Relational Decision Networks

William H. Hsu
Dept. of Computing & Information Sciences
Kansas State University
bhsu@cis.ksu.edu

Roby Joehanes
Dept. of Computing & Information Sciences
Kansas State University
robbyjo@cis.ksu.edu

Abstract

Decision-theoretic intelligent agents must function under uncertainty and be able to reason and learn about objects and relations in the context of action and utility. This paper presents a new relational graphical model (RGM), analogous to the probabilistic relational model (PRM), for representation of decisions under uncertainty. It first analyzes some basic properties of the representation and gives an adaptation of several decision network inference algorithms to these relational decision networks. It then describes some early experimentation with algorithms learning link structure in PRMs, discussing how these can be adapted to learning in decision networks. Finally, it considers the problem of representing dynamic relations in decision networks and sketches an extension of the dynamic PRM representation to include choice and utility.

1 INTRODUCTION

Uncertainty is a common feature of decision problems for which the *decision network* or *influence diagram* is currently one of the most widely-used graphical models. Decision networks represent the state of the world as a set of variables, and model probabilistic dependencies, action, and utility. Though they provide a synthesis of probability and utility theory, decision networks are still unable to compactly represent many real-world domains, a limitation shared by other propositional graphical models such as flat Bayesian ntworks and dynamic Bayesian networks. Decision domains can contain multiple objects and classes of objects, as well as multiple kinds of relations among them. Meanwhile, objects, relations, choices, and valuations can change over time. For example, a supply chain consists of multiple raw materials, components of manufactured goods, multiple goods, and a market that may consist of multiple redistributors and buyers.

Capturing such a domain in a decision network would require not only an exhaustive representing of all possible objects and relations among them, but also a combinatorially fast-growing space of choices and valuations. [JT99] This raises two problems. The first one is that the inference using such a dynamic decision network would likely exhibit near-pathological complexity, making the computational cost prohibitive. The second is that reducing the rich structure of domains such as supply-chain management and enterprise resource planning (ERP) to very large, "flat" decision network would make it much more difficult for human beings to comprehend. This paper addresses these two problems by introducing an extension of decision networks that captures the relational structure of some decision domains, and by adapting methods for efficient inference in this representation.

First-order formalisms that can represent objects and relations, as opposed to just variables have a long history in AI. Recently, significant progress has been made in combining them with a principled treatment of uncertainty. In particular, probablistic relational models, or PRMs, are an extension of Bayesian networks that allows reasoning with classes, objects, and relations. [FGKP99] The representation we introduce in this paper extends PRMs to decision problems in the same way that the decision networks extend Bayesian networks. We therefore call it the *relational decision network* or RDN. We develop two inference procedures for RDNs: the first based upon the traditional variable elimination algorithm developed by Shenoy [She92] and Cowell [Cow94], the second a more efficient one based upon an adaptive importance sampling-based algorithm [CD00, HGJ$^+$03].

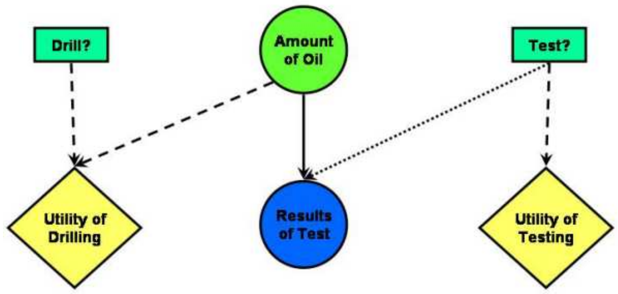

Figure 1: Decision network for the Oil Wildcatter problem. [She92]

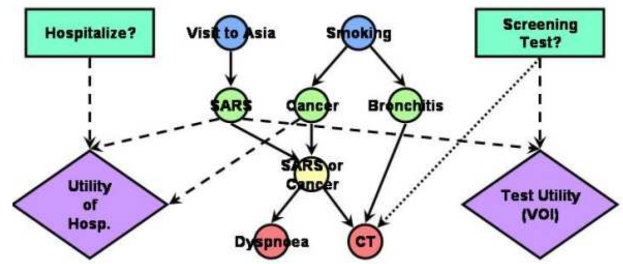

Figure 3: The DEC-Asia decision network.

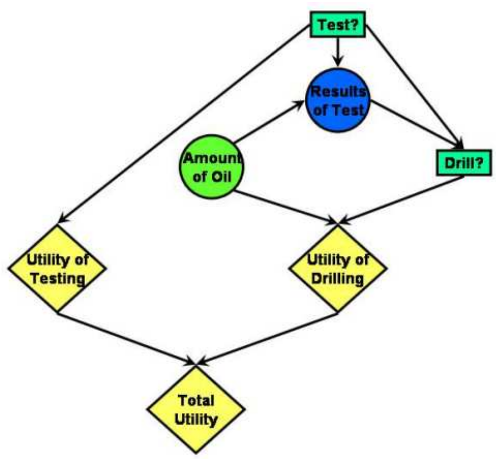

Figure 2: Influence diagram for the Oil Wildcatter problem. [She92]

2 DECISION NETWORKS AND INFLUENCE DIAGRAMS

Decision networks, introduced by Howard and Matheson [HM81], contain three kinds of nodes: *chance* (or *uncertainty*) nodes representing random variables as in a Bayesian network; *choice* (or *decision*) nodes) representing decisions to be made; and one or more *utility* nodes, each denoting a random variable ranging over utilities of the outcomes (as an aggregate of cost and benefit). These are depicted using circles, rectangles, and diamonds, respectively. [RN03, Nea04]

In a decision network, the edges are interpreted as follows [Nea04]:

1. Edges into chance nodes denote conditioning: The values of chance nodes are probabilistically dependent on the values of parents.
2. Edges into choice nodes denote sequence: The values of parents are known at the time the decision is made.
3. Edges into utility nodes denote deterministic functions of the values of parents.

Figure 1 depicts the decision network for the *Oil Wildcatter* problem. [She92, CDLS99] This graphical model represents the joint distribution and utility of a decision problem with boolean choice nodes and ternary chance nodes (Amount $\in$ Dry, Wet, Soaking; Result $\in$ Poor, Intermediate, Good).

Figure 2 shows an equivalent influence diagram, with links between choice (decision) nodes. These graphical models specify a decision sequence: (Test, Result, Drill, Oil), corresponding to a decision tree of orderings over choice, chance, and outcome utility nodes.

3 PROBABILISTIC RELATIONAL MODELS

Probabilistic relational models (PRMs) extend the flat (propositional) representation of the variables and the conditional dependencies among them to an object-relational representation. Before proceeding to discussion the decision network analogues of PRMs, we briefly review the PRM family and the relevant components of a PRM specification.

A *relational schema* is a set of classes $\mathcal{C} = C_1, C_2, \ldots, C_k$, where every class C is associated with a ground set of *propositional attributes* $\mathcal{A}(C)$ and a set of *relational attributes*, also known as its *reference slots*,

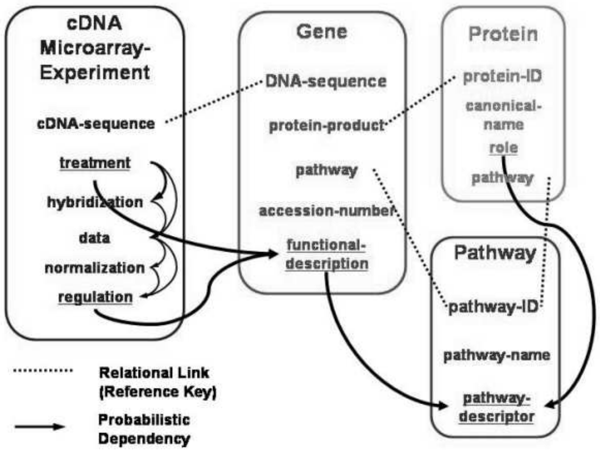

Figure 4: Probabilistic relational model.

$\mathcal{R}(C)$. Propositional attributes A are instance variables $C.A$ ranging over finite domains $V(C.A)$. Similarly, relational attributes $C.R$ each correspond to a set of member objects of a class and therefore range over the inclusion-exclusion (power set) $2^{C'}$ of some class C' from $\mathcal{C}$. That is, $C.R$ represents the set of members of the concept class C', so it is a selector function between C and $2^{C'}$. [SDW03]

As an example extending the DEC-Asia decision network above, the Patient schema might be used in to represent partial or total patient records, with classes corresponding to information about a patient's pulmonary medical history, smoking history, travel itinerary, and groups of people contacted. The propositional attributes of the medical history include the patient's age, previous contagious pulmonary diseases contracted, and currently extant diseases; the relational attributes might include the patient's membership in a list of quarantine subjects and *links* between patients denoting specific exposure incidents and contexts. Note that some of these are static and some, such as clusters of at-risk destinations and groups of people, may be dynamic relational attributes.

An *instantiation* of a schema is a set of objects, each belonging to some class $C \in \mathcal{C}$, where every object is specified using lists of its propositional and relational attributes and their nesting according to the (possibly hierarchical) schema. For example, an instantiation of the Patient schema might be a particular patient case history, with all available descriptive data and test results specified.

A *probabilistic relational model (PRM)* encodes a probability distribution over the set of possible instantiations I of some schema [FGKP99]. The *object skeleton* of an instantiation is the set of its member objects, abstracted over propositional attributes (i.e., with their values unspecified). The *relational skeleton* of an instantiation is the set of its member objects with all relational attributes specified for a **given** set of propositional attribute values. In the case of "known structure", the relational skeleton is provided as input and the PRM specifies a full conditional distribution for every attribute A of every class C using this skeleton. The underlying ground network consists of arcs into an attribute in C and parent sets consisting of attributes of C and other classes. The eligible classes that may contain parents of a node in C are those related by some *slot chain*, i.e., a composition of relational attributes. In order to implement this aggregation, a function mapping multiple attribute values into one is required. Aggregation functions could include descriptive statistics such as modes, median, and moments (mean and variance) of relational attributes. In the example from the previous section, the total or per capita average of diagnosed severe acute respiratory syndrome (SARS) cases in countries visited by a person is an example of an aggregate.

As a further example to illustrate slot chains, Figure 4 depicts a PRM for the domain of computational genomics, particular gene expression modeling from DNA hybridization microarrays. Slot chains can be traced using the reference keys (dotted lines). This PRM contains tables for individual microarrays or gene chips (admitting aggregation of chip objects into classes), putative gene function (where known or hypothesized), putative pathway membership (where known or hypothesized), and protein production (a target aspect of discovered function).

A PRM Π for a relational schema S is defined as follows. For every class C and every propositional attribute $A \in \mathcal{A}(C)$, we have:

1. A set of *parents* $Pa(C.A) = \{Pa_1, Pa_2, \ldots, Pa_i\}$, where each Pa_i has the form $C.B$ or $\gamma(C.\tau.B)$, where τ is a slot chain and $\gamma()$ is an aggregation function.

2. A *conditional probability function or table* for $P(C.A|Pa(C.A))$.

Let $\mathcal{O}$ be the set of objects in the relational skeleton of Π. The probability distribution over the instantiations I of S, over which the Π abstracts, is:

$$P(I) = \Pi_{obj \in \mathcal{O}} \Pi_{A \in \mathcal{A}(obj)} P(obj.A|Pa(obj.A))$$

This allows a PRM to be flattened into a large Bayesian network containing ground (propositional) chance nodes, with one variable for every attribute of every object in the relational skeleton of Π and belief functions (usually deterministic) for the aggregation operations. The latter are open-ended in form and omitted from the formula for brevity.

As Getoor *et al.* [GFKT02] and Sanghai *et al.* [SDW03] note, the most general case currently amenable to learning is where an object skeleton is provided and structure and parameter learning problems must be solved in order to specify a distribution over relational attributes. In the DEC-Asia domain, a PRM might specify a distribution over possible transmission vectors of a SARS-infected person (the itinerary, the locale of contamination, the set of persons contacted).

4 RELATIONAL DECISION NETWORKS

4.1 RDN Representation

We now extend decision networks to relational representations using a simple and straightforward synthesis of decision network and PRM specifications.

Definition 1 The relational decision schema S for a decision network B consists of three sets of classes $\mathcal{C}_{\mathcal{X}} = C_{X_1}, C_{X_2}, \ldots, C_{X_n}$, $\mathcal{C}_{\mathcal{D}} = C_{D_1}, C_{D_2}, \ldots, C_{D_m}$, and $\mathcal{C}_{\mathcal{O}} = C_{O_1}, C_{O_2}, \ldots, C_{O_l}$, where every class C_X is associated with a ground set of *propositional attributes* $\mathcal{A}(C_X)$ and a set of *relational attributes* $\mathcal{R}(C_X)$. Propositional decision attributes A are instance variables $C.A$ ranging over the finite chance, decision, and utility domains $V_{\mathrm{X}(C.A)}$, $V_{\mathrm{D}(C.A)}$, $V_{\mathrm{O}(C.A)}$. Relational attributes $C.R$ each correspond to a set of member objects of a class and, for all chance, decision, and outcome (utility) nodes, respectively, range over the power sets:

1. **Objects**: $2^{C_{X'}}$ of some class $C_{X'}$ from $\mathcal{C}_{\mathcal{X}}$
2. **Actions**: $2^{C_{D'}}$ of some class $C_{D'}$ from $\mathcal{C}_{\mathcal{D}}$
3. **Outcomes**: $2^{C_{O'}}$ of some class $C_{O'}$ from $\mathcal{C}_{\mathcal{O}}$

Thus the relational attributes $C.R$ can include distinguished member *action identifiers* and *outcome identifiers* specifying a representation for equivalence classes of decisions and outcomes. Note that the range of actions may be continuous (e.g., in intelligent control or continuous decision problems) and the range of *utilities* may also be continuous. C_D and C_O specify only membership in S.

Definition 2: A *relational decision network (RDN)* for a relational schema S is a PRM $M = (B, D, U)$ with distinguished decision and choice nodes, factored according to S, where B encodes all conditional probability functions among chance nodes and their parents (choice and chance nodes), D encodes all sequential decision functions, and U encodes all conditional outcome and utility functions.

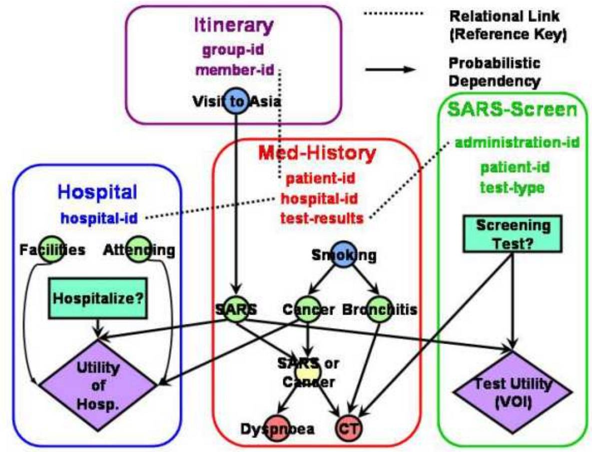

Figure 5: Relational decision network for the DEC-Asia domain.

When the decision network's object skeleton is not known (i.e., the set of decisions and outcomes is not fully pre-specified), the RDN includes a boolean *existence variable* for propositional attributes of C_X, C_D, and C_O, and a boolean *reference slot variable* for relational attributes of C_X, C_D, and C_O.

Figure 5 shows a DRN for the DEC-Asia domain.

4.2 RDN Inference: Sampling-Based Algorithms

Our DBN algorithms include two sampling algorithms: Likelihood Weighting and Adaptive Importance Sampling (AIS). [Guo03] For brevity, we refer interested readers to Cheng and Druzdzel (2000) [CD00] for detailed descriptions of these algorithms.

A desired joint probability distribution function P(X) can be computed using the chain rule for Bayesian networks, given above. The most probable explanation (MPE) is a truth assignment, or more generally, value assignment, to a query Q = X E with maximal posterior probability given evidence e. Finding the MPE directly using Equation (1) requires enumeration of exponentially many explanations. Instead, a family of exact inference algorithms known as clique-tree propagation (also called join tree or junction tree propagation) is typically used in probabilistic reasoning applications. The first MPE (*belief revision*) algorithm for DAG models was developed by Pearl [Pea88]. The first general maximum *a posteriori* (MAP or *belief updat-*

ing) algorithms were developed by Pearl [Pea88] and independently by Shachter [Sha86]. The most popular extant implementation of belief updating is the junction tree algorithm of Lauritzen and Spiegelhalter [CDLS99]. Although exact inference is important in that it provides the only completely accurate baseline for the fitness function f, the problem for general BNs is $\#\mathcal{P}$-complete (thus, deciding whether a particular truth instantiation is the MPE is NP-complete) [Co90, Wi02]. [SJJ96] Approximate inference refers to approximation of the posterior probabilities given evidence. One stochastic approximation method called importance sampling [CD00] estimates the evidence marginal by sampling query node instantiations.

5 DYNAMIC EXTENSIONS

5.1 Dynamic Bayesian Networks

Dynamic Bayesian networks (DBNs) extend flat Bayesian networks to model problems with a temporal component. In a decision network, the state at time t is represented using a set of random variables $X_t = \{X_{1,t}, X_{2,t}, \ldots, X_{n,t}\}$ The state at time t is dependent on thoes at previous time steps. Each state in the system is assumed to depend only on the immediately preceding state (i.e., the system observes the Markov property in the first degree). The representation must therefore capture the transition distribution $P(X_{t+1}|X_t)$. This can be done using a two-time-slice Bayesian network fragment (2TBN) B_{t+1}, containing variables from X_{t+1} whose parents are variables from X_t or X_{t+1}, and variables from X_t without any parents. The process is also usually assumed to be stationary, so that the transition equations for all time slices are identical: $B_1 = B_2 = \ldots = B_t = B_{\rightarrow}$. Thus a DBN network is fully specified using a pair of Bayesian networks $(B_0, B_{\rightarrow}$, where B_0 represents the initial distribution $P(X_0)$ and $B_{\rightarrow}$ is a two-time slice Bayesian network, which as discussed above defines the transition distribution $P(X_{t+1}|X_t)$.

5.2 Dynamic Probabilistic Relational Models

Sanghai *et al.* recently developed a hybrid of PRMs and DBNs for decision-theoretic troubleshooting, which are called *dynamic probabilistic relational models (DPRMs)*. The key innovation of this related work is that it is the first representation to support relational aggregates in a temporal model. This is achieved by extending the 2TBN representation to a 2TPRM where each time slice contains a PRM. This extension is straightforward, with the slight complication that in flattening (or "unrolling") a PRM into a ground representation.

DRNs admit straightforward extension to dynamic domains using the 2 time slice dynamic PRM representation presented by Sanghai et al..

6 CONCLUSIONS AND FUTURE WORK

We have described a new representation for decision networks that combines the compact abstraction of PRMs with utility theoretic graphical model. We have considered several continuations of this research, grouped into four categories: applications, scalability, comparison to other decision models, and improvements to the ordering algorithm. DRNs are currently being implemented for use in Bayesian Network tools in Java (BNJ).

Acknowledgements

This work was partially supported by the National Science Foundation under a cooperative agreement.

References

[CD00] J. Cheng and M. J. Druzdzel. AIS-BN: An adaptive importance sampling algorithm for evidential reasoning in large Bayesian networks. *Journal of Artificial Intelligence Research*, 13:155–188, 2000.

[CDLS99] R. G. Cowell, A. P. Dawid, S. L. Lauritzen, and D. J. Spiegelhalter. *Probabilistic Networks and Expert Systems.* Springer, New York, 1999.

[Cow94] Robert G. Cowell. Decision networks: a new formulation for multistage decision problems. Technical Report 132, Department of Statstical Science, University College London, 1994.

[FGKP99] Nir Friedman, Lise Getoor, Daphne Koller, and Avi Pfeffer. Learning probabilistic relational models. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-1999)*, pages 1300–1309, 1999.

[GFKT02] Lise Getoor, Nir Friedman, Daphne Koller, and Ben Taskar. Learning probabilistic models of link structure. *Journal of Machine Learning Research*, 3(12):679–707, 2002.

[Guo03] H. Guo. *Algorithm Selection for Sorting and Probabilistic Inference: A Ma-*

chine Learning-based Approach. PhD thesis, Kansas State University, 2003.

[HGJ+03] William H. Hsu, Haipeng Guo, Roby Joehanes, Benjamin B. Perry, and Julie A. Thornton. Bayesian network tools in java (bnj) v2. Distributed from URL: http://bndev.sourceforge.net, 2003.

[HM81] R.A. Howard and J.E. Matheson. Influence diagrams. In R.A. Howard and J.E. Matheson, editors, *Readings on the Principles and Applications of Decision Analysis*, volume II, pages 721–762. Strategic Decisions Group, Menlo Park, CA, 1981.

[JT99] E. Jorgensen and N. Toft. A Bayesian network based monitoring system for sow management. Technical Report 80, 1999.

[Nea04] R. E. Neapolitan. *Learning Bayesian Networks.* Prentice Hall, New York, 2004.

[Pea88] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufmann, San Mateo, CA, 1988.

[RN03] Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach.* Prentice Hall, NJ, 2 edition, 2003.

[SDW03] Sumit Sanghai, Pedro Domingos, and Daniel Weld. Dynamic probabilistic relational models. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-2003)*, pages 302–313, 2003.

[Sha86] R. Shachter. Evaluating influence diagrams. *Operations Research*, 34:871–882, 1986.

[She92] P. Shenoy. Valuation-based systems for Bayesian decision analysis. *Operations Research*, 40(3):463–484, 1992.

[SJJ96] L. K. Saul, T. Jaakkola, and M. Jordan. Mean field theory for sigmoid belief networks. *Journal of Artificial Intelligence Research*, 4:61–76, 1996.