Representation Requirements for Supporting Knowledge-Based Construction of Decision Models in Medicine

Tze-Yun Leong
MIT Laboratory for Computer Science
545 Technology Square, room 420
Cambridge, MA 02139
(leong@lcs.mit.edu)

This paper analyzes the medical knowledge required for formulating decision models in the domain of pulmonary infectious diseases (PIDs) with acquired immunodeficiency syndrome (AIDS). Aiming to support dynamic decision-modeling, the knowledge characterization focuses on the ontology of the clinical decision problem. A relevant set of inference patterns and knowledge types are identified.

Keywords: *knowledge-based systems, medical decision science, representation requirements.*

1 Introduction

In recent years, decision analysis has gradually been recognized as a powerful technique for selecting the optimal strategies in difficult clinical decision problems. Research in knowledge-based decision systems (KBDS) attempts to automate the decision analysis process using artificial intelligence techniques. In the dynamic decision- modeling approach, the decision models for different problems are dynamically constructed from a knowledge base [8]. To date, however, while much progress has been made in improving the algorithms for manipulating decision models, the automated model construction process remains to be formalized.

This paper characterizes the knowledge for supporting dynamic decision-modeling in medicine. Unlike previous efforts, instead of concentrating on the structural components of the *model* such as nodes, conditional probabilities, and influences, we focus on the ontological features of the decision *problem* such as contexts, classes of observed events, classes of available actions, classes of possible outcomes, temporal precedence, and probabilistic and contextual dependencies. Although the results reported here are yet to be tested in an implementation, this analysis exercise serves as a step toward realizing a uniform representation framework for supporting dynamic decision- modeling in clinical KBDS.

2 A Knowledge-Based Decision System

Figure 1 depicts the general system architecture on which the following analysis is based. The proposed KBDS consists of a *decision-maker* or *planner* which constructs a *decision model* by accessing the information in the *knowledge base*. The decision models in question are *qualitative probabilistic networks* (QPNs) [8]. A specific way of automatically generating QPNs is demonstrated in Wellman's SUDO PLANNER system [8]; our discussions, however, assume a general decision making process independent of any system or implementation.

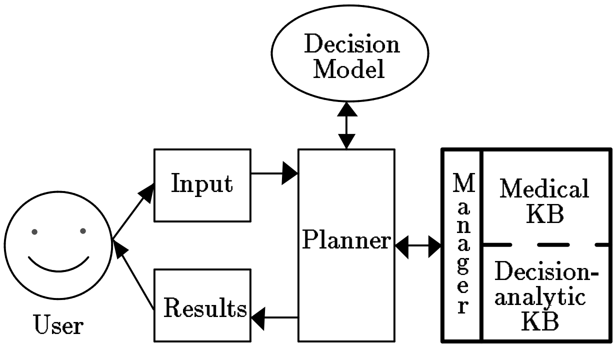

Figure 1: A Knowledge-Based Decision System.

3 An Example

An example decision problem in the domain of PIDs with suspected AIDS [7] is shown below:

The patient is a 29 year-old man with a history of intravenous (IV) drug abuse and a one-week history of low-grade fever, non-productive cough, and dyspnea. His chest X-ray (CXR) shows bilateral diffuse interstitial infiltrates. His arterial blood gas (ABG) shows mild hypoxemia on room air. The initial impression was pneumonia possibly due to opportunistic infection with suspected AIDS. The problem is to investigate whether or not to employ empiric therapy for pneumocystis carinii pneumonia (PCP), and how non-invasive diagnostic tests such as sputum examination and gallium scanning compare with invasive procedures such as bronchoscopic bronchoalveolar lavage (BAL) and bronchoscopic transbronchial biopsy (TBBx).

Given the above information, the ultimate goal for the KBDS is to construct a decision model as shown in Figure 2. The following section examines the corresponding

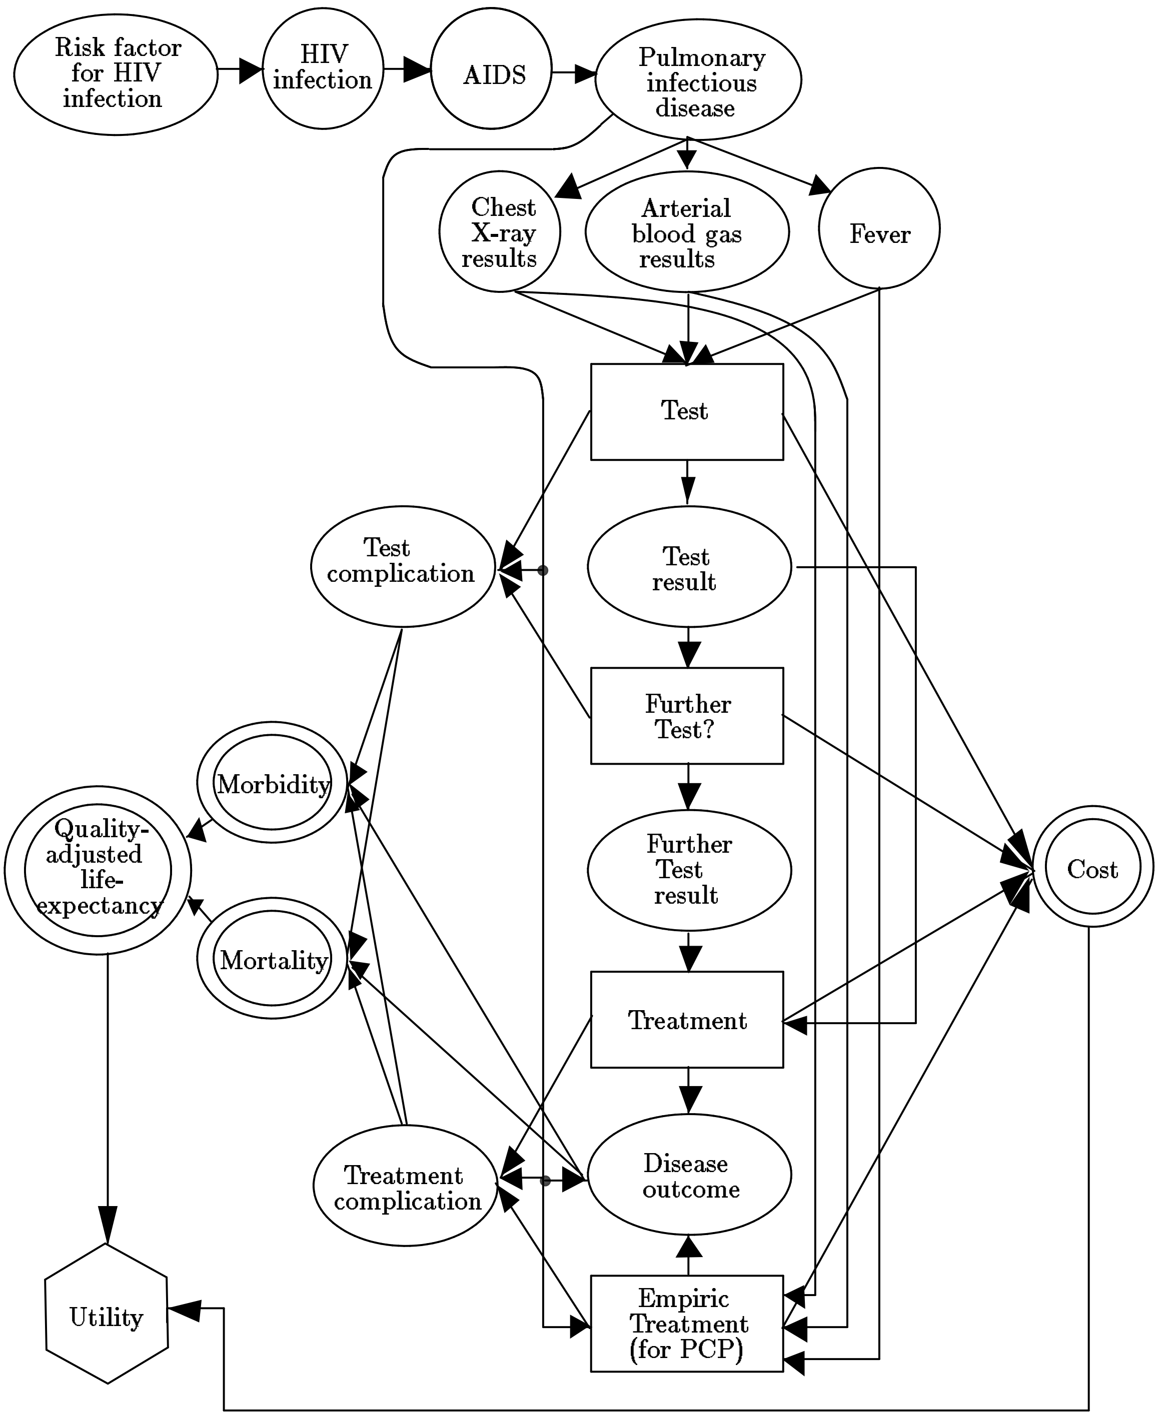

Figure 2: A QPN For The Example. The labels on the arcs are not shown.

decision making process. In each step of the process, the inferences involved and the representation support required are identified.

4 The Decision Making Process

The decision-analytic approach to decision making can be viewed as a five-step process:

4.1 Background Information Characterization

The process begins by differentiating the variables concerned, the actions available, and the possible outcomes involved in the input information. In the clinical setting, these events can be divided into six categories, as shown in Table 1 for our example.

Table 1: Characterized Background Information

Category	Concepts
General history	29 year old, male, IV drug abuse
Signs/Symptoms	Low-grade fever, dyspnea, non-productive cough
Laboratory findings	CXR: bilateral diffuse interstitial infiltrates ABG: hypoxemia on room air
Diseases	PIDs with suspected AIDS
Alternatives	Empiric therapy for PCP, sputum-examination, gallium scanning, BAL, and TBBx

Each event in Table 1 can be regarded as a *concept.* A concept is an *event* or a *random variable* in the probabilistic sense; it denotes an abstract description of an object, an attribute, a state of being or a process, depending on the circumstances. These input concepts can be categorized by asking questions like:

- Is low-grade fever a kind of sign or symptom?
- Is sputum examination a kind of alternative?

In general, the inference pattern for background characterization is captured in the query:

- (Q1) Does concept A relate to concept B in $< categorization >$?

where *categorization* is a grouping or classification of concepts by a partial-order, binary categorical relation, e.g., specialization (AKO), decomposition (PARTOF), etc.

To support such inferences, the concepts should be identifiable as belonging to different categories or classes. All concepts in a category share some specific characteristics in their descriptions.

The characterized background information, however, is still insufficient for formulating a decision model. For example, the different PIDs being considered are not explicitly stated. When necessary, the missing information must be derivable from the knowledge base.

4.2 Domain Context Establishment

The domain context serves as a basis for expectation and defines the scope of the applicable operations in a given problem. In the clinical setting, a context is usually indicated by a suspected disease, a syndrome, or a general diagnostic category, e.g., an acute respiratory disorder [3].

In our example, the clinical context is "PIDs with suspected AIDS." This context is established by simply identifying the suspected diseases in the input information. Therefore, the general inference pattern is captured in the query:

- (Q2) What are the concepts related to concept A in $< categorization >$?

where *categorization* is induced by a categorical relation as mentioned earlier.

As in characterizing background information, supporting domain context establishment again requires categorization or classification of the concepts involved.

Establishing the context should allow the planner to access the context-sensitive information. Hence, such context-sensitive knowledge must be expressible in the knowledge base.

4.3 Decision Problem Formulation

Guided by the characterized background information, a decision problem is formulated within the clinical context by identifying 1) the most relevant diseases/hypotheses involved, 2) the most relevant actions available, 3) the relative significance, possible outcomes, and complications of the concepts derived from 1) and 2), and their effects on each other, and 4) the evaluation criteria concerned.

Table 2 shows some of the relevant concepts in our example. These concepts are derivable by asking questions like:

- What are the most common pneumonias caused by AIDS-related PIs?
- What are the treatments of the pneumonias?

In general, the inference patterns are captured in the queries Q1, Q2, and:

- (Q3) What are the concepts related to concept A by $< interaction >$?

where Q1 and Q2 are defined as before, and Q3 is for inferring the interactions, i.e., the correlational/influential/causal relationships among the concepts. Some examples of interactions among the concepts are: "presence of HIV infection causes AIDS," "presence of treatment negatively-influences severity of disease," etc.

To support the above queries, the knowledge base must contain the relevant relationships among the con-

Table 2: Concepts Involved in Decision Problem

Risk-factor-of-HIV-infection:	**Disease-outcome:**	**Empiric-treatment-for-PCP:**
IV-drug-abuse	Cured	TMP-SMZ
HIV-infection	Improved	IV-pentamidine
AIDS	Not-improved	Aerosol-pentamidine
Pulmonary-infectious-disease:	Worsened	**Cost**
PCP	Death	**Morbidity**
Pul. toxoplasmosis	**Test:**	**Mortality**
Pul. TB	Sputum-examination	**Quality-adjusted-life-expectancy**
MAI-complex	Gallium-scanning	**Utility**
Pyogenic-bacterial-pneumonia	BAL	⋮
Legionellosis	TBBx	
Pul. cryptococcosis	⋮	
Other-PIDs		

cepts. In addition, the notion of varying degrees of significance for these relations in a particular context should be captured in the knowledge base. This would facilitate derivation of the most relevant information for the problem at hand.

4.4 Decision Model Construction

As mentioned, a decision model for our example is shown in Figure 2.

To construct such a decision model, its structure, e.g., nodes and links in a QPN, and preference model, e.g., evaluation criteria such as morbidity, mortality, and monetary costs associated with utilities, an must be inferrable from the knowledge base. Hence, the construction involves asking questions like:

- How are the observable effects of the alternatives relate to the chance events?
- What are the outcomes that affect the evaluation criteria?

The general inference patterns are captured in the queries Q3 and:

- (Q4) Does concept A relate to concept B by $< interaction >$?

where Q3 and the interactional relationships are as described earlier.

To support these queries, again the relevant interactions among the concepts must be expressible in the knowledge base. These interactions involve both domain concepts and decision-analytic concepts, e.g., "presence of disease positively-influences morbidity."

4.5 Decision Model Evaluation

Upon completion, the decision model is evaluated or solved by some procedure, e.g., graph reduction of an QPN, with respect to the evaluation criteria. The evaluation criteria assumed in our example are expected monetary cost and quality-adjusted life expectancy. Given a well-formed decision model, only procedural knowledge is needed in this step.

5 Summary of Representation Requirements

The above analysis shows that four types of general inference patterns, Q1-Q4, are involved in the automated decision analysis process. Three types of knowledge are required to support such inferences:

Categorical Knowledge: The categorical knowledge captures the definitional and structural relations among the clinical concepts. This type of knowledge should provide the system with the power of *abstraction* and *inheritance*.

Uncertain Knowledge: The uncertain knowledge captures the correlational, influential, or causal relations among the concepts. This type of knowledge should allow expression of the varying degrees of temporal and probabilistic dependency among the concepts.

A Contextual Notion: This contextual notion can be thought of as a focusing or conditioning mechanism in the probabilistic sense. It sets a boundary on the categorical and uncertain knowledge, enabling the planner to identify the relevant information in different situations. Differentiation of the relational significance among a set of concepts in a particular context should be expressible in the knowledge base. Moreover, the different contexts should be compositional and hierarchically definable.

6 Related Work

The major shortcomings of the static decision-modeling approach, i.e., treating pre-enumerated decision models or templates as knowledge bases, result from the rigidity of the knowledge bases. Such knowledge bases do not reflect the nature of the domain knowledge.

The different representations used in existing KBDS with the dynamic decision-modeling approach are not very satisfactory, either. The first order logic-like representations, such as those employed by Breese [1], and Goldman and Charniak [2], have no explicit hierarchical dimensions; limited contextual information are captured as conditional probabilities matrices in these frameworks.

Despite allowing explicit hierarchical and context-dependent domain descriptions, the representation framework in Wellman's SUDO-PLANNER [8] system has limited expressiveness; there is also no general mechanism for capturing contextual information in the whole framework.

Other relevant representation formalisms include those that incorporate an uncertainty model to a hierarchical representation framework. Some of these efforts attempt to accommodate the uncertainty models by re-interpreting the semantics of existing representations [5, 9], while others try to couple the two to form a coherent framework [6]. Again, however, none of these frameworks integrates context-sensitive categorical and uncertainty knowledge in a general way.

7 Discussion and Conclusion

By focusing on the ontology of a clinical decision problem in a complex domain, we have identified a set of inference patterns and knowledge types for supporting automated construction of decision models in medicine. The results show that to support dynamic decision-modeling, the structure of the knowledge base must reflect the nature of both the decision problem and the domain knowledge.

The brief survey on existing representations has shed some light on a design approach for integrating categorical and uncertain knowledge in a context-sensitive manner. We believe such an integration calls for a framework with 1) a terminological component for establishing the categories of structural concepts; 2) an assertional component for expressing the interactions among the concepts; and 3) a network interpretation for the concepts as the nodes and the relations as the arcs, thereby capturing the context notion by partitioning the network in different ways.

A more detailed exposition for such a design approach is described elsewhere [4]. Many interesting and complex research issues arise in the proposed representation design. Careful examination of these issues, we believe, will lead to the formalization of both the automated decision model construction process and the medical and decision-analytic knowledge involved.

Acknowledgments

I would like to thank Peter Szolovits and Mike Wellman for advice on the project, and Frank Sonnenberg for guidance in the medical domain and for providing the example case

This research was supported by the National Institutes of Health grant no. 5 R01 LM04493 from the National Library of Medicine.

References

[1] Jack Breese and Edison Tse. Integrating logical and probabilistic reasoning for decision making. In *Proceedings of the Workshop on Uncertainty in Artificial Intelligence*, pages 355–362, July 1987.

[2] Robert P. Goldman and Eugene Charniak. Dynamic construction of belief networks. In *Proceedings of the Sixth Conference on Uncertainty in Artificial Intelligence*, pages 90–97, 1990.

[3] Jerome P. Kassirer and Richard L. Kopelman. The critical role of context in the diagnostic process. *Hospital Practice*, pages 67–76, August 15, 1987.

[4] Tze-Yun Leong. Knowledge representation for supporting decision model formulation in medicine. TR 504, Massachusetts Institute of Technology, Laboratory for Computer Science, 545 Technology Square, Cambridge, MA, 02139, June 1991.

[5] Dekang Lin and Randy Goebel. Integrating probabilistic, taxonomic and causal knowledge in abductive diagnosis. In *Proceedings of the Sixth Conference on Uncertainty in Artificial Intelligence*, pages 40–45, 1990.

[6] Alessandro Saffiotti. A hybrid framework for representing uncertain knowledge. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, pages 653–658, Cambridge, Massachusetts, 1990. American Association for Artificial Intelligence, AAAI Press and The MIT Press.

[7] F. A. Sonnenberg. An intelligent decision system for lung diseases in AIDS. Grant proposal (No. 1R29LM04936-01A1) submitted to the National Library of Medicine, February 1989.

[8] Michael P. Wellman. *Formulation of Tradeoffs in Planning Under Uncertainty*. Pitman and Morgan Kaufmann, 1990.

[9] John Yen and Piero P. Bonissone. Extending term subsumption systems for uncertainty management. In *Proceedings of the Sixth Conference on Uncertainty in Artificial Intelligence*, pages 468–473, 1990.