures based on the data. Although there are quite a lot of people on that line, *e.g.* [2,5,12], it seems that dynamic decision models might be much more difficult to learn or refine. Another issue we are facing is that the table of conditional probabilities for the influence views are usually very large. In our case study, the table has about 4500 rows. How to store and retrieve the table of conditional probabilities efficiently is also a future research issue in our work.

ACKNOWLEDGMENTS

We would like to thank Susan Teo, Belinda Chang and Nancis Phoon of Singapore General Hospital for their generous assistance with the experiments, and Lau Aik Hiang for her effort in the data processing. The work is supported by an academic research grant RP950617 from the National University of Singapore and a postdoctoral fellowship from the National Science and Technology Board in Singapore.

REFERENCES

[1] J.R. Beck and S.G. Pauker. The markov process in medical prognosis. *Medical Decision Making*, 3:419–458, 1983.

[2] W. Buntine. Theory refinement on bayesian networks. In *the 7th Conference on Uncertainty in Artificial Intelligence*, pages 52–60. Morgan Kaufmann, 1991.

[3] R. Cowel, A. Dawid, and P. Sebastiani. A comparison of sequential learning methods for incomplete data. *Bayesian Statistics*, pages 581–588, 1996.

[4] A. Dempster, D. Laird, and D.B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Royal Statistical Society B*, 39:1–38, 1977.

[5] D. Heckerman, D. Geiger, and D. Chichering. Learning bayesian networks: The combination of knowledge and statistical data. Technical Report MSR-TR-94-09, Microsoft Research, 1994.

[6] S. Lauritzen. The em algorithm for graphical association models with missing data. *Computational Statistics and Data Analysis*, 19:191–201, 1995.

[7] Tze-Yun Leong. Dynamic decision modeling in medicine: A critique of existing f ormalisms. In *Proceedings of the Seventeenth Annual Symposium on Computer Applications in Medical Care*, pages 478–484. IEEE, 1993.

[8] Tze-Yun Leong. Multiple perspective reasoning. In *Proceedings of the 1996 International Conference on Knowledge Representation*. Morgan Kaufmann, 1996.

[9] G.M. Provan. Model selection for diagnosis and treatment using temporal influence diagrams. In *Selecting Models from Data*, volume 89 of *Lecture Notes in*

bution or may not get used to the elicitation procedure, and hence the answers from him might not be meaningful from the very beginning.

In our work, we used an "on-line" checking. The idea behind the on-line checking is actually a game between the expert and the modeler. Before interviewing the expert, we learn from the database the *sufficient statistics, i.e.* α'_{ij1} and α'_{ij2}, and the sufficient-statistics ratio, *i.e.*, $\alpha'_{ij1}/\alpha'_{ij2}$, for all i and j. The ratio of sufficient statistics serve as a checkpoint for the prior ratio, *i.e.* $\alpha_{ij1}/\alpha_{ij2}$. For example, if $\alpha'_{ij1}/\alpha'_{ij2}$ = 30/70, but the prior ratio $\alpha_{ij1}/\alpha_{ij2}$ = 80/20 based on the answers from the expert, then the modeler could interrupt the expert to see if any change for the α_{ij1} and α_{ij2} is necessary. This can sometimes lead the expert to reconsider his estimates.

6 LEARNING EXPERIMENTS

In our case study, we chose the data of Stage C patients (denoted as FU-DATA), and use Beta distributions as prior distributions for θ_{ijk}, since all the event variables are binary. Based on the elicitation method discussed in Section 5.2, we elicited all the Beta distributions associated with the events in our influence views. Let us take a dynamic local influence view, DLIV(*TR*, {S, *EOR*, *EOM*}, M&T, STAGE), from Figure 3. Some of the prior information for the DLIV is shown in Tables 2 and 3. Here, α_+ and α_- are the respective abbreviations for α_{ij1} and α_{ij2} in the elicitation question, and are provided by our expert. During the elicitation process, we also interrupt the expert when the ratio α_+/α_- differs from that of their corresponding sufficient statistics too much (recall discussions in Section 5.2).

After the prior distributions are obtained, we invoke the Bayesian learning system to learn the posterior distributions of θ_{ijk} and calculate their expectations. The learning results are output in a tabular format, as shown in Tables 2 and 3. Note that we have chosen six decision stages in our example.

Table 2: Prior Information and posterior probabilities

$E(\theta_+)$	$E(\theta_-)$	*STAGE*	*STATE*	*EOR*	*EOM*	*TR*	α_+	α_-
0.6200	0.3800	1	well	N	Y	+	30	70
0.6010	0.3990	2	well	N	Y	+	60	40
0.3836	0.6164	3	well	N	Y	+	60	40
0.6039	0.3961	4	well	N	Y	+	65	35
0.6041	0.3959	5	well	N	Y	+	70	30
0.6095	0.3905	6	well	N	Y	+	80	20

Table 3: Prior Information and posterior probabilities

$E(\theta_+)$	$E(\theta_-)$	*STAGE*	*STATE*	*EOR*	*EOM*	*TR*	α_+	α_-
0.9047	0.0953	1	well	Y	N	+	95	5
0.9000	0.1000	2	well	Y	N	+	90	10
0.9000	0.1000	3	well	Y	N	+	90	10
0.8932	0.1068	4	well	Y	N	+	85	15
0.9038	0.0062	5	well	Y	N	+	80	20
0.9019	0.0081	6	well	Y	N	+	80	20

In Table 2, θ_+ and θ_- respectively denote:

- Pr(*TR*=+|*EOR*=N,*EOM*=Y,**well**, M&T, *STAGE*, FU-DATA)
- Pr(*TR*=-|*EOR*=N, *EOM*=Y,**well**, M&T,*STAGE*,FU-DATA),

and $E(\theta_+)$ and $E(\theta_-)$ respectively denote the posterior expectations of θ_+ and θ_-. Similarly in Table 3, θ_+ and θ_- respectively represent:

- Pr(*TR*=+|*EOR*=Y,*EOM*=N,**well**, M&T, *STAGE*, FU-DATA)
- Pr(*TR*=-|*EOR*=Y, *EOM*=N,**well**, M&T,*STAGE*,FU-DATA),

and their corresponding expectations are denoted by $E(\theta_+)$ and $E(\theta_-)$.

The learned parameters can be verified by cross-validation, or by verifying and conducting sensitivity analyses on the optimal course of action derived from the resulting dynamic decision model.

7 DISCUSSION AND CONCLUSION

Dynamic decision analysis has received considerable study in recent years. Compared to its static counterpart, dynamic decision analysis is much more complicated, because many model parameters (e.g. probabilities) vary with time. The domain experts may not be able to assess probabilities with sufficient precision needed in dynamic decision modrior distribution of θ_{ij1} and θ_{ij2} is given by:

$$D\langle\Theta_{ij}|D, \alpha_{ij1}, \alpha_{ij2}\rangle = \frac{\Gamma\left(\sum_{k=1}^{2}(\alpha_{ijk} + \alpha'_{ijk})\right)}{\prod_{k=1}^{2}\Gamma(\alpha_{ijk} + \alpha'_{ijk})} \times \prod_{k=1}^{2} \theta_{ijk}^{\alpha_{ijk} + \alpha'_{ijk} - 1}$$

When the posterior distribution is calculated, the Bayes estimate (i.e. expectation) of θ_{ijk} is calculated by:

$$E\langle\theta_{ijk}|D, \alpha_{ij1}, \alpha_{ij2}\rangle = \frac{\alpha_{ijk} + \alpha'_{ijk}}{\alpha_{ij} + \alpha'_{ij}}$$

where $\alpha_{ij} = \alpha_{ij1} + \alpha_{ij2}$ and $\alpha'_{ij} = \alpha_{ij1} + \alpha_{ij2}$.

5•2 Assessing Prior Distributions

Now we discuss how to assess the exponents for Beta distributions, especially when decision stages are taken into consideration. The assessment issue was addressed by several authors from a static perspective. Based on our experience and previous works by other authors [15,16,5], we have devised a procedure for eliciting from the expert the exponents, *i.e.* α_{ij1} and α_{ij2}, associated with a binary variable X_i in an influence view.

The elicitation phase centers around *dynamic local influence structures* (DLIV) in influence views. A dynamic local influence view, DLIV(X_i, π_i, T, A), consists of an event (called the center), its parents, a decision stage, and an action. Figure 7 shows a DLIV.

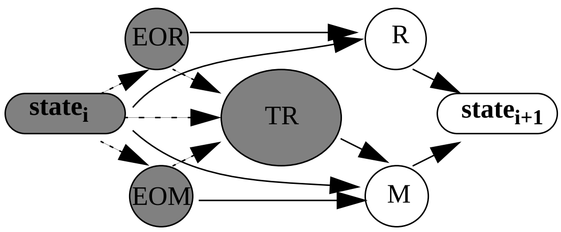

Figure 7: A DLIV with X_i=*TR*, *T*=i, and *A*=M&T

Suppose DLIV(X_i, π_i, T, A) is a dynamic local influence view structure. We fix a value for each of π_i, X_i and A, and these values serve as part of a "context" to ask the expert questions in. We then let T vary *increasingly* and for each value of T, we ask the expert to assess 1) a size of a sample which is roughly equivalent to his prior knowledge about DLIV(X_i, π_i, T, A) with the fixed values, and 2) the times that X_i takes its fixed value in the sample. This request is a *dynamic* version of the "equivalent sample size" technique (e.g. see [15,16,5]). The continuity in the values of T proves to be very useful for the expert to offer the exponents. We notice that the expert often thinks for moments before offering exponents for the first decision stage. But for other stages, he can give exponents immediately. As our expert told us, this is so because the exponent with respect to the initial stage provides a useful reference point to orient himself.

We also notice that the expert gets sometimes stuck in the assessment procedure, because he does not know the equivalent sample sizes for some pieces of his prior information. In two recent studies, we find that experts often like to fix the sample size to be 100 or 10. In that case, the elicitation question becomes "Assuming that you are given 100 cases, what is the number of cases in which X_i takes X_{i1} and π_i takes π_{ij}?" Suppose that the expert tells us the number is about α_{ij1}. Then, we roughly have:

$$\alpha_{ij2} = 100 - \alpha_{ij1}$$

To increase the precision of the exponents in some situations, the number 100 can be enlarged to 1000.

We can not expect the estimates offered by the experts are always reasonable. This is especially true if the expert is not familiar with probability theory, or the expert is not confident in the way his prior information is assessed. Thus, some sort of checking is necessary. Checking the estimates can be done after the whole elicitation is completed. But this "off-line" checking might not be efficient. The main reason is that the expert may misunderstand the meaning of a Beta distrience structure, the event variables in their structure may have different probability distributions.

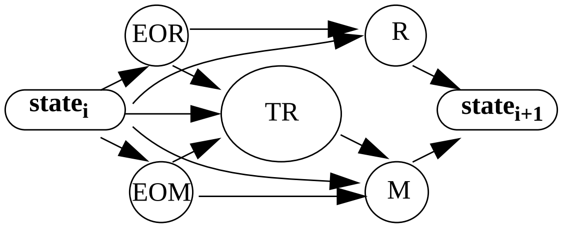

Figure 6: Influence view for M&ST and M&AT

We do not use the specific actions and their influence views, because we find that patients with metastatic cancers have little specific treatment information kept in the databases under study. That is, we found two data-sparse areas corresponding to AT and ST. We need to group the two areas together to form a data-rich one.

Nevertheless, the action M&T and its influence view are still meaningful in the sense that it offers a general guideline for the doctor at a given decision stage, and the doctor will take other aspects of patients into account when determining which specific treatment (*e.g.* AT or ST) to prescribe.

5 THE LEARNING FRAMEWORK

5•1 The Learning Mechanism

In the Bayesian learning community, the Dirichlet or Beta distributions are commonly used as prior distributions for model parameters. This strategy simplifies the learning process [12,5, 11], because of mathematical properties of those distributions[16]. However, the question is whether Dirichlet or Beta distributions are "rich" enough in the sense that they can capture a variety of an individual's prior information. Fortunately, as noted by Winkler [16], these distributions can possess "different locations, dispersions, shapes, and so on".

We first present the general definitions of variables in the learning framework. Let X_i be an event with r_i possible values. We use π_i to denote the set of parents of X_i , and assume that the values of π_i can be ordered somehow (theoretically, this can always be done since influence views are acyclic graphs). We use θ_{ijk} to parameterize the probability that X_i takes its k-th value (denoted by X_{ik}), given π_i taking its j-th value (denoted by π_{ij}), an action *A*, and a decision stage *T*. Note that we omit *A* and *T* from θ_{ijk} for simplicity, but one should understand θ_{ijk} based on its context. Let Θ_{ij} denote the collection of those θ_{ijk}.

We assume that the variables Θ_{ij} have a Dirichlet distribution with exponents $\alpha_{ij1}, \ldots, \alpha_{ijr_i}$ .The general formula of the Dirichlet distribution is then:

$$\langle\Theta_{ij}|\alpha_{ij1}, \ldots, \alpha_{ijr_i}\rangle = \frac{\Gamma\left(\sum_{k=1}^{r_i}\alpha_{ijk}\right)}{\prod_{k=1}^{r_i}\Gamma(\alpha_{ijk})} \times \prod_{k=1}^{r_i}\theta_{ijk}^{\alpha_{ijk}-1}$$

where $\Gamma()$ is the gamma function: $\Gamma(x+1)=x\Gamma(x)$ for positive real number x.

In our case study, all the event variables are binary (i.e. all r_i are 2). In this case, the Dirichlet distribution degenerates to a Beta distribution:*

$$D\langle\Theta_{ij}|\alpha_{ij1}, \alpha_{ij2}\rangle = \frac{\Gamma\left(\sum_{k=1}^{2}\alpha_{ijk}\right)}{\prod_{k=1}^{2}\Gamma(\alpha_{ijk})} \times \prod_{k=1}^{2}\theta_{ijk}^{\alpha_{ijk}-1}$$

When given a database D, we calculate the posterior distribution of θ_{ij1} and θ_{ij2} as follows. Let α'_{ijk} be the number of the cases of D in which X_i=X_{ik} and $\pi_i=\pi_{ij}$. Then the poste-

*The following discussions will mainly be based on the Beta distributions. The main conclusions and insights, however, should apply to the general Dirichlet distributions.

variable whose value is missing.

Two association rules are given below. They fill in a missing value of *TR* (test result) if metastasis (*M*) or local recurrence (*R*) are detected. They are used in the data processing for our case study.

If *M* = Y, **then** *TR* = +
If *R* = Y, **then** *TR* = +

It is well-accepted that rule elicitation is generally hard, especially when the problem domain is complicated and domain experts can not articulate their knowledge easily (it is usually the case). But fortunately, we elicit the association rules along the way of decision model building, and the model building itself provides a sound basis for rule elicitation.

4•2 Implicit Values

When modeling a dynamic decision problem, we may make or introduce a number of new variables whose values may not be directly available in the database at hand. In our case study, for instance, we introduced two binary abstract variables, i.e. *EOR* and *EOM*, in influence views depicted in Figures 3, 4 and 5. *EOR* is a variable representing whether evidence of local recurrence is present (e.g. change of bowel habit, and loss of appetite), while *EOM* is a variable representing evidence related to metastasis (e.g. back pain, loss of appetite, and lymph node involvement). In other words, *EOR* (*EOM*) is Y if the patient has symptoms/signs which may suggest local recurrence (metastasis). Therefore, the values of *EOR* and *EOM* are implicit in our database.

Again, deducing the values for the new variables is based a set of association rules extracted from the expert. Two rules for *EOM* are shown below:

If *BACK-PAIN* = Y, **then** *EOM* = Y
If *LYMPH-NODES* = INVOLVED, **then** *EOM* = Y

Furthermore, actions in DynaMoL influence views can be simple and compound (*e.g.* M&T). In the latter case, we may not directly have their instances in the database, and therefore have to develop domain-specific inquiries which extract the action instances from the data. For example, to find a instance of M&T, we have to select from the original database those cases where tests are ordered for detecting metastasis and if metastasis is detected, proper treatment is planned for patients.

4•3 Sparse Data

Another common problem of raw data is *data spareness*. Data sparseness refers to data-poor areas in a database. In such areas, learning performance often degenerates.

One way out of the data-spareness problem is simply to ignore it, and thus the modeler does not qualitatively model those parts of the problem where quantitative modeling may rely on data-poor areas. This strategy may not always work, because the precise locations of data-poor areas are not easy to determine, or the database may have so many data-poor areas that a trivial qualitative model of the problem may result.

Another way, which we usually take, is through *data abstraction*. When a data-poor area is detected, we group it with a data-rich area or another data-poor area together so that it becomes data-rich.

Data abstraction goes hand in hand with *model abstraction*. Consider Figure 3 again. The influence view shown in Figure 3 is an abstraction of the two influence views for two more specific actions, namely M&AT and M&ST. AT and ST mean adjuvant and surgical treatments, respectively; they are specializations of M&T. These two actions have an identical influence structure, as depicited in Figure 6. Notice that although "M&ST" and "M&AT" have the same influthe possible event variables that affect the state transitions, and the links the probabilistic dependences. For example, *EOM* probabilistically influences *TR*, which in turn influences *M*.

For an event variable *X* in an influence view, those who directly influence *X* are called parents of *X*. For example, in Figure 3, *EOR*, *EOM*, and $\mathbf{state}_i$ are the three parents of *TR*.

3•3 The Learning Problem

We adopt a Bayesian approach to assessing conditional probabilities for event variables in the influence views. We use Dirichlet or Beta prior distributions as the prior distributions.

Our learning problem consists of two tasks. First, we have to elicit prior information from the domain expert. This prior information is needed for assessing Dirichlet or Beta prior distributions of events in influence views. In dynamic decision making, the expert typically has different prior information at different decision stages, which makes prior information elicitation more time-consuming. However, we note that when events (*e.g.* cancer status) evolve continually over time, the burden of prior information elicitation could be much reduced, since the "continuity" len-provides a time scale for the expert to assess the prior information at different decision stages. For more details about prior information elicitation, see Section 5.

Second, we have to learn the posterior probabilities for event variables in influence views. The posterior probabilities of an event variable *X* are conditional on a action chosen, a decision stage at which the action is taken, and the parent event variables of *X*, as exemplified below:

$$\Pr(TR|EOR, EOM, STATE, ACTION, STAGE, DATASET).$$

4 DATA PREPROCESSING

There are a few tricky issues regarding raw data that must be handled before conditional probabilities can be learned. In dynamic decision modeling, we summarize a number of key issues in the following subsections, and discuss feasible methods for handling them.

4•1 Missing Values

a common phenomenon in a real-life database is that values of some variables may be missing in a database. Two general "missing mechanisms" can be distinguished. When there are no rules to account for the missingness, we say that the missing values are *missing at random*; otherwise, we say that they are *missing systematically.*

In the machine learning community, several techniques have been developed to learning probabilistic networks from incomplete data [6,4,13,3, 11]. These methods assume that the unknown values are missing at random [11]. This assumption is not reasonable in real-life domains. In decision-making domains, unknown values are missing usually due to decisions made. For example, when the doctor decides not to prescribe a test for a patient, the result of that test is certainly missing in the database. But it is not missing at random, but due to the doctor's decision.

In this paper, we adopt a rule-based approach to handling the missing-value issue. The approach fills in the missing values by deducing them based on a set of *association rules*. An association rule consists of two parts:

- **If**: Specifying the condition where it is applicable. The condition is represented by a conjunction of simple terms of the form (*variable* **rel** *value*), where *value* is either a value interval (*e.g.* (0.5, 3.5]), a set of values (*e.g.* {back-pain, bone-pain}), or a single value (*e.g.* yes), and **rel** is one of the two set relators **in** and **notin**, or the six numeric relators **=**, >=, >, <=, < and **!=**.
- **Then**: Recommending a value for a

semi-Markov decision process, provides a concise formulation of the decision problem; it also admits various solution methods. The translation convention supports automatic transformations among the different graphical representations.

3•1 Transition View

The transition view corresponds directly to the Markov state transition diagram. Given an action, the transition view depicts the possible state transitions. Figure 2 shows a transition view for an action or strategy "M&T", which denotes performing a diagnostic test to detect metastasis, and treating accordingly if cancer is detected. In the figure, the nodes denote the states, and the arcs the possible transitions given the action. The possible transitions at any decision stage are governed by a set of transition probabilities.

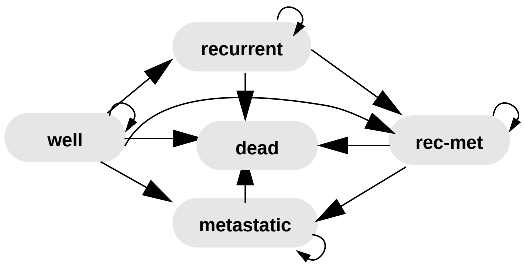

Figure 2: Transition view for action "M&T".

To solve for the optimal course of action, we have to assess transition probabilities for the alternative actions. It is, however, usually very difficult to assess such numbers directly. Therefore, the effects of the action are elaborated in the Influence View to facilitate reasoning and assessment of the probabilities.

3•2 Influence View

Given an action, the influence view shows the possible event variables that affect transitions from one state to another. In other words, an influence view is a refinement of a transition view. The event variables correspond to the chance nodes in an influence diagram; the influence view, therefore, is also analogous to a slice of a dynamic influence diagram[13], including all the chance nodes relevant to a specific decision stage.

Figures 3, 4, and 5depict the influence views for the three actions "M&T", "R&T", and "RM&T" in our example, respectively. Table 1 specifies all the event variables involved in the influence views.

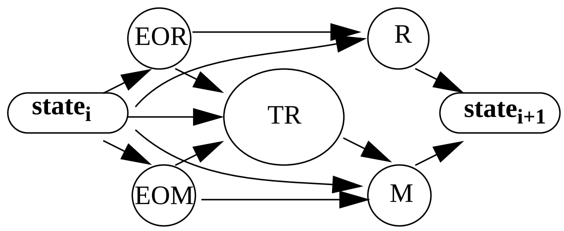

Figure 3: Influence view for action "M&T".

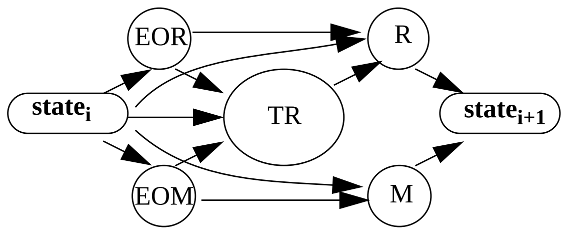

Figure 4: Influence view for action "R&T"

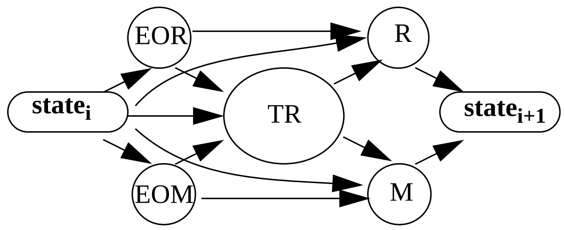

Figure 5: Influence view for action "RM&T"

Table 1: Event Variables Specification

Variable	Possible values	Meaning
EOR	Y, N	evidence of recurrence
EOM	Y, N	evidence of metastasis
TR	+, -	test results
R	Y, N	Is recurrence detected?
M	Y, N	Is metastasis detected?

Figure 3 shows the influence structure for the action "M&T". The circular nodes represent objective probabilities may not be easily calculated to support decision modeling; the recording formats, the measurement assumptions, and the processing errors associated with the data may complicate such derivations.

In this work, we examine the critical issues in automated learning of probabilistic parameters from large medical database. Our discussions are based on the DynaMoL framework and a case study in the follow-up of patients who have undergone colorectal cancer surgery. We present a Bayesian method for learning conditional probabilities from data for influence views, a key decision model in DynaMoL, analyze how to elicit prior probabilities from the domain expert, and discuss several important issues on preparing and processing raw data for application in dynamic decision modeling.

2 FOLLOWING UP PATIENTS WITH COLORECTAL CANCERS

To facilitate exposition, we consider a dynamic decision problem of detecting an optimal course of action in a follow-up program of colorectal cancer patients. After colorectal cancer surgery, patients are followed up so that any recurrence and/or metastases can be detected early. Based on the Dukes' classification system, colorectal cancers have four possible stages: A, B, C, or D. Stages A and B colorectal cancers are localized and have a good prognosis after curative resection. Stage C cancers are regional and are at a high risk of recurrence and metastasis. Stage D cancers are late-stage or metastatic cancers and have a poor prognosis even after surgery. In our case study, we focus on Stage C cancers.

At any decision stage, a Stage C patient's health state can be classified into three general states: **dead**, **well**, or **cancerous**; and **cancerous** can be further refined into three subclasses: **recurrent**, **metastatic**, **rec-met** (a special term for both recurrent and metastatic). The classification is shown in Figure 1.

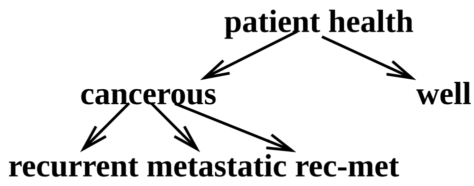

Figure 1: Patient's health state hierarchy

At each decision stage of a follow-up program for a patient, a doctor has three alternatives: 1) decide to detect if the patient has metastatic cancers; if detected, prescribe treatment (i.e. M&T); 2) decide whether to detect if the patient has recurrent cancers; if detected, prescribe treatment (i.e. R&T); 3) decide to detect if the patient has both metastatic and recurrent cancers; if detected, prescribe treatment (i.e. RM&T).

In managing the follow-up of colorectal cancer patients, a series of diagnostic tests are performed to detect possible recurrence, metastasis, or both recurrence and metastasis of the cancer; treatment is prescribed if cancer is detected. The decision is to determined the optimal course of diagnostic tests, over a sequence of decision stages, that would lead to the most cost-effective treatment outcomes.

3 THE DynaMoL FRAMEWORK

The DynaMoL framework has four major components: a dynamic decision grammar, a graphical presentation convention, a formal mathematical representation, and a translation convention. The decision grammar supports problem formulation with multiple interfaces. The presentation convention, in the tradition of graphical decision models, governs parameter visualization and specification in multiple perspectives; two graphical representations are currently included: Transition View and Influence View. The mathematical representation, in terms of a

Learning Conditional Probabilities for Dynamic Influence Views

Cungen Cao and Tze-Yun Leong
Medical Computing laboratory
Department of Information Systems and Computer Science
National University of Singapore, Lower Kent Ridge Road, Singapore 119260
{caocg,leongty}@iscs.nus.sg

Adrian Pheng Kheong Leong and Francis Seow Choen
Department of Colorectal Surgery
Singapore General Hospital, Outram Road, Block 6, Singapore 169608

Abstract

Dynamic decision making concerns problems in which both time and uncertainty are explicitly considered. A major challenge in applying decision analysis in dynamic decision problems is to elicit, estimate, and specify the numerous time-dependent conditional probabilities in the models. Based on the DynaMoL (a Dynamic decision Modeling Language) framework, we examine the critical issues in automated learning of numerical parameters from large medical databases. In this paper, we present a Bayesian method for learning conditional probabilities from data for influence views, a key decision-modeling facility in DynaMoL, analyze how to elicit prior probabilities from the domain expert, and discuss several important issues on processing and preparing raw data for application in dynamic decision modeling

1 INTRODUCTION

Many challenging decision problems in medicine and pharmacology involve explicit consideration of time and uncertainty. For instance, in patients who have undergone colorectal cancer surgery, an important decision problem is to determine the optimal follow-up schedule over 5 to 10 years. Similarly, the cost-effectiveness of a new drug for AIDS can be evaluated by comparing alternative treatment efficacies over the course of disease progression, in terms of estimated CD4 cells counts.

In recent years, decision analysis techniques are increasingly being applied to model and analyze dynamic decision problems in medicine [1,14,9,10]. Dynamic decision analysis or modeling frameworks are based on structural and semantical extensions of conventional decision models, e.g., decision trees and influence diagrams, with the mathematical definitions of finite-state stochastic processes. Recently, [7] has identified semi-Markov decision processes (SMDPs) as the common theoretical basis of existing dynamic decision modeling formalisms. A new framework called DynaMoL (for a Dynamic decision Modeling Language) has subsequently been proposed to integrate the graphical capabilities of the existing frameworks, and the concise properties and varied solutions of the mathematical formulations [8].

The mathematical formulations render the requirements for building complete or well-formed dynamic decision models more explicit. Nevertheless, assessing the relevant probabilistic parameters remains a very challenging task. Subjective assessments from domain experts may be adequate in some cases. When the decision situations are complex or the decision dimensions are large, however, the practicality of the modeling approach is limited by the lack of realistic estimations. On the other hand, given a large set of data,