

**LABORATORY FOR
COMPUTER SCIENCE**



**MASSACHUSETTS
INSTITUTE OF
TECHNOLOGY**

MIT/LCS/TR-504

**KNOWLEDGE REPRESENTATION
FOR SUPPORTING DECISION
MODEL FORMULATION
IN MEDICINE**

Tze-Yun Leong

June 1991

545 TECHNOLOGY SQUARE, CAMBRIDGE, MASSACHUSETTS 02139

This blank page was inserted to preserve pagination.

Knowledge Representation for Supporting Decision Model Formulation in Medicine

by

Tze-Yun Leong

May, 1991

© Massachusetts Institute of Technology, 1991

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

Knowledge Representation for Supporting Decision Model Formulation in Medicine

by
Tze-Yun Leong

This report is a modified version of a thesis submitted to the Department of Electrical Engineering and Computer Science on August 15, 1990, in partial fulfillment of the requirements for the degree of Master of Science

Abstract

Clinical decision making involves a large, complex, and ever-changing body of knowledge. Characterizing such knowledge illuminates the representational and computational requirements for automated clinical decision analysis. This work analyzes the medical knowledge required for formulating decision models in the domain of pulmonary infectious diseases (PIDs) with acquired immunodeficiency syndrome (AIDS). Based on the analysis, a knowledge representation framework is proposed. The framework is evaluated by showing how it supports decision model formulation for an example case.

Aiming to support constructive decision-modeling, the knowledge characterization focuses on the ontological features of the decision problem such as contexts, classes of observed events, classes of available actions, classes of possible outcomes, and probabilistic and contextual dependency. A relevant set of inference patterns and knowledge types are identified. These results are incorporated into a representation design that integrates categorical and uncertain knowledge in a context-sensitive manner.

Keywords: Knowledge representation, automated decision analysis, categorical and uncertain knowledge, context-sensitivity, qualitative probabilistic networks.

Thesis Supervisor: Peter Szolovits

Contents

1	Introduction	1
1.1	Background and Motivations	1
1.1.1	Knowledge-Based Decision Systems	1
1.1.2	Static versus Constructive Decision-Modeling	2
1.1.3	Supporting Clinical Decision-Modeling	3
1.2	Objectives	3
1.3	A Hypothetical System	4
1.4	Overview of Thesis	5
2	The Domain	7
2.1	Pulmonary Infectious Diseases with AIDS	7
2.2	An Example Case	9
3	Clinical Decision Analysis	11
3.1	Background Information Characterization	11
3.2	Clinical Context Establishment	13
3.3	Decision Problem Formulation	13
3.4	Decision Model Construction	16
3.5	Decision Model Evaluation	16
3.6	Representation Requirements	16
3.6.1	Categorical Knowledge	16
3.6.2	Uncertain Knowledge	18
3.6.3	A Contextual Notion	18

7.4	Contextual Effects on Concept Categorizations	53
8	The Medical Knowledge Base	57
8.1	The Contents	57
8.1.1	Physiological-states	57
8.1.2	Infections	58
8.1.3	Diseases	58
8.1.4	Tests	59
8.1.5	Treatments	59
8.2	The Structure	59
8.3	Constructing The Medical Knowledge Base	59
9	Supporting Decision Model Formulation	63
9.1	General Queries Format	63
9.2	Supporting Background Characterization	65
9.3	Supporting Clinical Context Establishment	66
9.4	Supporting Decision Problem Formulation	66
9.5	Supporting Decision Model Construction	72
9.6	Supporting Decision Model Evaluation	75
9.7	Formulating Other Decision Models	76
10	Conclusions	77
10.1	Achievements	77
10.2	Limitations	78
10.3	Future Directions	79
10.3.1	Temporal Representation	79
10.3.2	Changes of Representation in an Evolving KB	80
10.3.3	Integrating Decision Analytic Knowledge	81
10.4	Summary	81
A	The Medical Knowledge Base	83

List of Tables

2.1	Differential Diagnosis of AIDS-related Pulmonary Infectious Diseases	8
2.2	Diagnostic Tests for Pulmonary Infections	8
2.3	Treatments for AIDS-related Pulmonary Infections	9
3.1	Characterized Background Information	12
3.2	Concepts Involved in Decision Problem	15
5.1	Indirect Effects of Interactions.	38
7.1	Synergistic effects on interactions.	51

List of Figures

1-1	A knowledge-based clinical decision system.	5
3-1	A QPN for the Example Case.	17
4-1	The Meaning Triangle	25
4-2	The Definability Spectrum	26
5-1	A simplified example of a concept: PC-infection.	30
5-2	A simplified attributive concept example: Severity of disease	32
5-3	Part of the interaction-model of PC-infection	37
7-1	Complication of disease	49
7-2	Enabling effect example.	50
7-3	Blocking effect example.	52
7-4	An example of one-to-many contextual effects.	53
8-1	A concept as a context.	60
9-1	Fragment of the medical KB showing the etiology of AIDS.	68
9-2	Fragment of the medical KB showing opportunistic pneumonias as complications of AIDS.	70
9-3	Fragment of the medical KB showing the outcomes of the pneumonias.	71
9-4	Fragment of the medical KB showing the results and the complications of BAL for Pneumocystis-carinii infection.	73
9-5	Fragment of the medical KB showing the treatments and their complications for PCP.	74

Chapter 1

Introduction

1.1 Background and Motivations

Clinical decision making is a very challenging task. The multitude of problems, the patient-specificity, and the uncertainty involved all contribute to the intricacy of the process. In recent years, decision analysis has gradually been recognized as a powerful technique for selecting the optimal strategies in difficult clinical decision problems [19, 28]. By explicating the important factors in formulating a *decision model*, this method also helps the clinicians gain better insights into the problems.

Nevertheless, clinical decision analysis is still a very expensive process. Each step in the procedure—problem formulation, preference assessments, decision model construction, probability assessments, and decision model evaluation—is very time-consuming and skill-intensive. The large, complex, and ever-changing body of medical knowledge involved further complicates the approach.

1.1.1 Knowledge-Based Decision Systems

In attempting to automate the decision analysis process in knowledge-intensive domains, efforts in developing *knowledge-based decision systems*¹ have emerged. These systems combine artificial intelligence, particularly knowledge-based systems, and decision analysis techniques to provide decision assistance for a wide range of problems in a domain. It is hoped that by capturing the relevant knowledge in the knowledge bases, a well-trained analyst or a domain expert would seldom, if at all, be needed in the automated decision process. Consequently, the cost of using the decision-analytic method in decision making could be greatly reduced [13].

Some existing knowledge-based decision systems have knowledge bases in the form of decision models. In these systems, notably the PATHFINDER system by Heckerman

¹Other names include *intelligent decision systems* [15] and *normative expert systems* [13].

1.1.3 Supporting Clinical Decision-Modeling

In addition to capturing the relevant domain and decision-analytic relations, a knowledge base for supporting clinical decision-modeling should satisfy the following criteria:

- *Reusability*: Given the high costs for constructing a large knowledge base, we may want to use it for other purposes, *e.g.*, differential diagnosis, therapeutic planning, *etc.*
- *Modifiability*: Complete knowledge may not be available when the knowledge base is first put into use; information should be easily incorporated incrementally into the knowledge base. Moreover, sometimes we may wish to combine the knowledge from different domains.
- *Explanation support*: The decisions made by an automated procedure are credible only if it uses the domain knowledge correctly; explanations are indispensable for clarifying how the domain knowledge is being used.

Therefore, the structure of the knowledge base must reflect the nature of the domain knowledge as well as the decision-analytic knowledge. Moreover, the semantics of the underlying knowledge representation should be expressive, precise, and adaptive; in particular, it should not be restricted by the specific evaluation mechanisms for the decision models, *e.g.*, folding back of a decision tree, graph reduction of an influence diagram, *etc.*

Judging from these observations, the constructive approach is more likely to be effective than the static approach to clinical decision-modeling. Besides the advantages mentioned earlier, the constructive approach allows the development of flexible, multi-purpose knowledge bases. These features outweigh the relatively elaborate model-construction process required.

Some decision systems mentioned earlier have addressed various issues in constructive decision-modeling, including how to construct decision models correctly and efficiently from the knowledge base. To support constructive clinical decision-modeling, however, both the domain and the decision-analytic knowledge need to be carefully analyzed and formalized. So far, very little effort has been made in this direction.

1.2 Objectives

This project aims to characterize the knowledge for supporting constructive decision-modeling in medicine. Characterizing such knowledge illuminates the representational and computational requirements for automating clinical decision analysis. The decision models in question may be decision trees, influence diagrams, or *qualitative*

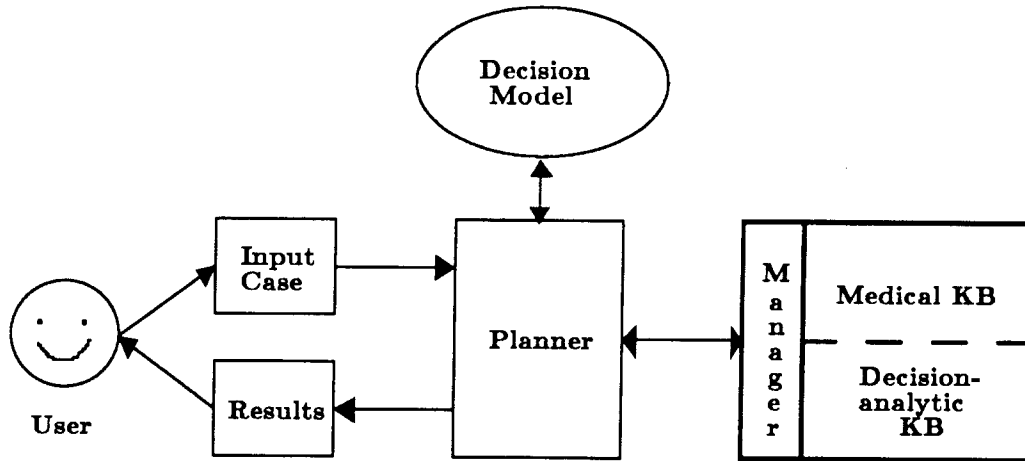


Figure 1-1: A knowledge-based clinical decision system.

1.4 Overview of Thesis

This introductory chapter has briefly discussed the nature of knowledge-based decision systems, how they might assist clinical decision making, and how we plan to analyze and formalize the representation requirements. The remainder of this thesis is organized as follows. Chapter 2 describes the medical domain addressed in this work, namely, pulmonary infectious diseases with AIDS, and explains why it is interesting from the decision-analytic perspective. An example case is also included for illustration. Chapter 3 details the analysis of the clinical decision analysis process and characterizes the different types of medical knowledge involved. Chapter 4 compares the knowledge requirements found in the analysis with some existing representations. The motivations behind our own representation formalism are then presented. Our representation design is documented in Chapter 5 and Chapter 6. Chapter 5 describes the representation of concepts; the interactions among these concepts are discussed in greater details in Chapter 6. Chapter 7 illustrates the contents of the medical KB. Chapter 8 shows a comprehensive example of how the medical KB supports clinical decision-modeling and postulates how it could be used in general. Finally, Chapter 9 lists the accomplishments, limitations, and future directions of this work.

Chapter 2

The Domain

2.1 Pulmonary Infectious Diseases with AIDS

Acquired immune deficiency syndrome (AIDS) is a disease caused by human immunodeficiency virus (HIV) infection. The defective immune system of AIDS patients predisposes them to a variety of infections and diseases that usually cause less or little harm in immunocompetent hosts. Pulmonary diseases, especially pulmonary infectious diseases (PIDs), are seen in nearly all AIDS patients [9, 16]. Depending on the immunosuppression status, different pulmonary infections (PIs)¹ occur with different frequencies; multiple simultaneous infections are common. The resulting PIDs are usually very serious and may be rapidly fatal if not treated properly [31]. For instance, *Pneumocystis carinii* pneumonia (PCP) is the most frequently diagnosed AIDS-related opportunistic PID; it occurs at least once in about 70% of patients, with a 20-30% mortality rate per episode [9]. Table 2.1 shows some common AIDS-related PIDs and their relative occurrence frequencies² [9]. The relative frequencies reflect the percentages of the individual PIDs, with respect to all the PIDs, that might occur in an AIDS patient. These frequencies add up to more than 100% because simultaneous multiple diagnoses are possible.

A number of diagnostic tests exist for PIs with suspected AIDS, ranging from the less efficient but non-invasive procedures, *e.g.*, sputum examination and gallium scanning, to the more efficient but invasive techniques, *e.g.*, bronchoscopic bronchoalveolar lavage (BAL) and bronchoscopic transbronchial biopsy (TBBx). Different diagnostic modalities are chosen depending on the situation; factors considered include the general condition of the patient, presence of other diseases, recent medical interventions, *etc.* If the initial test results are not adequate, further testing may be necessary. Table 2.2 shows a list of common diagnostic tests for the PIs with corresponding

¹Although the word *infection* is commonly used to denote both the infection itself and the disease it causes, we are distinguishing the two for clarity and easier illustrations.

²A list of medical terms referenced in this thesis can be found in Appendix B.

Diagnosis	Treatment	Essential Complications
PCP	Trimethoprim-Sulfamethozazole	Rash, fever, leukopenia, hemolytic anemia
	Intravenous Pentamidine	Neutropenia, leukopenia, renal dysfunction
	Aerosol Pentamidine (<i>prophylactic</i>)	Cough, bronchospasm
Pul. toxoplasmosis	Pyrimethamine-Sulfadiazine	Rash, fever, neutropenia
Pul. TB	Antibiotics regimen	Isoniazid hepatitis
MAI complex	Antibiotics regimen	Isoniazid hepatitis
Pyo. bac. pneumonia	Antibiotics regimen	Allergic reactions, renal dysfunction
Legionellosis	Erythromycin	Gastrointestinal upset
Pul. cryptococcosis	Amphotericin B-Flucytosine	Fever, chills, renal dysfunction, liver dysfunction
Oth. fungal infections	Amphotericin B	Fever, chills, renal dysfunction
CMVP	No standard treatment	
HSP	Acyclovir	

Table 2.3: Treatments for AIDS-related Pulmonary Infections

tating rapid development and evaluation of safer and more cost-effective management strategies [31]. The diversity and uncertainty involved in diagnosis and therapy make PIDs with suspected AIDS a challenging domain for decision analysis. The richness of the underlying knowledge is ideal for examining the representational requirements for supporting automated decision making.

2.2 An Example Case

An example case in the domain of PIDs with suspected AIDS, a simplified version of a case presented in [31], 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. He has no known drug allergies and has never been tested for HIV infection. The initial impression was pneumonia possibly due to opportunistic infection in this patient with suspected AIDS. The problem is to investigate whether or not to employ empiric therapy for PCP, and how non-invasive diagnostic tests such as sputum examination and gallium scanning compare with invasive procedures such as BAL and TBBx.

In the following chapters, we will examine the decision-making process with ref-

Chapter 3

Clinical Decision Analysis

The decision-analytic approach to the clinical decision making process can be viewed as follows:

Background characterization --> Context establishment
--> Problem formulation --> Model construction --> Model evaluation

Given the background information, we establish the domain context¹ in which the problem is formulated. We then identify the relevant concepts: variables, actions, and outcomes to form a decision model. The decision model is then evaluated with respect to some criteria, *e.g.*, life expectancy, expected monetary value, *etc.*

With reference to the example case presented in Section 2.2, we shall now examine each step of the decision making process, with emphasis on analyzing the required representational support.

3.1 Background Information Characterization

Given the background information, we have to characterize the variables concerned, the actions available, and the complications involved. These events can be divided into the following categories:

- *General history*: Relevant general conditions and previous experience.
- *Signs and Symptoms*: Conditions observed by the physician or reported by the patient.
- *Laboratory findings*: Conditions described by laboratory findings.

¹This is different from the *decision context* [4, 15] which refers to all the assumptions, constraints, variables, and alternatives considered in the decision problem.

3.2 Clinical Context Establishment

Studies show that “framing the problem,” *i.e.*, selecting the context or task environment in which the problem is to be solved, is an important aspect of clinical decision making [18]. The context, also known as the problem space or decision frame [26, 7, 35], is selected with only a few clues [17]; it serves as a basis for expectation, and sets the stage for asking further questions and evaluating available options in the decision making process. Sometimes several contexts may be involved.

In the clinical setting, a context is usually indicated by a suspected disease, a syndrome, *i.e.*, a set of signs and symptoms that convey special meanings, or a general diagnostic category, *e.g.*, an acute respiratory disorder [18]. In this case, the clinical context is “PIDs with suspected AIDS.” This context is established by simply identifying the suspected diseases in the input information. Given the characterized background information, supporting clinical context establishment again requires categorization or classification of the concepts involved.

Note, however, that the whole purpose of establishing a context is to allow access to the context-sensitive information. In particular, the PIDs considered with suspected AIDS are normally opportunistic infectious diseases; some of the non-opportunistic infectious diseases such as tuberculosis (TB) are aggravated with AIDS, while others are rarely to be found. Moreover, due to the varying nature of different kinds of immunosuppression, AIDS and other immunodeficiencies may induce similar opportunistic infections at different frequencies [5]. For instance, though usually associated with AIDS (see Table 2.1), PCP occurs at a much lower frequency in patients with bone marrow transplants. Therefore, such context-sensitive knowledge must be expressible in the knowledge base.

3.3 Decision Problem Formulation

Guided by the characterized background information, a decision problem is formulated within the clinical context by identifying:

- all (or the most important) diseases/hypotheses that may be involved;
- the relative significance of all these concepts;
- all the possible outcomes/complications of these concepts;
- all the actions available;
- the effects of the actions on the concepts and their outcomes and possible complications; and
- the evaluation criteria.

Risk-factor-of-HIV-infection: IV-drug-abuse	Empiric-treatment-for-PCP: TMP-SMZ IV-pentamidine Aerosol-pentamidine
HIV-infection	
AIDS	
Pulmonary-infectious-disease: PCP Pul. toxoplasmosis Pul. TB MAI-complex Pyogenic-bacterial-pneumonia Legionellosis Pul. cryptococcosis Other-PIDs	Treatment: TMP-SMZ (for PCP) IV-pentamidine (for PCP) Aerosol-pentamidine (for PCP) Pyrimethamine-sulfadiazine (for pul. toxoplasmosis) Antibiotics (for pul. TB, MAI-complex, and Pyo. bac. pneumonia) Erythromycin (for legionellosis) Amphotericin-B-Fluocytosine (for pul. cryptococcosis) Amphotericin-B (for other fungal PIDs) Acyclovir (for HSP)
Symptom/sign: Low-grade-fever Non-productive-cough Dyspnea	Treatment-complication: (Each complication may be for various treatments.) Allergic reactions Bronchospasm Chill Cough Fever Gastrointestinal-upset Hemolytic-anemia Isoniazid hepatitis leukopenia Liver dysfunction Neutropenia Rash Renal dysfunction
ABG-result: Mild-hypoxemia-on-room-air	Cost
CXR-result: Bilateral-diffuse- interstitial-infiltrates	Morbidity
Disease-outcome: Cured Improved Not-improved Worsened Death	Mortality
Test: Sputum-examination Gallium-scanning BAL TBBx	Quality-adjusted-life-expectancy
Test-result: Positive: (for each PID) Negative (for each PID)	Utility
Test-complication: Death (for BAL and TBBx) Fever (for BAL) Worsening-oxygenation (for BAL) Pneumothorax (for TBBx) Hemorrhage (for TBBx)	

Table 3.2: Concepts Involved in Decision Problem

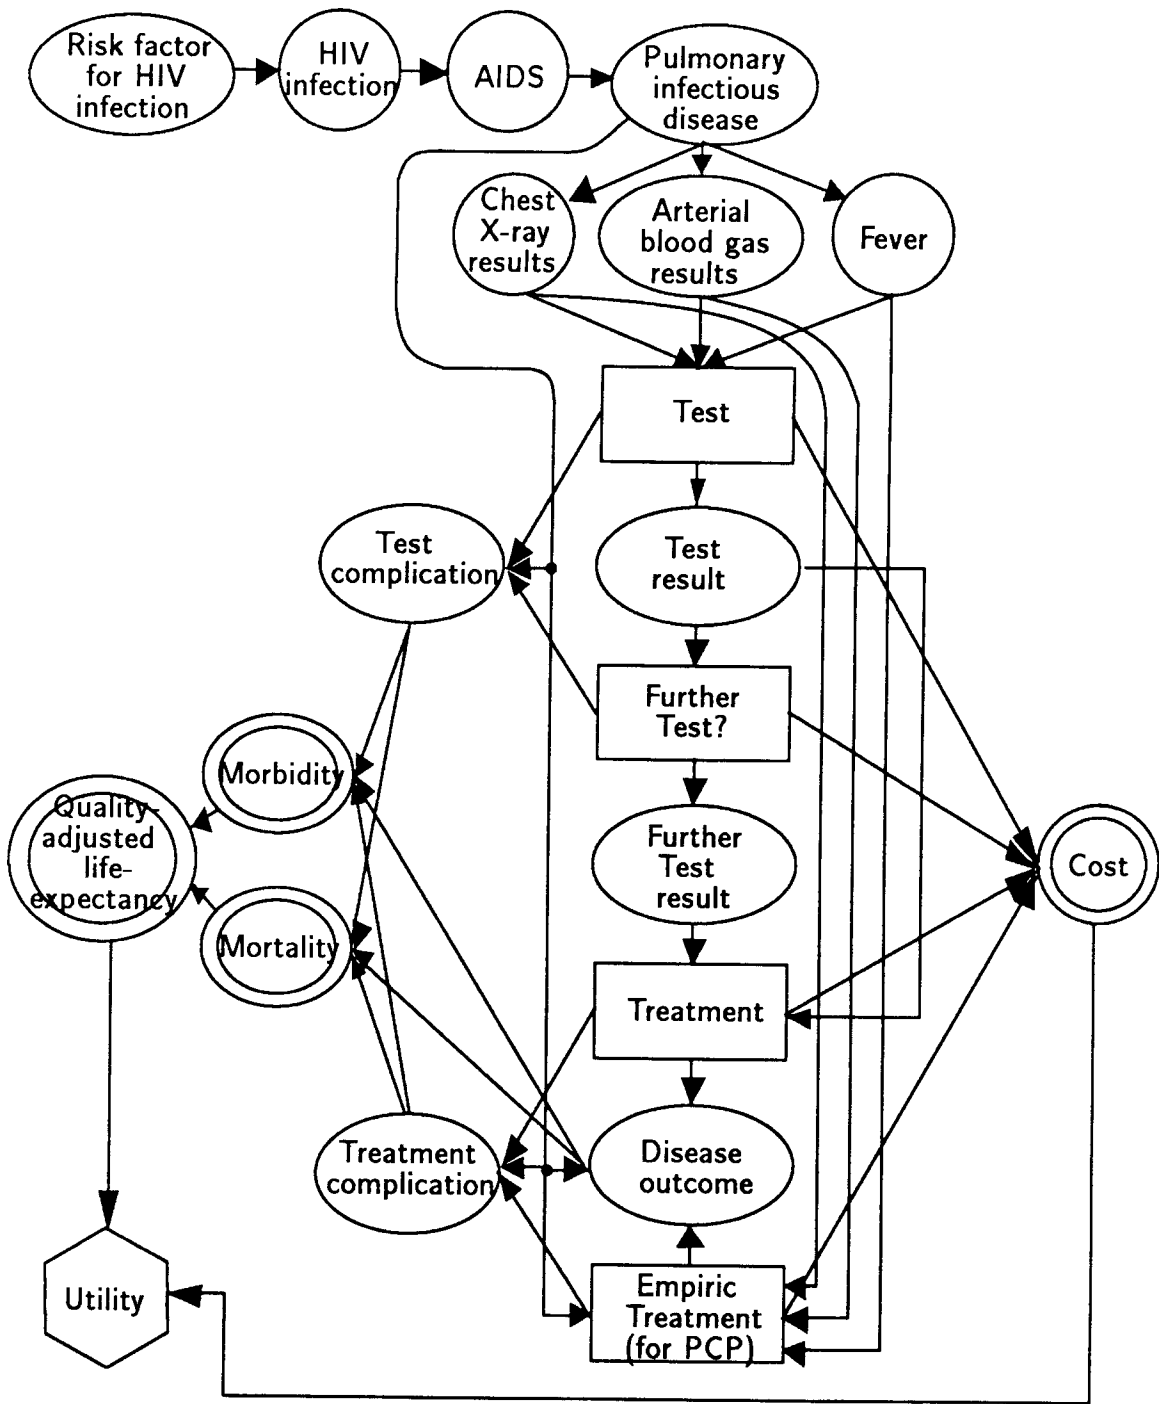


Figure 3-1: A QPN for the Example Case. The evaluation criteria are expected monetary cost and quality-adjusted life-expectancy, *i.e.*, a measure of time remaining in a patient's life, taking into account the inconveniences caused by the illness (morbidity).

treatments, outcomes, and other entities related to the decision problem. Guided by this set of requirements, we shall evaluate some relevant representation formalisms in the next chapter. We shall also propose a representation design approach aimed at fulfilling the above requirements.

Chapter 4

Representation Design Approach

4.1 Related Work

Many existing representation frameworks have some desirable features for our requirements. These include formalisms such as semantic networks, frame-based languages, term subsumption languages descended from the KL-ONE framework [1], and systems such as OWL[11, 33] and NETL [8]. All these frameworks support hierarchical representation of categorical information, and some allow limited expression of context-dependent information. None of these frameworks, however, satisfy all the requirements concluded in Chapter 3. Since each formalism has unique assumptions and semantics, it is difficult to adapt or combine the frameworks for our purposes. Hence, we have decided to design a new representation by assimilating the desirable features of existing ones.

In this section, we shall first examine the representations in some existing knowledge-based decision systems. A brief survey of some other relevant frameworks is then presented; the discussions will be based only on the expressiveness of the frameworks. The lessons learned in these exercises serve as the basis for our own design motivations and approach, to be described later in the chapter.

4.1.1 Representations in Existing Knowledge-Based Decision Systems

In Chapter 1, we have briefly discussed why the static approach to clinical decision-modeling is inadequate. The major shortcomings of this approach result from the rigidity of the knowledge bases. The underlying representations of these knowledge bases cannot accommodate most of our representational requirements. In particular, the decision model-like semantics does not allow hierarchical representation of clinical concepts and their context-sensitive interactions. Such knowledge bases, therefore, do not reflect the domain structure of clinical medicine.

commodating these frameworks to support approximate reasoning, *i.e.*, finding out facts that are not absolutely true or false, but *believed* to a certain degree, have only emerged recently. Some of these efforts attempt to accommodate the uncertainty models by re-interpreting the semantics of existing representations, while others try to couple the two to form a coherent framework.

For instance, in the network representation developed by Lin and Goebel [22], both subsumption and causal relationships are expressible. Probabilistic interpretations are given to parts of the causal network, called the *scenarios*. These scenarios can be considered as contexts with different probability distributions. Although the scenarios are not hierarchically arranged, their probabilistic rankings are preserved across the subsumption relationships. Nevertheless, this network formalism does not allow the properties, and hence the nature of each node or event to be explicitly represented.

Another example of adapting an existing representation to accommodate uncertain information is the work done by Yen and Bonissone [40]. This work attempts to generalize the semantics of term subsumption languages with an approximate reasoning model, such as fuzzy logic or possibility theory, to support plausible inferences. Term subsumption languages, which are mainly descendants of the KL-ONE [1] family, are formalisms for defining *terms* or *concepts*¹ in terms of their logical necessary and sufficient conditions. The concepts in these languages are related by the subsumption relationship; properties can be inherited in the subsumption hierarchies. On the other hand, non-definitional relations among the concepts are not expressible in these frameworks. There is also no general notion of context-dependent definitions. The generalized semantics by Yen and Bonissone allows inferences to measure the degree to which an instance satisfies a terminological expression. For example, instead of an absolute answer such as “Disease-X is an infectious-disease”, an approximate answer such as “Disease-X is likely (with probability 0.6) an infectious-disease” is derivable from a term subsumption framework. The underlying logical definitions of the concepts, however, remain unchanged.

A more general framework that integrates a categorical representation formalism with an uncertainty model is developed by Saffiotti [29]. Uncertain knowledge is represented in this hybrid framework in two components, one dealing with absolute or categorical knowledge and one dealing with the uncertainty of this knowledge. Any formal representation formalism and uncertainty model may constitute the two components in the framework, *e.g.*, first-order logic with Dempster-Shafer theory, term subsumption language with probability theory, *etc.* We believe this work is an important step toward the theoretical foundations of integrating categorical and uncertain knowledge. The expressiveness and hence the usefulness of the framework, however, depend solely on the component formalisms. It is this kind of integration that we wish to explore in our own design.

¹Term subsumption languages are also called *concept-based languages*.

All the above phenomena can be modeled or described as some discrete structures, called *concepts*² in a representation framework. Conceptual modeling has gained much attention in philosophy, cognitive psychology, and artificial intelligence. Two major views that have motivated our own approach are as follows:

View 1 *Concepts as intensional descriptions.*

In [39], Woods argues for the need for representing *intensional* concepts. Different concepts may have the same extension, *i.e.*, a referent that exists in a world, *e.g.*, both of the concepts “morning star” and “evening star” refer to the planet Venus, and some concepts may have no extension at all, *e.g.*, the concepts “4-sided triangle” and “positive integer less than 0.” Hence, a concept corresponds to an *intension* in *The Meaning Triangle* shown in Figure 4-1 [32, 27]. Such intensions cannot be identified with the predicates in first-order logic, nor with the classes in set theory. Instead, it is useful to identify intensions with the notion of abstract *descriptions* [39].




Figure 4-1: The Meaning Triangle

Specifically, an intensional concept can be regarded as a description that can be analyzed and reasoned with. This description has three major characteristics, as abridged below from [39, page 5-6]:

1. *A description can be satisfied BY something (in the way that a predicate is satisfied by values for its arguments);*
2. *A description can be satisfied IN a situation and can be used to characterize the entity so recognized; and*
3. *A description can be used as a structured plan for creating or reasoning about something that does not yet exist or may never exist.*

The description of a concept usually includes all other concepts that are associated with it. These associated concepts are called the *properties* or *characteristics* of the concept.

²So far we have been using the word “concept” in a very general way. From now on we will adhere to its more formal meaning as characterized in this chapter and defined in the following one.

Proposition 3 *The invariant and context-dependent interactions specified among the concepts adequately reflect the phenomena being modeled.*

Based on the above propositions, we design a framework for describing the classes of concepts and the types of interactions among them in a context-sensitive manner. Our descriptions of the concept hierarchies are very much influenced by the assumptions and the approaches in NIKL [25] and OWL [11, 33], although our presentation is very different. Our descriptions of the interactions among the concepts are based mainly on Wellman's work on the QPN formalism. In the current work, we shall only describe the intended interpretations of the relevant interactions; the interested reader should refer to [37] for the theoretical basis of the QPN framework. Our context-sensitive representation is inspired by the ideas in Hendrix's partitioned network [14].

In the following chapters, we will discuss the overall design of a concept and the different types of interactions among the concepts. we will also describe how to represent context-dependent information in this framework. Due to the large number of issues involved, the descriptions in the following chapters may indeed seem incomplete. we shall leave the intricate details and illustrations to the comprehensive example in Chapter 9.

Chapter 5

Representation of Concepts

5.1 Overview

In our framework, a *concept* is an intensional description of the relational/causal interpretation of an object, a state, a process, or an attribute of these phenomena. In other words, a concept reflects the salient features of the underlying phenomenon through a set of interactions, *i.e.*, correlational/influential/causal relationships with other concepts. These relevant concepts are called the *properties* of the concept being described. For example, the description of the concept **disease**¹ includes properties such as **severity**, **manifestation**, and **treatment**, as well as interactions² such as “**presence-of-disease causes presence-of-manifestation-of-disease**” and “**presence-of-treatment-of-disease alleviates severity-of-disease.**”

The description of a concept is constrained by a set of *categorizers*. A categorizer is a categorical or *class* relationship; it is a binary relation that specifies the properties and the interactions of a concept in terms of those of another concept. By imposing a partial order on the related concepts, a categorizer establishes a unique *perspective* for describing each concept. For example, a concept can be described as “a kind of” another concept or “a part of” another concept. Some common categorizers include the *specialization* (AKO) relation, the *decomposition* (PARTOF) relation, and the *equivalence* (EQV) relation. All the concepts related by a categorizer are said to be in a *categorization*; some categorizations have hierarchical interpretations, while others are more naturally seen as networks.

The categorizers establish some general perspectives for describing a concept. For example, a **pulmonary-infection** is “a kind of” **infection** in general. The description of a concept in these general perspectives is further constrained by a set of *contexts*. A context can be thought of as a “meta-categorizer;” it is a bi-

¹All concepts defined in our framework will be referenced in **typewriter type style**.

²The actual interaction formats will be described in Section 5.4; all interaction formats presented preceding that section are for illustration only.

of a concept in more details; we shall also compare our design decisions with similar features in existing representations whenever appropriate. In this discussion, we are only dealing with concepts of *types* or *classes*; we have not worked on how to represent concepts of *individuals*.

5.2 Denotation of A Concept

More formally, a concept can be denoted by a tuple $(a \# b)$, read “ a of b ,” where a is the *basic identity* of the concept, and b is the context in which it is defined; both a and b are concepts themselves. The basic identity is the most accurate general description of the concept. The context relates to the denoted concept through a CXT relation; it specifies the condition in which the description of the denoted concept is valid, and allows this description to vary, if necessary, from the basic identity. There is a special concept, denoted as \top , which is defined to be itself; any concept defined in the context of \top is in the universal context, *i.e.*, valid in general. For instance, the concept PC-infection shown in Figure 5-1 is denoted as $(\text{PC-infection} \# \top)$.

The notation above allows concepts to be “chained” to form a new concept, analogous to the “role-chaining” notion in KL-ONE [1] and NIKL [25]. The new concept is in the form of: $((\dots(c_1 \# c_2) \# \dots) \# c_n)$, read: “ c_1 of c_2 of \dots of c_n ”. For example, $((\text{duration} \# \text{treatment}) \# \text{disease})$ and $((\text{severity} \# \text{complication}) \# \text{AIDS})$ are read as **duration-of-treatment-of-disease** and **severity-of-complication-of-AIDS** respectively⁴. The chaining expression is associative, *i.e.*, it can also be denoted as $(c_1 \# (\dots \# (c_{n-1} \# c_n) \dots))$. When the meaning is obvious, we will use the shorthand notation $(c_1 \#^* c_n)$ to denote the same concept above. For example, $((\text{presence} \# \text{existence-status}) \# \text{disease})$ can be denoted as $(\text{presence} \#^* \text{disease})$.

5.3 Properties of A Concept

The properties of a concept include its inherent qualities, characteristics, and other relevant concepts that constitute its description *e.g.*, **size** (of a tumor) and **diagnostic-test** (of a disease). Each property is a concept itself. From now on we shall adopt the phrase *base-concept* to denote the concept we are describing, and the phrase *property-concept* to denote the concept corresponding to the referred property. Declaring the properties in a concept can thus be seen as establishing associations with the relevant property-concepts.

Each property of a concept has a list of *values*. These values are concepts related to the property-concepts in a categorization, usually the specialization (AKO)

⁴From now on we shall reference a concept with its denotation or its name interchangeably.

concept **test-of-pulmonary-infection** is derived from the property **test of pulmonary-infection**. This new concept, called a *derived-concept*, is related to the property-concept by a categorizer, but its description is constrained by the base-concept that it is derived from. Each derived-concept is related to the base-concept by a context (CXT) relation, *i.e.*, the derived-concept is defined in the context of the base-concept. For example, the derived-concept **test-of-pulmonary-infection** is a specialization of **test-of-infection**, but the values of its properties **specificity**, **sensitivity**, and **complication**, and hence the corresponding descriptive interactions are specific to **pulmonary-infection**. We shall discuss more about expressing context-dependent information in Chapter 7

The properties of a concept in our framework are analogous to the *roles* in term subsumption languages and the *slots* in frame-based languages. The difference is that the properties alone do not completely describe a concept; they serve only as *indices* to the interactions that constitute the meaning of a concept. These interactions are expressed in terms of the corresponding derived-concepts.

5.4 Interactions of Concepts

An interaction is a correlational/influential/causal relationship between two or more concepts. There are many such relevant behavioral relationships in the decision-modeling context. In one extreme, the interactions can be described in English words such as “causes,” “alleviates,” “indicates,” *etc.*; in another extreme, they can be expressed as numeric conditional probabilities between two or more concepts. As mentioned, the interactions in the description of a concept are specified in terms of its derived-concepts, *i.e.*, the concepts derived from the properties of the concept being described. For instance, the properties **treatment** (of a **disease**) and **severity** (of a **disease**) are involved in the interaction “**presence-of-treatment-of-disease** alleviates **severity-of-disease**.”

As shown in Figure 5-1, interactions can be specified among all the derived-concepts involved in the description. When the values of a property are specified, the interactions are expressed in terms of the more specific values. For example, the property **treatment** of PCP (a specialization of **disease**) has the values **TMP-SMZ-treatment** and **pentamidine-treatment**, hence the relevant interactions in the description of PCP are “**presence-of-TMP-SMZ-treatment-of-PCP** alleviates **severity-of-PCP**” and “**presence-of-pentamidine-treatment-of-PCP** alleviates **severity-of-PCP**.”

Moreover, if we do know all the properties of the derived-concepts, and the properties of those properties, we can “chain” the descriptions as well. For example, the derived-concept **risk-factor-of-PC-infection** also has the property **existence-status**, hence the value **presence** (of **existence-status** of **risk-factor** of PC-infection) can be involved in the description of PC-infection.

Definition 5.1 (Association) Let \mathcal{C} be the set of all concepts in a decision model, $\forall a, b, x, y \in \mathcal{C}$, where 1) $x \neq a$, and x relates to b via some direct or indirect interactions; and 2) $y \neq b$, and y relates to a via some direct or indirect interactions:

$$a \overset{a}{\longleftrightarrow} b \iff Pr(b|a, x) \neq Pr(a) \text{ or } Pr(a|b, y) \neq Pr(a).$$

The arrow-heads in the above notations are present only to indicate the conditioning direction at the time of the encoding. There are no causal or temporal implications.

Precedence Links

The *precedence links* indicate temporal orders with unknown type of probabilistic influence. In other words, they are associational links with known temporal precedence.

Definition 5.2 (Precedence) Let \mathcal{C} be the set of all concepts in a decision model, $\forall a, b \in \mathcal{C}$:

$$a \overset{p}{\rightarrow} b \iff a \text{ temporally precedes } b$$

The arrow-heads in the above notation denotes the direction of temporal precedence between two concepts.

Influential Links

The *influential links* denote conditional probabilistic dependencies with unknown temporal precedence among some concepts. The concepts are viewed here as probabilistic random variables; the values of these random variables are the instances of the concepts⁵. There are two types of influences, corresponding to the *positive-* and *negative-influences* in the QPN formalism.

In a nutshell, a positive-/negative-influence is defined as follows: For two *binary* concepts A and B, *e.g.*, **presence-of-infection** and **presence-of-fever**, if A positively-/negatively-influences B, then the presence of A increases/decreases the probability of the presence of B, with all other things being unchanged. If A and B are *continuous* concepts, *e.g.*, **severity-of-disease** and **degree-of-morbidity**, then higher values of A increase/decrease the probability of higher values of B, with all other things being unchanged.

⁵All instances involved in the definitions of interactions, and later in Chapter 6, of categorizers, are *concepts* of instances. These are not to be confused with the referents or extensions of the concepts in a world.

causal relationship, and for $A \xrightarrow{c/i} B$, whether A is needed just to initiate or inhibit B, or A should be present while B occurs or is absent [23]. Finer granularity of the causal/inhibitive links could be derived in the future when necessary.

5.4.2 Indirect Interactions

The spectrum of interactions reflect the varying certainty of information at the time of encoding. To support dynamic decision modeling, a knowledge base should allow a decision maker to select the relevant concepts and derive the direct or indirect relationships among them. There are two forms of interactional indirections: *interaction chains* and *parallel interactions*. For example, with reference to Figure 5-3, we have to derive the interaction between the **presence-of-risk-factor-of-PC-infection** and the **presence-of-PC-infection** if we do not care about the **pathogen** and the **infective-route** involved.




Figure 5-3: Part of the interaction-model of PC-infection

Table 5.1 defines the indirect effects of the interactions. The definitions are consistent with those for the \otimes operator for combining *influence chains* and the \oplus operator for combining *parallel influences* in QPN. The corresponding operations are commutative, associative, and distributive, just like ordinary multiplication and addition [37, page 268].

Chapter 6

Categorization of Concepts

As mentioned in Chapter 5, a categorizer is a binary relation that groups concepts, according to their descriptions, into a categorization. By knowing the position of a particular concept with respect to another concept in a categorization, we can infer the description of the former from the latter. The descriptive inference in a categorization is called *inheritance*.

A concept can be involved in multiple categorizations. Multiple categorizer-specifications in a concept description is by default interpreted as a conjunctive form; the relations in different categorizations are orthogonal. For example, specifying $\text{AKO}(\mathbf{A},\mathbf{B})$, $\text{AKO}(\mathbf{A},\mathbf{C})$, $\text{PARTOF}(\mathbf{A},\mathbf{D})$ in \mathbf{A} indicates that \mathbf{A} is both AKO B and AKO C in the specialization categorization, and \mathbf{A} is PARTOF D in the decomposition categorization.

A concept can also be characterized with a disjunction of similar or different categorizers by explicitly applying the special relation OR . For example, specifying $\text{AKO}(\mathbf{A},\mathbf{B})$, $(\text{OR}(\text{AKO}(\mathbf{A},\mathbf{C}) \text{AKO}(\mathbf{A},\mathbf{D})))$ indicates that \mathbf{A} is both AKO B and AKO C , or both AKO B and AKO D in the specialization categorization. \mathbf{A} is said to be in two different categorizational *tracks*. The properties and interactions inherited in different tracks (of the same or different categorizations) are orthogonal. In other words, the descendants (in a hierarchical categorization) of a multi-tracked concept can only access its description in a single track. The choice among the tracks, however, can be delayed until necessary. For instance, referring to the previous example, we can have the relation $\text{AKO}(\mathbf{E},\mathbf{A})$ without specifying the track of \mathbf{A} that \mathbf{E} is in. \mathbf{E} will then simply inherit all the unambiguous properties and interactions, together with the disjunctive specification of specialization relations in \mathbf{A} .

In the following discussions, let \mathcal{C} be the set of all concepts. Let Ω be the set of categorizers. Let $\mathcal{O}_\omega \subseteq \mathcal{C}$ be the set of concepts in a categorization related by categorizer $\omega \in \Omega$.

For now, we assume the properties and the interactions that a concept inherits in the AKO categorization have no conflict nor inconsistency. In other words, we assume *monotonic* inheritance in the current framework. All properties and interactions inherited by a concept are additive with respect to its parent concepts.

6.1.2 The Generalization Relation

The generalization (**GEN**) relation is the dual of the specialization relation. The categorization of the generalization relation is exactly the same as that of the specialization relation, except that when it is viewed as a directed graph, the directions of the links are reversed.

Definition 6.2 (Generalization) For all $a, b \in \mathcal{C}$, and for $GEN \in \Omega$ where $GEN \subseteq \mathcal{C} \times \mathcal{C}$:

1. $GEN \stackrel{def}{=} \{ (a, b) | b \subset a, \text{ i.e., } \forall \beta, \beta \in b \implies \beta \in a \}$.
2. Let $gen : \mathcal{C} \longrightarrow 2^{\mathcal{C}}$ be a function defined on GEN :
 $gen(a) = \{ b | (a, b) \in GEN \}$.

Two major properties are observed for the **GEN** categorizer:

1. $a \in \mathcal{O}_{GEN} \iff \exists b, (a, b) \in GEN \text{ or } (b, a) \in GEN$.
2. The GEN relation is irreflexive, asymmetric, and transitive.

The generalization relation, incidentally, is not very useful in our framework. A particular concept may be a generalization of numerous other concepts, along different dimensions; since the generalization (and hence specialization) dimensions are not explicitly expressed in the categorization, it is unclear how to infer the description of the general concept from its specializations. Similarly, it is also difficult to infer the descriptions of concepts which generalize a particular concept in a multiply-connected hierarchy.

We postulate that the generalization relation may be useful in some special situations, *e.g.*, when a single generalization dimension, say, **pathogen** is identified, we could find out what other types of **infection** are possible besides **protozoal-infection** by calling the functions $gen(\text{protozoal-infection})$ and $ako(\text{infection})$ along the **pathogen** dimension. We have not, however, looked into these issues in the current work.

inherited from different component concepts, the inherited values of each property are usually combined, modified, or filtered to arrive at the correct values for the aggregate concept. Different relations, *e.g.*, maximization, minimization, union, magnitude-comparison, *etc.*, govern the derivation of the correct values because of the different part-whole relations between the aggregate and the component concepts. Presently, we have not been able to explore much on these issues. We postulate, however, that a comprehensive characterization of the inheritance behavior would involve explicit representation of the decompositional dimensions that relate the aggregate concept to its components.

6.2.2 The Aggregation Relation

The aggregation (**CONTAIN**) relation is the dual of the decomposition relation. The categorization of the aggregation relation is the same as that of the decomposition relation, except that when it is viewed as a directed graph, the directions of the links are reversed.

Definition 6.4 (Aggregation) For all $a, b \in \mathcal{C}$, and for $CONTAIN \in \Omega$ where $CONTAIN \subseteq \mathcal{C} \times \mathcal{C}$:

1. $CONTAIN \stackrel{def}{=} \{(a, b) | \forall \beta, \exists \alpha, \beta \in b \implies \beta \in \alpha, \alpha = \{\beta_1, \beta_2, \dots\} \in a\}$.
2. Let $contain : \mathcal{C} \rightarrow 2^{\mathcal{C}}$ be a function defined on $CONTAIN$:
 $contain(a) = \{b | (a, b) \in CONTAIN\}$.

Two major properties are observed for the **CONTAIN** categorization:

1. $a \in \mathcal{O}_{CONTAIN} \iff \exists b, (a, b) \in CONTAIN$ or $(b, a) \in CONTAIN$.
2. The **CONTAIN** relation is irreflexive, asymmetric, and transitive.

The properties and interactions of the concepts are downward inheritable² in the aggregation categorization. Since a general aggregation relates numerous component concepts decomposed along different dimensions, it is difficult to infer the description of the aggregate concept from the component concepts. Therefore, we assume an *exhaustive* definition of the aggregation relation in the current framework. For example **wasting-syndrome** **CONTAIN** only **fever**, **diarrhea**, **night-sweat**, and **weight-loss**. In this definition, an aggregate concept simply inherits all the properties and interactions of its component concepts. For example, **wasting-syndrome** has all the properties of its component concepts. The inheritance is again assumed to be additive and non-contradictory with respect to the component concept. The shortcomings of adopting this assumption are the same as those mentioned in Section 6.2.1.

²Assuming that the *direction* of the categorization is from the aggregate concept to its decomposed concepts, *i.e.*, opposite to that in the **PARTOF** categorization.

Definition 6.6 (Structural-copy) For all $a, b \in \mathcal{C}$, and for $SC \in \Omega$ where $SC \subseteq \mathcal{C} \times \mathcal{C}$:

1. $SC \stackrel{def}{=} \{(a, b) | \exists \alpha, \alpha \in a \implies \alpha \in b, \text{i.e., } a \cap b \neq \emptyset\}$
2. Let $sc: \mathcal{C} \rightarrow 2^{\mathcal{C}}$ be a function defined on SC :
 $sc(a) = \{b | (a, b) \in SC\}$.

Two major properties are observed for the SC categorizer:

1. $a \in \mathcal{O}_{SC} \iff \exists b, (a, b) \in SC \text{ or } (b, a) \in SC$.
2. The SC relation is irreflexive, asymmetric, and transitive.

Intuitively, the SC relation provides a means for different concepts to share description under different constraints or situations, *e.g.*, in the presence of another **disease**. These extra constraints or situations are usually captured in the context (CXT) relations of the involving concepts, to be described in Chapter 7

6.5 Relationships Among Different Categorizations

In the definitions above, monotonic inheritance is assumed in each categorization of concepts. The properties and interactions are inherited independently from the different categorizations in a concept description. We now further assume that monotonic inheritance is applicable across all the categorizations in the framework. In other words, all properties and interactions inherited and specified are consistent and complete, with respect to the available information in the knowledge base, in each concept description. Consequently, the indirect interactions described in Section 5.4.2 are only relevant in the description of each concept.

Chapter 7

Context-dependent Representation of Concepts

The context (CXT) relation between two concepts constrains the description of one concept to the context, or perspective, of the description of the other concept; the properties and interactions of the former are specified, categorized, and interpreted with respect to the latter. This arrangement leads to two important consequences: 1) the same phenomenon can be represented as different concepts in different situations; and 2) context-dependent information can be hierarchically arranged.

7.1 The Context Hierarchy

The formal definition of a context relation is yet to be worked out. We shall try to describe its intended semantics as detailed as possible in the following sections. The context relation is informally defined as follows:

Definition 7.1 (Context) *Let \mathcal{C} be the set of all concepts. Let Ω be the set of categorizers. Let $\mathcal{H}_{CXT} \subseteq \mathcal{C}$ be the set of concepts related by the context relation $CXT \subseteq \mathcal{C} \times \mathcal{C}$, read “in the context of.” For all $a, b \in \mathcal{C}$:*

1. $CXT \stackrel{\text{def}}{=} \{(a, b) | \forall \alpha, a = (\alpha \# b), \exists \omega, \omega \in \Omega, \text{ s.t. } \omega(\alpha, \beta) \text{ where } \beta \text{ is a property of } b\}$.
2. Let $\text{cxt} : \mathcal{C} \rightarrow \mathcal{C}$ be a function defined on CXT :
 $\text{cxt}(a) = b$ where $(a, b) \in CXT$.

Three major properties are observed for the CXT relation:

1. $\forall c \in \mathcal{C}, (c, \top) \in CXT$.
2. $\mathcal{C} \equiv \mathcal{H}_{CXT}$.


```

(defconcept COMPLICATION-OF-DISEASE

  (;;; Contexts
   (#CXT disease))

  (;;; Categorizers
   ($AKO complication-of-process)
   ($OR ($SC disease)
         ($SC physiological-state))

  (;;; Properties
   (!PROCESS disease))

  (;;; Interactions
   (affects (presence #* disease) (presence #* *))))

```

Figure 7-1: Complication of disease

a complication-of-disease. The context hierarchy thus established enables us to express facts like: “tuberculosis is aggravated in patients associated with AIDS.”

7.3 Contextual Effects On Concept Interactions

Expressing the interactions of concepts in terms of temporally ordered qualitative influences, even though simple, intuitive, and useful in general, cannot capture many interesting patterns. For example, to express facts like: “PCP is the opportunistic infectious disease complication most frequently associated with AIDS,” “a positive test-result of a visualization-test for *Pneumocystis carinii* infection will show the presence of *Pneumocystis carinii*,” or “opportunistic infections can only happen when the enabling-factors are present,” we need to model the *contextual effects* on the interactions themselves. Contextual effects on the interactions characterize changes in the interactional patterns among some concepts in the presence of some other interactions. These changes usually affect only the probabilistic component of the interactions; the temporal order of these interactions remains intact.

There are two kinds of contextual effects on the interactions of concepts: *many-to-one* and *one-to-many*.

7.3.1 Many-To-One Effects

Given an interaction between two concepts A and C, i_{ac} , the presence of an interaction between another concept B and C, i_{bc} , may sometimes alter the nature of both interactions. In other words, the combined effect of both A and B on C may be different from a simple addition of the two independent effects. There are three types of

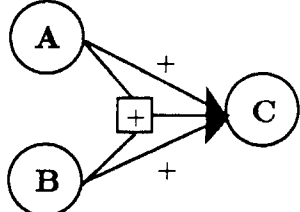
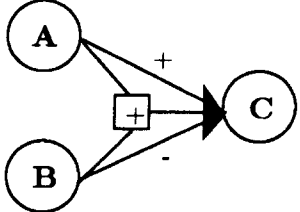
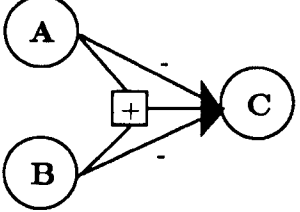
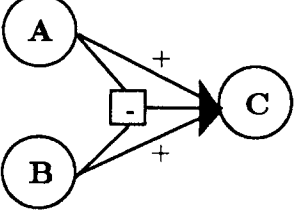
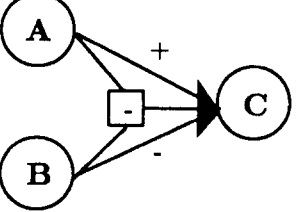
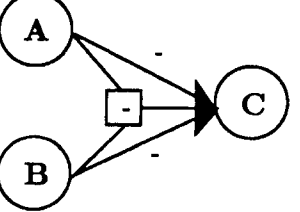
Case	Description
<p>a)</p> 	<p>Positive effect of i_{ac} is larger for higher values of B; positive effect of i_{bc} is larger for higher values of A; combined positive effect is higher than the addition of both positive effects.</p>
<p>b)</p> 	<p>Positive effect of i_{ac} is larger for lower values of B; negative effect of i_{bc} is larger for higher values of A.</p>
<p>c)</p> 	<p>Negative effect of i_{ac} is larger for lower values of B; Negative effect of i_{bc} is larger for lower values of A.</p>
<p>d)</p> 	<p>Positive effect of i_{ac} is larger for lower values of B; positive effect of i_{bc} is larger for lower values of A; combined positive effect is lower than the addition of both positive effects</p>
<p>e)</p> 	<p>Positive effect of i_{ac} is larger for higher values of B; negative effect of i_{bc} is larger for lower values of A.</p>
<p>f)</p> 	<p>Negative effect of i_{ac} is larger for higher values of B; Negative effect of i_{bc} is larger for higher values of A.</p>

Table 7.1: Synergistic effects on interactions. i_{ac} denotes the interaction between concepts A and C, while i_{bc} denotes the interaction between concepts B and C.

1 the strongest, is used to indicate the strength of the individual interaction in the context. The relative strengths of the individual interactions can be approximated by normalizing the individual strengths to 1, with respect to the context.

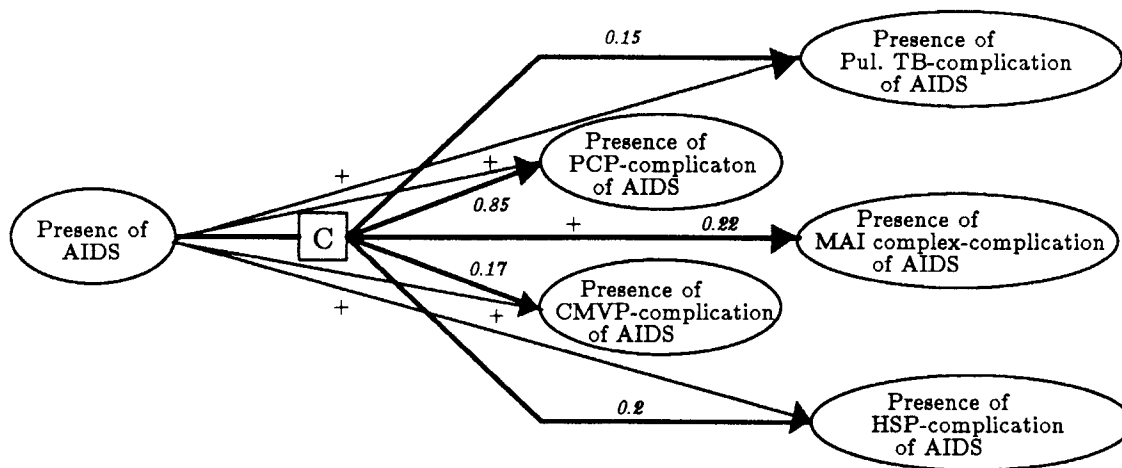


Figure 7-4: An example of one-to-many contextual effects. Underlying positive-influences from Presence of AIDS to every other concept omitted in diagram.

Relative conditioning on the interactions allows us to impose an approximate “order of significance” on some relevant interactions. It is different from actual probabilistic conditioning because we still do not understand how to assign probabilities to concept types; probabilistic distributions are normally in terms of concept instances. Although our scale is somewhat arbitrary and the assignment process strongly subjective, the approximation can still greatly constrain the information for forming a decision model. For example, we could now identify the more common AIDS-related opportunistic-PID and worry about the most common ones first in decision making. Moreover, the assignment of the individual strengths can be guided by whatever information is available, *e.g.*, the relative frequencies of the AIDS-related pulmonary infections in Table 2.1, without requiring complete and accurate probability distributions.

Relative conditioning effects are also inherited in the categorizations similar to the inheritance of interactions. We have yet to work out the details of the indirect relative conditioning effects on the interactions.

7.4 Contextual Effects on Concept Categorizations

Since all concepts are defined in some contexts, the categorizers applicable in the description of a concept are also context-dependent. As shown in Figure 7-1, the

concept is denoted as a pair: (*genus specializer*) in the AKO hierarchy, where *genus* is the general class characterization of the concept, and *specializer* is the specialization dimension, *e.g.*, (leg human). Derivative subclassification classifies a concept with respect to another concept according to the generality of both the *genus* and the *specializer*, *e.g.*, AKO((leg human), (limb animal)). This is very similar to, but not as general as our context-dependent denotation and interpretation of a concept. Moreover, derivative subclassification was not generalized to other categorizations, *e.g.*, facts like PARTOF((phalanx # hand),(finger # human))², which can be expressed in our representation, are not handled in the OWL framework.

The context-dependent categorization specifications are orthogonal in different categorizations. For example, our current framework cannot infer the relationship between (bronchoscopy-in-BAL-test # pulmonary-infection) and (BAL # infection), where PARTOF((bronchoscopy-in-BAL-test # \top),(BAL # \top)) and AKO((pulmonary-infection # \top),(infection # \top)). The interrelations among the different categorizations and their implications on the expressiveness of the framework will be addressed in the future.

²A *phalanx* is a bone of the finger.

Chapter 8

The Medical Knowledge Base

From the representation framework presented so far, we could construct a medical KB in the automated clinical decision making system as shown in Figure 1-1. In Appendix A, we have encoded in a Lisp-like language a small medical KB for supporting formulation of our example decision model shown in Figure 3-1. Although we have not actually implemented this medical KB, much insights could still be gained by interpreting its contents in terms of our representation framework. In this chapter, we examine the contents and the structure of this medical KB. We shall first describe the domain knowledge involved, discuss the structure of the KB abstractly, and then suggest how it can be constructed.

8.1 The Contents

The medical KB contains the domain knowledge required to formulate the decision model shown in Figure 3-1 for the example case presented in Chapter 2. There are five major types of concepts in the medical KB: **physiological-state**, **infection**, **disease**, **test**, and **treatment**. The attributes associated with these concepts, *e.g.*, **existence-status**, **severity** (of disease), **sensitivity** and **specificity** (of test), *etc.*, are the main types of attributive-concepts represented.

8.1.1 Physiological-states

The physiological-states, which are a kind of states, include all the manifestations, *i.e.*, signs, symptoms, laboratory findings, and complications that are applicable in the example case. Some examples are **fever**, **cough**, **neutropenia**, and **bronchospasm**. We also classify concepts for describing the general background such as **sex**, **IV-drug-abuse-history**, and **hemophiliac** as physiological-states because, in accordance with the definition, they are the *status* of part of the world at some time. A physiological-state can have other physiological-states as its own manifestations,

8.1.4 Tests

A diagnostic test is an action in which the existence status of a state or a process is revealed by observing the test results. In a way, the concept `test` is also a role-playing concept because it only makes sense in the context of a state or a process whose existence is to be tested.

We usually associate the following properties with a test: sensitivity, which is a measure of how accurate the test is to confirm an infection or a disease, specificity, which is a measure of how accurate the test is to rule out a disease, complications, mortality rate, which is a measure of how often death results from performing the test, and monetary costs. A test is also commonly classified as a non-invasive or an invasive action; such a description bears significant implications on the above properties.

The medical KB contains the following subclasses of tests: visualization tests, antigen tests, antibody tests, and by-product tests. The individual subclasses of tests include HIV antibody test, which is an antibody test, and all those listed in Table 2.2, which are all visualization tests.

8.1.5 Treatments

A treatment for a disease alleviates the severity of the disease. There are two types of treatments: medical and surgical, but only medical treatments are considered for now. The medical treatments included in the KB are those for the pneumonias.

8.2 The Structure

Abstractly, the medical KB is a huge network comprising numerous multiply-interconnected nodes. Each node is a concept and each link is one of the interactional links, the categorizers, or the context relation. Imagining that each link type has a different color, the context hierarchy, the different categorizations, and the different sets of interactions in each concept description can be easily distinguishable. Figure 8-1 illustrates part of the abstract visualization of the medical KB.

8.3 Constructing The Medical Knowledge Base

As part of a complete knowledge representation system, the medical KB is constructed by defining concepts via the KB-manager. Besides interpreting our representation language and organizing the information in the KB, the KB-manager also functions as an interface which interprets all accessing requests and derives the answers.

The actual organization of the medical KB is decided by the KB-manager. A possible construction process is as follows:

Whenever a concept is defined, the KB-manager derives all the concepts related to it and establish the appropriate relationships among them. For example, the concept **PC-infection** can be defined as:

```
(defconcept PC-INFECTIOIN

  (;;; Contexts
    (#CXT nil))

  (;;; Organizers
    ($AKO protozoal-infection opportunistic-infection))

  (;;; Properties
    (!PATHOGEN pneumocystis-carinii)
    (!ENABLING-FACTOR immunosuppression))

  (;;; Interactions
    ()))
```

The KB-manager will incorporate the knowledge into the KB, derive concepts such as **pathogen-of-PC-infection**, **infection-route-of-PC-infection**, *etc.* from the definition and the organizations, derive more concepts from these concepts, and so forth. If more specific information is provided for the derived concepts, the KB-manager will integrate the information into the KB, making changes and compromises when necessary.

Chapter 9

Supporting Decision Model Formulation

We shall now examine how our medical KB, with the structure as suggested in Section 8.2, supports formulating the decision model shown in Figure 3-1. We shall also briefly discuss how other decision models can be formulated similarly.

9.1 General Queries Format

Given the input case, the planner analyzes the information and then accesses the KB to construct a decision model. To do this, the planner must go through the first four steps in the decision making process mentioned in Chapter 3: background information characterization, clinical context establishment, decision problem formulation, and decision model construction. In each step, a series of queries are processed. Incidentally, only four types of queries of the following formats are required:

- Q1: What are the concepts related to **A** by *<category>*?
- Q2: Does **A** relate to **B** by *<category>*?
- Q3: What are the concepts that directly *<interact>* (with) **A**?
- Q4: Does **A** *<interact>* (with) **B**?

We shall now discuss how the medical KB provides answers for the queries above. In the following discussions:

1. Let \mathcal{C} be the set of all concepts.
2. Let $\Omega = \{AKO, GEN, PARTOF, CONTAIN, EQV, SC\}$ = the set of all categorizers.

Q4: Does A <interact> (with) B?

To find out whether two concepts A and B are involved in an interaction, let $i_0 \in \mathcal{I}$ be the interaction in question.

$$\text{Answer}_{Q4} = \begin{cases} \text{yes} & \text{if } (A, B) \in i_0 \\ \text{no} & \text{otherwise.} \end{cases}$$

An example Q4 query is: Does AIDS influence **opportunistic-pulmonary-infection**? The answer is: yes, and the relationship noted is: **presence-of-AIDS** positively-influences **presence-of-opportunistic-pulmonary-infection**.

In the following sections, we shall show that these four types of queries are adequate for supporting decision model formulation. Throughout the following discussions, we assume the planner knows when to use the right concept in the queries. For example, if we want to find out the causes of a disease, say AIDS, the planner will process the query as: What (are the concepts that) cause **presence-of-existence-of-AIDS**? The answer found, in this case **presence-of-existence-of-HIV-infection**, could also be registered simply as **presence-of-HIV-infection** or **HIV-infection**. We further assume that the planner can derive the indirect relationships of the concepts according to Table 5.1.

9.2 Supporting Background Characterization

Before the input information is characterized by the planner, we assume a front-end processor that “filters” out the relevant concepts for further processing. For example, the concepts **low-grade-fever**, **sputum-examination**, *etc.* as shown in Table 3.1 are identified by the processor and “fed” to the planner. The planner will then create a *case-specific buffer* (CSB), which holds the concepts relevant to the current case.

As mentioned in Section 3.1, the input information has a general pattern; every concept can be characterized as a *history-finding*, a *sign or symptom*, a *laboratory-finding*, a *disease*, an *alternative*, a *complication*, or an *outcome*. The planner has to “understand” all these categories before it can categorize the relevant concepts. This “understanding” can be implemented by an algorithm that matches the concepts with the right categories by asking specific questions. For instance, the planner will employ query Q2 to characterize a **low-grade-fever** as a **sign** or **symptom** by asking questions like:

- Does **low-grade-fever** relate to **history-finding** by specialization?
- Does **low-grade-fever** relate to **sign** by specialization?
- Does **low-grade-fever** relate to **symptom** by specialization? *etc.*

by the user. Given the problem statement, the clinical context, and the CSB, the planner derives the missing information as described in Section 3.3.

Specifically, the planner has to find out:

- I1: What are the causes of AIDS?
- I2: What are the input concepts that are evidence to the presence of AIDS?
- I3: What are the most common pneumonias caused by AIDS-related PIs?
- I4: What are the input concepts that are evidence to the presence of pneumonias?
- I5: What are the outcomes of the pneumonias?
- I6: What are the observable results of the diagnostic-tests of the PIs?
- I7: What are the complications of the diagnostic-tests of the PIs?
- I8: What are the treatments of the PIDs?
- I9: What are the complications of the treatments of the PIDs?

The information is derived as follows:

I1: What are the causes of AIDS?

This question can be formulated in the form of Q3: What are the concepts that directly cause **presence-of-AIDS**? In this case only **presence-of-HIV-infection** is found. The concept **HIV-infection** is then returned to the CSB as the answer.

The influence link found between the two concepts mentioned are derived from the **etiology** property of AIDS. Figure 9-1 shows fragment of the medical KB that supports the derivation.

I2: What are the input concepts that are evidence to the presence of AIDS?

To answer this question, the planner has to find out 1) the concepts that influence the presence of AIDS and its causes, 2) the concepts that the presence of AIDS and its causes influence, and 3) if these concepts are explicitly mentioned in the background information. We assume here the planner remembers all previous results derived, *i.e.*, it knows that **HIV-infection** is the only direct cause of AIDS,

Again with reference to Figure 9-1, the planner will ask a series of additional queries as follows:

1. (Q3) What are the concepts that directly influence presence-of-AIDS or presence-of-HIV-infection? The answers are: risk-factor-of-HIV-infection, pathogen-of-HIV-infection, and infective-route-of-HIV-infection.
2. (Q3) What are the concepts that are directly influenced by presence-of-AIDS and presence-of-HIV-infection? The answers are: wasting-syndrome, immunosuppression, opportunistic-infection, opportunistic-neoplasm, and their specializations.
3. (Q2) Is each concept C, which is a sign or a symptom in the CSB, a kind of, part of, or equivalent to wasting-syndrome, immunosuppression, opportunistic-infection, opportunistic-neoplasm, HIV-infection, risk-factor-of-HIV-infection, pathogen-of-HIV-infection, or infective-route-of-HIV-infection? Here only fever is found to be part of wasting-syndrome and IV-drug-abuse is found to be a risk-factor-of-HIV-infection. Both concepts are marked in the CSB, highlighting their significance for consideration in the decision model.

I3: What are the most common pneumonias caused by opportunistic PIDs with suspected AIDS?

This question can be answered by finding out 1) what are the pneumonias caused by opportunistic infections and 2) if these pneumonias are complications of AIDS. With reference to Figure 9-2, which shows part of the medical KB with the relevant concepts, a set of 4 queries are used to derive the answer:

1. (Q1) What are the concepts that specialize pneumonia? The answers found are all the specializations of pneumonia, including opportunistic-pneumonia.
2. From the answers found so far, (Q4) what are the concepts that are caused by presence-of-opportunistic-infection? The only answer found is: opportunistic-pneumonia.
3. (Q2) Does opportunistic-pneumonia relate to complication-of-AIDS by specialization? The answer is "yes".
4. From the answers found so far, *i.e.*, opportunistic-pneumonia in this case, (Q1) What are the concepts that specialize complication-of-AIDS? The answers found are all the subclasses of opportunistic-pneumonia, including the individual diseases such as PC-pneumonia. The planner will select the most relevant concepts in the context of AIDS based on the "significance degrees" as described in Section 7.3.2.

I4: What are the input concepts that are evidence to the presence of pneumonias?

The answer to this question is derived similarly to that for I2. In essence, the planner uses Q3 and Q2 to find out 1) the concepts that influence the presence of pneumonias and its causes, 2) the concepts that the presence of pneumonias and its causes influence, and 3) if these concepts are explicitly mentioned in the background information. The answers found in this case are the CXR results, the ABG results, and **fever**.

I5: What are the outcomes of the pneumonias?

Query Q1 is used in this case: What are the concepts that specialize outcome-of-pneumonia? The answers found, which are inherited from outcome-of-disease, are: cured, improved, not-improved, worsened, and death. Figure 9-3 shows part of the medical KB that supports the derivation.

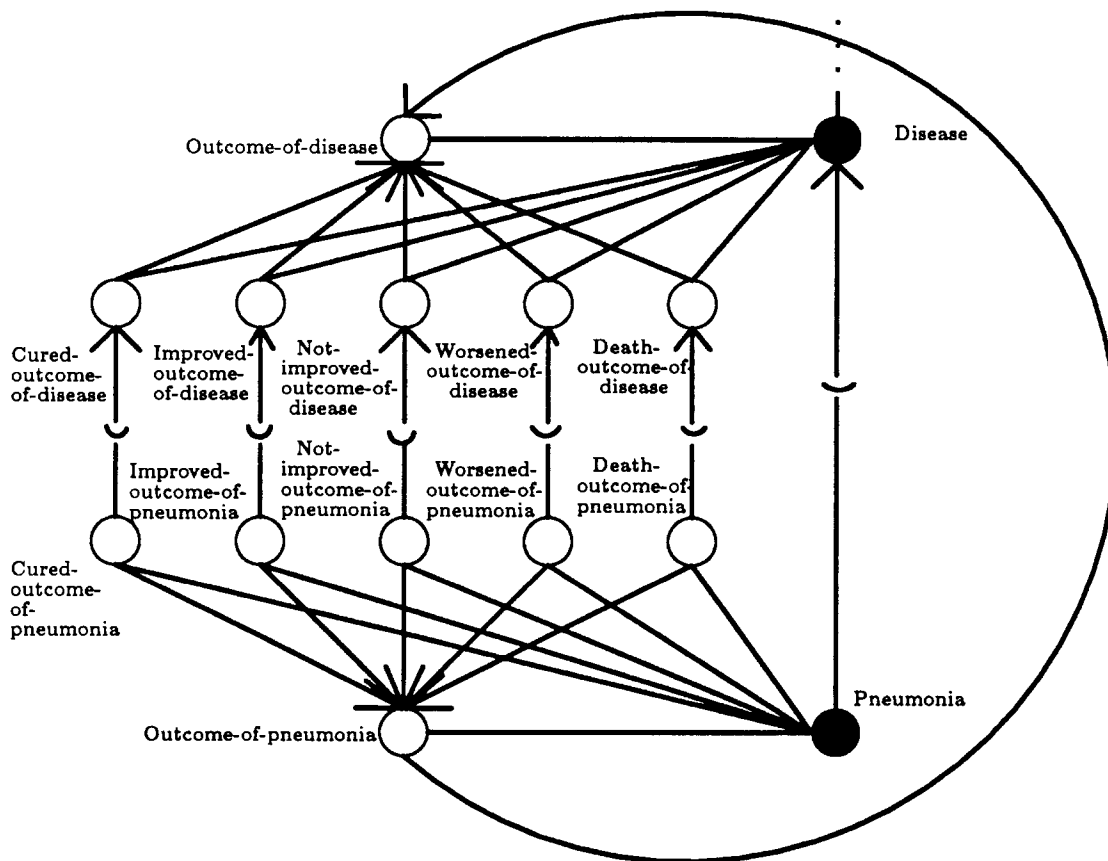


Figure 9-3: Fragment of the medical KB showing the outcomes of the pneumonias.

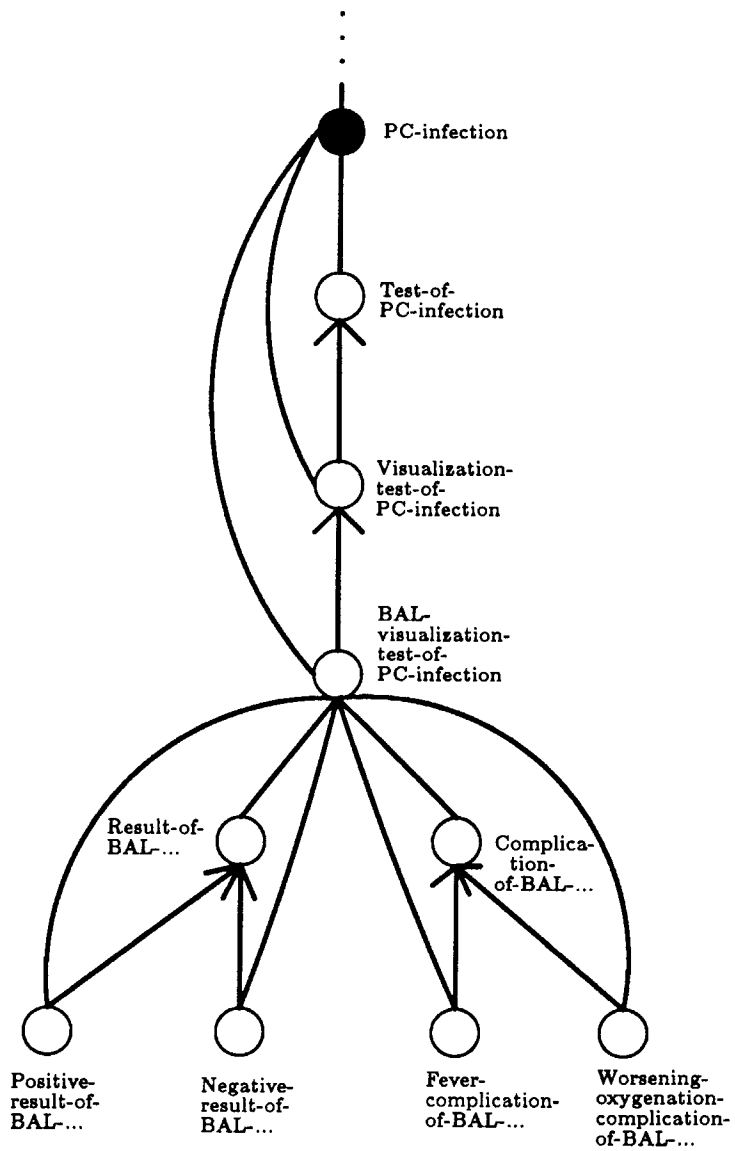


Figure 9-4: Fragment of the medical KB showing the results and the complications of BAL for Pneumocystis-carinii infection.

provided in the background information and solicited when necessary. The process knowledge, which guides the planner to derive the missing information, construct and evaluate the decision model, and provide the explanations, is built into the planner. Our knowledge representation framework captures the domain relations and the probabilistic knowledge structure¹ required for formulating a decision model. The preference knowledge, which ties the evaluation criteria with utility functions, can be solicited from the user; this solicitation can be guided by the knowledge in the decision-analytic KB shown in Figure 1-1.

Similar to the procedure described in Section 9.4, we assume that the planner would derive the relevant decision-analytic knowledge from the corresponding KB. The medical knowledge and the decision-analytic knowledge are then integrated to construct a decision model. Some of the knowledge required for the construction are as follows:

- *Evaluation criteria.* In our example case, which also represents a large class of clinical decision problems, the evaluation criterion is **quality-adjusted-life-expectancy**, which is in turn a function of **monetary-cost**, **morbidity**, and **mortality**. We assume these concepts are derived from the decision-analytic KB.
- *Categorical and probabilistic relations among all the relevant medical concepts and the evaluation criteria.* These are recorded in the CSB as the concepts are derived. For example, the probabilistic relationship between the **dosage-of-drug-treatment** and the **severity-of-disease** is captured in the influence link between the two concepts. The qualitative influence can be instantiated directly in the decision model if a qualitative influence diagram is desired; otherwise, guided by the nature of each influence link, numerical probabilities are solicited from the user.
- *Preference functions of the evaluation criteria.* The relationships between the evaluation criteria and the **utility-value** concept are established. These are again derived from the decision-analytic KB.

9.6 Supporting Decision Model Evaluation

As noted earlier, this step in the decision making process is supported only by the decision-analytic KB.

¹ *vs.* actual numerical probabilities.

Chapter 10

Conclusions

In this work, we have characterized the medical knowledge for supporting decision model formulation in the domain of PIDs with AIDS. Based on our analysis of the clinical decision making process on an example case, we have designed a framework to represent the relevant knowledge. Our design is very much influenced by many existing representation formalisms, including the frame-based languages, semantic-nets, and hierarchical languages such as OWL, KL-ONE, NIKL, and NETL. However, besides being motivated by and hence more suited to the decision-analytic setting, we believe our framework is unique in some ways. We shall now discuss the merits and the limitations of our work, as well as outline a future agenda.

10.1 Achievements

Motivated by the need for dynamically generated decision models, we have categorized the medical knowledge involved for formulating such models into the following:

- Concepts involved in clinical decision making, *e.g.*, physiological-states, infections, diseases, diagnostic-tests, treatments, *etc.*
- Interactions among the concepts, *e.g.*, presence of treatment affects severity of disease, presence of HIV infection causes presence of AIDS, *etc.*
- Categorizations that allow descriptions of the concepts to be inherited from each other, *e.g.*, pulmonary infection is a kind of infection, bronchoscopy is part of a bronchoscopic bronchoalveolar lavage, *etc.*
- Context-dependent descriptions of the concepts, their interactions, and their categorizations.

In addition, we have also characterized the notion of a “context” as something that 1) sets a boundary on the relevant information, 2) allows significance differentiation

clinical problem solving processes such as diagnosis and therapy planning. We need a comprehensive analysis of the expressiveness of the formalism to really understand the contributions of our achievements.

- We have only been able to deal with concept *types* in our framework. Throughout our discussion, we have assumed concept *instances* are handled in the CSB automatically. To have a complete representation framework, however, we need to explicate the similarities and the differences of the concept types and the concept instances.
- Some of the representation issues addressed were not fully explored. In particular, the context (CXT) relation is not formally defined. Moreover, we have recognized the importance of expressing the different degrees of significance in contextual effects in Section 7.3.2, but have not been able to work out all the details of this construct. Another important issue left incomplete is the indirect effects of qualitative synergies on our definitions of influences.
- The assumption of monotonic inheritance in the categorizations is unrealistic. Most existing hierarchical representation systems are deficient because they cannot handle non-monotonic inheritance consistently and correctly. Our framework is subject to the same deficiency.
- By concentrating on the epistemological aspects of knowledge representation, we have totally left out the computational implications in this work. Although it is appropriate to first concentrate on expressiveness, the computational requirements need to be addressed eventually, especially with our ultimate objective of building multi-purpose large KBs.

10.3 Future Directions

Our immediate future agenda is to address the issues raised in the previous section. To study our language design more closely, we also plan to implement a prototype of the medical KB and the KB-manager in our decision making system. In addition, we will attempt to explore the following issues:

10.3.1 Temporal Representation

Except for the temporal component in an interaction, we have worked very little on how to represent temporal information. Temporal considerations, however, are a crucial part of clinical decision making. For instance, the status or severity of a disease after treatment is very much affected by its status or severity before treatment. Often, different stages of a disease in time have different manifestations and implications. Other examples include the temporal effects on treatment and diagnostic tests, the

10.3.3 Integrating Decision Analytic Knowledge

Finally, to complete the picture of the knowledge representation support for our clinical decision making system, we need to address the issue of 1) how to represent the decision analytic knowledge and 2) integrate it with the medical knowledge. Although it might turn out to be very different from what we have explored so far, the approach we would use to examine this topic should be the same: characterize the knowledge, find out what formalisms are available, and adapt useful representational features to build a coherent and semantically clear framework.

10.4 Summary

In this work, we have identified and examined some of the knowledge representation issues in automated clinical decision analysis. As can be concluded from the discussions above, we still have a very long way to go before we could fully and efficiently automate the process. Nevertheless, we believe we have characterized some important features that would guide us towards our objective. The representation framework we have developed is still in a very premature state. Most of the ideas put into our framework design, however, are actually the generalizations of various existing representation formalisms. By experimenting and improving on these ideas, we hope, and are optimistic that our framework will become an effective and comprehensive test-bed for future efforts in this direction.

Appendix A

The Medical Knowledge Base

This appendix encodes the medical KB for formulating the example decision model in a Lisp-like syntax. The structure of the KB is as described in Chapter 8. This encoding is not complete but it contains all the information needed in the example. The implicitly derivable concepts are not included.

In the following encoding, most concepts are defined in the universal context; context-dependent definitions are in the form of A-of-B. The CXT relations in universally defined concepts are not shown. Moreover, the denotation (A B) is shorthand for (A #* B).

```
(defconcept CONCEPT ())

(defconcept ATTRIBUTIVE-CONCEPT ()
  (($AKO concept))
  (!TARGET-CONCEPT))
  ())

;;; Some general concepts.

(defconcept OBJECT ()
  (($AKO generic-concept))
  ()
  ())

(defconcept STATE ()
  (($AKO generic-concept))
  ()
  ())

(defconcept PROCESS ()
  (($AKO generic-concept))
  (!SETTING))
  ())
```

```

    ())

(defconcept HIV-INFECTION ()
  (($AKO lentiviral-infection))
  ((!PATHOGEN human-immunodeficiency-virus))
  ())

(defconcept CMV-INFECTION ()
  (($AKO viral-infection opportunistic-infection))
  ((!PATHOGEN cytomegalovirus)
  (!ENABLING-FACTOR immunodeficiency))
  ())

(defconcept HS-INFECTION ()
  (($AKO viral-infection))
  ((!PATHOGEN herpes-simplex))
  ())

; Protozoal infections.

(defconcept PROTOZOAL-INFECTION ()
  (($AKO infection))
  ((!PATHOGEN protozoan))
  ())

(defconcept PC-INFECTION ()
  (($AKO protozoal-infection)
  ($AKO opportunistic-infection))
  ((!PATHOGEN pneumocystis-carinii)
  (!ENABLING-FACTOR immunodeficiency))
  ())

(defconcept TG-INFECTION ()
  (($AKO protozoal-infection))
  ((!PATHOGEN toxoplasma-gondii))
  ())

; Bacterial infections.

(defconcept BACTERIAL-INFECTION ()
  (($AKO infection))
  ((!PATHOGEN bacterium))
  ())

(defconcept MYCOBACTERIAL-INFECTION ()
  (($AKO bacterial-infection))
  ((!PATHOGEN mycobacterium))
  ())

(defconcept MTB-INFECTION ()
  (($AKO mycobacterial-infection))
  ((!PATHOGEN mycobacterial-tuberculosis))
  ())

```

```

; AIDS.

(defconcept AIDS ()
  (($AKO infectious-disease))
  (!!ETIOLOGY HIV-infection)
  (!!MANIFESTATION wasting-syndrome immunosuppression)
  (!!COMPLICATION opportunistic-infection opportunistic-neoplasm)
  (!!TREATMENT azt))
  ())

; Pulmonary infectious diseases.

(defconcept PULMONARY-DISEASE ()
  (($AKO disease))
  (!!LOCATION lung))
  ())

(defconcept PULMONARY-INFECTIOUS-DISEASE ()
  (($AKO pulmonary-disease)
  ($AKO infectious-disease))
  ())

(defconcept PNEUMONIA ()
  (($AKO pulmonary-infectious-disease))
  (!!MANIFESTATION fever cough))
  ())

(defconcept CMV-PNEUMONIA ()
  (($AKO pneumonia))
  (!!ETIOLOGY CMV-infection))
  ())

(defconcept HS-PNEUMONIA ()
  (($AKO pneumonia))
  (!!ETIOLOGY HS-infection))
  ())

(defconcept PC-PNEUMONIA ()
  (($AKO pneumonia))
  (!!ETIOLOGY PC-infection))
  ())

(defconcept PULMONARY-TOXOPLAMOSIS ()
  (($AKO pneumonia))
  (!!ETIOLOGY TG-infection))
  ())

(defconcept PULMONARY-TUBERCULOSIS ()
  (($AKO pneumonia))
  (!!ETIOLOGY MTB-infection))
  ())

(defconcept MAI-COMPLEX ()

```

```

()

(defconcept COUGH ()
  (($AKO physiological-state))
  ()
  ())

(defconcept RASH ()
  (($AKO physiological-state))
  ()
  ())

(defconcept LEUKOPENIA ()
  (($AKO physiological-state))
  ()
  ())

(defconcept NEUTROPENIA ()
  (($AKO physiological-state))
  ()
  ())

(defconcept ALLERIC-REACTION ()
  (($AKO physiological-state))
  (!ALLERGY)
  (!MANIFESTATION))
  ())

(defconcept GASTROINTESTINAL-UPSET ()
  (($AKO physiological-state))
  ()
  ())

(defconcept BRONCHOSPASM ()
  (($AKO physiological-state))
  (!BRONCHOSCOPE)
  (!MANIFESTATION cough))
  ((primarily-causes (presence bronchoscope) (presence bronchospasm))))

(defconcept HEMOLYTIC-ANEMIA ()
  (($AKO physiological-state))
  ()
  ())

(defconcept HEPATITIS ()
  (($AKO physiological-state))
  (!MANIFESTATION))
  ())

(defconcept LIVER-DYSFUNCTION ()
  (($AKO physiological-state))
  (!MANIFESTATION))
  ())

```

```

()
()

(defconcept VISUALIZATION-TEST ()
  (($AKO test))
  ()
  ())

(defconcept BY-PRODUCT-TEST ()
  (($AKO test))
  ()
  ())

(defconcept ARTRIAL-BLOOD-GAS ()
  (($AKO test))
  (!!NATURE non-invasive))
  ())

(defconcept CHEST-X-RAY ()
  (($AKO test))
  (!!NATURE non-invasive))
  ())

(defconcept SPUTUM-EXAMINATION ()
  (($AKO visualization-test))
  (!!NATURE non-invasive)
  (!SENSITIVITY low)
  (!SPECIFICITY high)
  (!COST low))
  ()))

(defconcept GALLIUM-SCANNING ()
  (($AKO visualization-test))
  (!!NATURE non-invasive)
  (!SENSITIVITY high)
  (!SPECIFICITY low)
  (!COST moderate))
  ())

(defconcept BAL ()
  (($AKO visualization-test))
  (!!NATURE invasive)
  (!SENSITIVITY high)
  (!SPECIFICITY high)
  (!COMPLICATION fever hypoxemia)
  (!MORTALITY low)
  (!COST moderate))
  ())

(defconcept TBBX ()
  (($AKO visualization-test))
  (!!NATURE invasive)
  (!SENSITIVITY high)
  (!SPECIFICITY high)

```

```

    ())

(defconcept AMPHOTERICIN-B-FLUCYTOSINE-TREATMENT ()
  (($AKO drug-treatment))
  (!!DRUG amphotericin-B-flucytosine))
  ())

(defconcept ACYCLOVIR-TREATMENT ()
  (($AKO drug-treatment))
  (!!DRUG acyclovir))
  ())

;;; Some role-playing clinical generic-concepts.

(defconcept PATHOGEN-OF-INFECTION
  (#CXT infection))
  ($EQV microorganism)
  ($AKO pathogen))
  ()
  ((positively-influences (presence pathogen) (presence infection))))

(defconcept INFECTIVE-ROUTE-OF-INFECTION
  (#CXT infection))
  ($EQV route)
  ($AKO infective-route))
  ()
  ((positively-influences (presence infective-route) (presence infection))))

(defconcept RISK-FACTOR-OF-PROCESS
  (#CXT process))
  ($EQV process)
  ($AKO risk-factor))
  ()
  ((positively-influences (presence risk-factor) (presence process))))

(defconcept ENABLING-FACTOR-OF-PROCESS
  (#CXT process))
  ($EQV process)
  ($AKO enabling-factor))
  ()
  ((enables (presence enabling-factor) (presence state))
   (enables (presence enabling-factor) (presence process))))

(defconcept ENABLING-FACTOR-OF-OPPORTUNISTIC-INFECTION
  (#CXT opportunistic-infection))
  ($AKO enabling-factor-of-process)
  ($EQV immunodeficiency))
  ()
  ())

(defconcept ETIOLOGY-OF-PROCESS
  (#CXT process))
  ($EQV process))
  ((primarily-causes (presence etiology) (presence process))))

```

((!MILD)
(!MODERATE)
(!SEVERE))
()

(defconcept LOCATION ()
 ((SAED attributive-concept))
 ()
 ())

(defconcept NATURE ()
 ((SAED attributive-concept))
 ()
 ())

(defconcept NATURE-OF-ACTION
 ((SCTF action))
 ()
 ((!INVASIVE)
 (!NON-INVASIVE))
 ())

Appendix B

Glossary

This appendix contains a list of simplified explanations for the medical terms referenced in the thesis.

Acyclovir An antiviral drug for treatment of infections caused by Herpesvirus.

AIDS Acquired Immune Deficiency Syndrome. A disease caused by HIV infection and resulting in suppression of the body immune response.

Amphotericin-B-Flucytosine An antibiotic-antifungal drug regimen for treating fungal infections.

Antibiotic A substance derived from or produced by some microorganisms to destroy other microorganisms.

Antibody A blood protein produced in the body to destroy or neutralize an antigen; there is a specific type of antibodies produced for each type of antigens.

Antibody test A test to measure the body's immune response to a particular organism by examining the antibodies produced against the organism.

Antigen A foreign and potentially harmful substance in the body; antibodies are produced in the body when antigens are detected.

Antigen test A test to measure the body's immune response to a particular organism by examining the antigens produced by the organism.

Arterial blood gas A test for determining the acidity-alkalinity (pH) and the concentration of oxygen, carbon dioxide, and bicarbonate in the arterial blood.

Bacterium A single-celled microorganism which may be cause harmful effects to the body.

Gallium scanning A visualization test for observing internal organs or presence of organisms in a specimen by injecting the radioactive gallium into the bloodstream or the specimen.

Gastrointestinal upset Discomfort in the digestive system, *i.e.*, the esophagus, stomach, and small and big intestines.

Hemolytic anemia Reduction in the number of red blood cells caused by destruction of these cells.

Hemorrhage Abnormal internal or external discharge of blood.

Herpesvirus A virus that causes latent infections in humans and animals.

Herpes simplex A type of herpesvirus.

HIV Human immunodeficiency virus. The virus responsible for causing AIDS.

Hypoxemia Low concentration of oxygen in the blood.

Immunodeficiency, immunosuppression A decrease in the body's resistance to infections and other diseases; malfunction of the immune system.

Infection Invasion of the body by harmful organisms called pathogens. A infection may lead to inflammation, fever, and sometimes an *infectious-disease*.

Infective route The passage through which a microorganism enters the body to cause infection.

Isoniazid hepatitis Inflammation of the liver caused by the antibacterial drug, isoniazid.

Legionella A pyogenic bacterium that causes pulmonary infections.

Legionellosis A disease caused by infection with legionella.

Leukopenia Reduction in the number of leukocytes, *i.e.*, white blood cells in the blood.

Liver dysfunction Malfunction of the liver.

Manifestation Observable conditions of a disorder.

Mycobacterium A rodlike bacterium.

Mycobacterium avium-intracellulare A mycobacterium which is usually harmless but can cause pulmonary infection in an immunosuppressive host.

Mycobacterium avium-intracellulare complex A disease caused by infection with mycobacterium avium-intracellulare.

Tuberculosis A disease caused by infection with mycobacterium tuberculosis.

Virus A minute particle that is capable of replicating in living cells; the smallest type of infectious agents.

Visualization test A direct visualization test to detect the presence of an organism.

Wasting syndrome A collection of symptoms including fever, night sweat, weight loss, and diarrhea.

Weight loss Decrease in body weight; weight loss without deliberate weight reduction is a symptom of many diseases.

Bibliography

- [1] Ronald J. Brachman and James G. Schmolze. An overview of the KL-ONE knowledge representation system. *Cognitive Science*, 9:171–216, 1985.
- [2] 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.
- [3] John S. Breese. Knowledge representation and inference in intelligent decision systems. Research Report 2, Rockwell International Science Center, Palo Alto, April 1987.
- [4] John S. Breese. Construction of belief and decision networks. Draft manuscript, Rockwell International Science Center, 1989.
- [5] Terence S. Dermody and Daniel R. Lucey. Pulmonary infiltrates in the immunocompromised host. In Mark D. Aronson and Thomas L. Delbanco, editors, *Manual of Clinical Evaluation: Strategies for Cost-Effective Care*, chapter 38, pages 278–288. Little Brown, 1988.
- [6] Jon Doyle and Ramesh S. Patil. Two dogmas of knowledge representation: language restrictions, taxonomic classifications, and the utility of representation services. TM 387b, Massachusetts Institute of Technology, Laboratory for Computer Science, 545 Technology Square, Cambridge, MA, 02139, September 1989.
- [7] H. J. Einhorn and R. M. Hogarth. Behavioral decision theory: Process of judgment and choice. *Annual Review of Psychology*, 32(53), 1981.
- [8] Scott E. Fahlman. *NETL: A System for Representing and Using Real-World Knowledge*. The MIT Press, Cambridge, MA, 1979.
- [9] Harrison Farber. Pulmonary manifestatons. In Howard Libman and Robert A. Witzburg, editors, *Clinical Manual for Care of the Adult Patient with HIV Infection*, pages 93–99. Department of Medicine, Boston City Hospital, Boston, MA, 1990.
- [10] 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.

- [27] C. K. Odgen and I. A. Richards. *The Meaning of Meaning*. Harcourt, Brace, and World, New York, 8th edition, 1946.
- [28] S. G. Pauker and J. P. Kassirer. Medical progress: Decision analysis. *New England Journal of Medicine*, 316:250–258, 1987.
- [29] 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.
- [30] Ross D. Shachter. Evaluating influence diagrams. *Operations Research*, 34:871–882, 1986.
- [31] 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.
- [32] John F. Sowa. *Conceptual Structures: Information Processing in Mind and Machine*. Addison-Wesley, 1984.
- [33] P. Szolovits, L. Hawkinson, and W. A. Martin. An overview of OWL, a language for knowledge representation. TM 86, Massachusetts Institute of Technology, Laboratory for Computer Science, 545 Technology Square, Cambridge, MA, 02139, 1977. Also in Rahmstorf, G., and Ferguson, M., Eds., *Proceedings of the Workshop on Natural Language Interaction with Databases*, International Institute for Applied Systems Analysis, Schloss Laxenburg, Austria, January 10, 1977.
- [34] P. Szolovits, J. P. Kassirer, W. J. Long, A. J. Moskowitz, S. G. Pauker, R. S. Patil, and M. P. Wellman. An artificial intelligence approach to clinical decision making. TM 310, Massachusetts Institute of Technology, Laboratory for Computer Science, 545 Technology Square, Cambridge, MA, 02139, September 1986.
- [35] Amos Tversky and Daniel Kahneman. The framing of decisions and the psychology of choice. *Science*, 211:453–458, 1981.
- [36] Michael P. Wellman. *Formulation of Tradeoffs in Planning Under Uncertainty*. Pitman and Morgan Kaufmann, 1990.
- [37] Michael P. Wellman. Fundamental concepts of qualitative probabilistic networks. *Artificial Intelligence*, 44(3):257–304, 1990.
- [38] Michael Paul Wellman. *Formulation of Tradeoffs in Planning Under Uncertainty*. PhD thesis, Massachusetts Institute of Technology, July 1988.
- [39] William A. Woods. Understanding subsumption and taxonomy: A framework for progress. Technical Report TR-19-90, Center for Research in Computing Technology, Harvard University, Cambridge, MA 02138, 1990.

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

[38] S. G. Parker and J. P. Kassner. *Medical progress: Decision analysis*. Harvard Journal of Medicine, 316:270-272, 1987.

[39] ... A hybrid framework for representing ... knowledge ... of the right ... Cambridge, MA: MIT Press and the MIT ...

[30] Rose D. Spencer. *Evaluating influence diagrams*. ... 1988.

[31] E. A. Sonnenberg. *An intelligent problem solver for lung diseases in AIDS*. Grant proposal (NSF IR901M04938-01A) submitted to the National Library of Medicine, February 1989.

[32] John F. Howal. *Conceptual structures of diagnosis in AIDS and malaria*. Addison-Wesley, 1984.

[33] P. Szolovits, L. Levinson, and W. A. Ruml. *A review of OWL: a language for knowledge representation*. ... Also in Rahmsdorf, G. and ... Proceedings of the Workshop on Natural Language Processing with Databases, International Institute for Applied Systems Analysis, Schloss Lavandru, Austria, January 10, 1977.

[34] P. Szolovits, J. P. Kassner, W. J. Long, S. J. Stoszkowitz, S. G. Parker, R. S. Patti, and M. P. Wellman. *An artificial intelligence approach to clinical decision making*. TM 310, Massachusetts Institute of Technology, Laboratory for Computer Science, 345 Technology Square, Cambridge, MA 02139, September 1988.

[35] Amos Tversky and Daniel Kahneman. *The framing of decisions and the psychology of choice*. Science, 211:453-458, 1981.

[36] Michael P. Wellman. *Formulation of methods in planning*. Ph.D. thesis, MIT, 1990.

[37] Michael P. Wellman. *Fundamental concepts of qualitative probabilistic networks*. Artificial Intelligence, 44(3):287-304, 1990.

[38] Michael Paul Wellman. *Formulation of methods in planning*. Ph.D. thesis, Massachusetts Institute of Technology, July 1987.

[39] William A. Woods. *Understanding subscription and economy: A framework for progress*. Technical Report TR 19-90, Center for Research in Computing Technology, Harvard University, Cambridge, MA, 02138, 1990.

- [11] Lowell B. Hawkinson. The representation of concepts in OWL. In *Proceedings of the Fourth International Joint Conference on Artificial Intelligence*, 1975.
- [12] D. E. Heckerman, E. J. Horvitz, and B. N. Nathwani. Update on the Pathfinder project. In *Symposium on Computer Applications in Medical Care*, pages 203–207. IEEE Computer Society Press, November 1989.
- [13] David E. Heckerman. Probabilistic similarity networks. *Networks*, forthcoming.
- [14] Gary G. Hendrix. Encoding knowledge in partitioned networks. In Nicholas V. Findler, editor, *Associative Networks*. Academic Press, 1979.
- [15] Samuel Holtzman. *Intelligent Decision Systems*. Addison-Wesley, 1989.
- [16] Philip C. Hopewell and John M. Luce. Today's practice of cardiopulmonary medicine. *Chest*, pages 104–112, January 1985.
- [17] J. P. Kassirer and G. A. Gorry. Clinical problem solving: A behavioral analysis. *Annals of Internal Medicine*, 89:245–255, 1978.
- [18] Jerome P. Kassirer and Richard L. Kopelman. The critical role of context in the diagnostic process. *Hospital Practice*, pages 67–76, August 15 1987.
- [19] Jerome P. Kassirer, Alan J. Moskowitz, Joseph Lau, and Stephen G. Pauker. Decision analysis: A progress report. *Annals of Internal Medicine*, 106:275–291, 1987.
- [20] Frank C. Keil. *Concepts, Kinds, and Cognitive Development*. MIT Press, Cambridge, MA, 1989.
- [21] Joseph A. Kovacs and Henry Masur. Opportunistic infections. In Vincent T. DeVita, Jr., Samuel Hellman, and Steven A. Rosenberg, editors, *AIDS: Etiology, Diagnosis, Treatment, and Prevention*, chapter 12, pages 199–225. J. B. Lippincott Co., Philadelphia, second edition, 1988.
- [22] 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.
- [23] Perry L. Miller and Paul R. Fisher. Causal models in medical artificial intelligence. In *Symposium on Computer Applications in Medical Care*, pages 17–22, 1987.
- [24] Marvin Minsky. *The Society of Mind*. Simon and Schuster, New York, 1986.
- [25] M. G. Moser. An overview of NIKL, The New Implementation of KL-ONE. Technical Report 5421, Bolt, Beranek and Newman, Inc., 1983.
- [26] Allen Newell and Herbert A. Simon. *Human Problem Solving*. Prentice-Hall, 1972.

Tuberculosis A disease caused by infection with mycobacterium tuberculosis.
Virus A minute particle that is capable of replicating in living cells; the smallest
type of infectious agent.

Visualization test A direct visualization test to detect the presence of an organism.
Wasting syndrome A collection of symptoms including fever, night sweat, weight
loss, and diarrhea.

Weight loss Decrease in body weight; weight loss without deliberate weight reduc-
tion is a symptom of many diseases.

Mycobacterium tuberculosis A mycobacterium that causes infections.

Neutropenia Reduction in the number of neutrophils, a type of white blood cells, in the blood.

Night sweat Copious sweating during sleep. A symptom for many diseases.

Opportunistic infection An infection that occurs in an immunosuppressive host.

Opportunistic neoplasm A benign or malignant tumor, *i.e.*, abnormal growth, that occurs in an immunosuppressive host.

Pathogen A microorganism, such as a bacterium, that parasitizes a living organism and produces some harmful effects.

Pentamidine A drug for treating pneumocystis-carinii pneumonia.

Pneumocystis carinii A protozoan commonly found in humans. May cause severe pneumonia in an immunosuppressive host.

Pneumonia Inflammation of the lung with pus-filled alveoli (air sacs), resulting in solidness of the lung.

Pneumothorax A collection of air or gas in the chest.

Protozoan The simplest, most primitive type of single-celled animal.

Pulmonary infection Infection of the lungs.

Pyogenic bacterium A bacterium that cause the formation of pus.

Pyrimethamine-sulfadiazine A two-drug regimen for treating toxoplasmosis.

Rash A temporary eruption of spots or reddening of the skin, sometimes accompanied by itching or fever.

Renal dysfunction Malfunction of the kidney.

Sputum examination A visualization test to detect the presence of organisms by examining induced sputum, *i.e.*, saliva mixed with mucus coughed up from the respiratory tract.

Test A procedure done to obtain information for rule out or confirm disorders.

Trimethoprim-sulfamethoxazole A two-drug antibacterial regimen.

Toxoplasma gondii A parasitic protozoan.

Toxoplasmosis A disease caused by infection with toxoplasma gondii.

Treatment A procedure done to alleviate or cure disorders; there are two types of treatment, medical and surgical.

- Bronchoscopic Bronchoalveolar lavage** A test to detect the presence of organisms in which the specimen is obtained by washing out (lavage) of the alveoli (air sacs) in the lung via a hollow tube, called bronchoscope, that extends through the bronchus into the lung.
- Bronchoscopic Transbronchial biopsy** A test to detect the presence of organisms in which the specimen is obtained by cutting a piece of the lung tissue (biopsy) via a tube, called bronchoscope, that extends through the bronchus into the lung.
- Bronchospasm** Abnormal narrowing of the bronchi, *i.e.*, wind pipes, by muscular contraction.
- By-product test** A test to detect the presence of an organism by observing the presence of its by-products, *e.g.*, a toxin.
- Chest X-ray** A routinely performed test to observe upper body organs via electromagnetic waves that can penetrate opaque body mass.
- Chill** A shivering attack accompanied by chattering teeth, pale skin, goose bumps, and a cold feeling. It usually precedes a fever caused by an infection.
- Complication** A disease or an adverse effect that occurs in the presence of another disease, a test, or a treatment.
- Cough** A forceful or sometimes violent exhalation.
- Cryptococcus neoformans** A single-celled yeastlike fungus which can cause infections in humans.
- Cryptococcosis** A disease caused by infection with *cryptococcus neoformans*.
- Cytomegalovirus** A virus commonly found in humans which produces symptoms similar to the common cold. May cause severe effects in immunosuppressive hosts.
- Diarrhea** Increased fluidity, frequency, or volume of bowel movement.
- Disease** A disorder with specific cause, called etiology, and recognizable signs and symptoms. A disease is considered a developing process in this work.
- Enabling factor** A condition that enables an infection or other disorders to occur.
- Erythromycin** An antibiotic used to treat various infections.
- Etiology** The cause of a disease.
- Fever** Rise in body temperature above the normal.
- Fungus** A simple parasitic plant which lacks chlorophyll. May be beneficial or harmful to humans, depending on the species.

(GLEN:)
(STANSON:)
(REVER:)
(C)

(MONTAGI squashed)
(MONTAGI squashed)
(C)
(C)

(MONTAGI squashed)
(MONTAGI squashed)
(C)
(C)

(MONTAGI squashed)
(MONTAGI squashed)
(C)
(REVER:)
(REVER:)
(C)

```

(defconcept MANIFESTATION-OF-PROCESS
  ((#CXT process))
  (($EQV process))
  (!!OBSERVABILITY true))
  ((primarily-causes (process presence) (manifestation presence))))

(defconcept COMPLICATION-OF-CLINICAL-PROCESS
  ((#CXT clinical-process))
  (OR ($EQV physiological-state)
    ($EQV disease))
  ((positively-influences (presence clinical-process) (presence complication))))

(defconcept COMPLICATION-OF-DISEASE
  ((#CXT disease))
  (($AKO complication-of-clinical-process))
  ()
  ())

(defconcept COMPLICATION-OF-TEST
  ((#CXT test))
  (($AKO complication-of-clinical-process))
  ()
  ())

(defconcept COMPLICATION-OF-TREATMENT
  ((#CXT treatment))
  (($AKO complication-of-treatment))
  ()
  ())

(defconcept RESULT-OF-ACTION
  ((#CXT action))
  (($EQV state))
  (!!OBSERVABILITY true)
  (!STATE))
  ())

;;; Some attributive concepts.

(defconcept EXISTENCE ()
  (($AKO attributive-concept))
  (!!PRESENCE)
  (!ABSENCE))
  ())

(defconcept SEVERITY ()
  (($AKO attributive-concept))
  ()
  ())

(defconcept SEVERITY-OF-DISEASE
  ((#CXT disease))
  (($AKO severity))

```

```

      (!COMPLICATION pneumothorax hemorrhage)
      (!MORTALITY moderate)
      (!COST moderate))
    ())

;;; Treatments.

(defconcept TREATMENT ()
  (($AKO clinical-process)
   ($AKO action))
  (!DISEASE))
  ((negatively-influences (presence treatment) (severity disease))))

(defconcept MEDICAL-TREATMENT ()
  (($AKO treatment))
  ()
  ())

(defconcept DRUG-TREATMENT ()
  (($AKO medical-treatment))
  (!DRUG)
  (!DOSAGE)
  (!DURATION)
  (!ADMINISTRATIVE-MODE))
  ((negatively-influences (dosage drug-treatment) (severity disease))
   (negatively-influences (duration drug-treatment) (severity disease))))

(defconcept TMP-SMZ-TREATMENT ()
  (($AKO drug-treatment))
  (!DRUG TMP-SMZ)
  (!DISEASE PC-pneumonia)
  (!ADMINISTRATIVE-MODE oral))
  ())

(defconcept PENTAMIDINE-TREATMENT ()
  (($AKO drug-treatment))
  (!DRUG pentamidine)
  (!DISEASE PC-pneumonia)
  (!ADMINISTRATIVE-MODE oral aerosol))
  ())

(defconcept PYRIMETHAMINE-SULFADIAZINE-TREATMENT ()
  (($AKO drug-treatment))
  (!DRUG pyrimethamine-sulfadiazine))
  ())

(defconcept ANTIBIOTICS-TREATMENT ()
  (($AKO drug-treatment))
  (!DRUG antibiotics))
  ())

(defconcept ERYTHROMYCIN-TREATMENT ()
  (($AKO antibiotics))
  (!DRUG erythromycin))

```

```
(defconcept RENAL-DYSFUNCTION ()
  (($AKO physiological-state))
  (!MANIFESTATION))
  ())
```

```
(defconcept HYPOXEMIA ()
  (($AKO physiological-state))
  (!MANIFESTATION))
  ())
```

```
(defconcept PNEUMOTHORAX ()
  (($AKO physiological-state))
  (!MANIFESTATION))
  ())
```

```
(defconcept HEMORRHAGE ()
  (($AKO physiological-state))
  ())
  ())
```

;;; Syndromes

```
(defconcept SYNDROME ()
  ($CONTAIN physiological-state))
  ())
  ())
```

```
(defconcept WASTING-SYNDROME ()
  ($AKO syndrome)
  ($CONTAIN fever diarrhea night-sweat weight-loss))
  ())
  ())
```

;;; Tests.

```
(defconcept TEST ()
  (($AKO clinical-process action))
  (!RESULT)
  (!CLINICAL-PROCESS)
  (!SENSITIVITY)
  (!SPECIFICITY)
  (!COMPLICATION)
  (!MORTALITY)
  (!COST))
  ((positively-influences (presence clinical-process) (presence result))
  (positively-influences (presence test) (presence result))))
```

```
(defconcept ANTIBODY-TEST ()
  (($AKO test))
  ())
  ())
```

```
(defconcept ANTIGEN-TEST ()
  (($AKO test))
```

```

(($AKO pneumonia))
(!ETIOLOGY MAI-infection))
())

(defconcept PULMONARY-PYOGENIC-BACTERIAL-PNEUMONIA ()
  (($AKO pneumonia))
  (!ETIOLOGY pyogenic-bacterial-infection))
  ())

(defconcept PULMONARY-LEGIONELLOSIS ()
  (($AKO pneumonia))
  (!ETIOLOGY legionella-infection))
  ())

(defconcept FUNGAL-PNEUMONIA ()
  (($AKO pneumonia))
  (!ETIOLOGY fungal-infection))
  ())

(defconcept PULMONARY-CRYPTOCOCCOSIS ()
  (($AKO pneumonia))
  (!ETIOLOGY CM-infection))
  ())

;;; Physiological states.

(defconcept PHYSIOLOGICAL-STATE ()
  (($AKO state))
  ()
  ())

(defconcept FEVER ()
  (($AKO physiological-state))
  ()
  ())

(defconcept DIARRHEA ()
  (($AKO physiological-state))
  ()
  ())

(defconcept NIGHT-SWEAT ()
  (($AKO physiological-state))
  ()
  ())

(defconcept WEIGHT-LOSS ()
  (($AKO physiological-state))
  ()
  ())

(defconcept CHILL ()
  (($AKO physiological-state))
  ()
  ())

```

```
(defconcept MAI-INFECTIOIN ()
  (($AKO mycobacterial-infection)
   ($AKO opportunistic-infection))
  ((!PATHOGEN mycobacterium-avium-intracellulare)
   (!ENABLING-FACTOR immunodeficiency))
  ())
```

```
(defconcept PYOGENIC-BACTERIAL-INFECTIOIN ()
  (($AKO bacterial-infection))
  ((!PATHOGEN pyogenic-bacterium))
  ())
```

```
(defconcept LEGIONELLA-INFECTIOIN ()
  (($AKO pyogenic-bacterial-infection))
  ((!PATHOGEN legionella))
  ())
```

; Fungal infections.

```
(defconcept FUNGAL-INFECTIOIN ()
  (($AKO infection))
  ((!PATHOGEN fungus))
  ())
```

```
(defconcept CN-INFECTIOIN ()
  (($AKO fungal-infection))
  ((!PATHOGEN cryptococcus-neoformans))
  ())
```

;;; Diseases.

```
(defconcept DISEASE ()
  (($AKO clinical-process))
  ((!ETIOLOGY)
   (!SEVERITY)
   (!LOCATION)
   (!MANIFESTATION)
   (!COMPLICATION)
   (!OUTCOME improved unchanged worsened death)
   (!TEST)
   (!TREATMENT))
  ((primarily-causes (presence etiology) (presence disease))
   (primarily-causes (presence disease) (presence manifestation))
   (positively-influences (presence complication) (severity disease))
   (positively-influences (severity disease) (outcome disease))
   (negatively-influences (presence treatment) (severity disease))))
```

; Infectious diseases.

```
(defconcept INFECTIOUS-DISEASE ()
  (($AKO diseases))
  ((!ETIOLOGY infection))
  ())
```



```

(defconcept ACTION ()
  (($AKO process))
  (!!AGENT)
  (!NATURE))
  ())

;;; Concepts in the clinical-setting.

(defconcept CLINICAL-PROCESS ()
  (($AKO process))
  (!!SETTING clinic))
  ())

;;; Infections.

(defconcept INFECTION ()
  (($AKO process))
  (!!PATHOGEN)
  (!INFECTIVE-ROUTE)
  (!RISK-FACTOR)
  (!ENABLING-FACTOR)
  (!LOCATION)
  (!MANIFESTATION inflammation infectious-disease)
  (!TEST))
  ((positively-influences (presence risk-factor) (presence pathogen))
   (positively-influences (presence risk-factor) (presence infective-route))
   (positively-influences (presence risk-enabling) (presence risk-factor))
   (positively-influences (presence risk-factor) (presence infection))
   (positively-influences (presence pathogen) (presence infection))
   (positively-influences (presence infective-route) (presence infection))
   (enables (presence enabling-factor) (presence infection))
   (primarily-causes (presence infection) (presence manifestation))))

(defconcept OPPORTUNISTIC-INFECTION ()
  (($AKO infection))
  (!!ENABLING-FACTOR immunodeficiency)
  (enables (presence immunodeficiency)
           (presence opportunistic-infection)))

; Viral infections.

(defconcept VIRAL-INFECTION()
  (($AKO infection))
  (!!PATHOGEN virus))
  ())

(defconcept RETROVIRAL-INFECTION ()
  (($AKO viral-infection))
  (!!PATHOGEN retrovirus))
  ())

(defconcept LENTIVIRAL-INFECTION ()
  (($AKO retroviral-infection))
  (!!PATHOGEN lentivirus))

```

10.3.3 Integrating Decision Analytic Knowledge

Finally, to complete the picture of the knowledge representation support for our clinical decision making system, we need to address the issues of (1) how to represent the decision analytic knowledge and (2) integrate it with the medical knowledge. Although it might turn out to be very different from what we have explored so far, the approach we would use to examine this topic should be the same: characterize the knowledge, find out what formalisms are available, and adapt useful representational features to build a coherent and semantically clear framework.

10.4 Summary

In this work, we have identified and examined some of the knowledge representation issues in automated clinical decision analysis. As can be concluded from the discussions above, we still have a very long way to go before we could fully and efficiently automate the process. Nevertheless, we believe we have characterized some important features that would guide us towards our objective. The representation framework we have developed is still in a very premature state. Most of the ideas put into the framework design, however, are actually the generalizations of various existing representation formalisms. By experimenting and improving on these ideas, we hope, and are optimistic that our framework will become an effective and comprehensive test-bed for future efforts in this direction.

temporal effects on the observable outcomes of considering “waiting” as an alternative, *etc.*

In [37], Wellman devised a method in QPN to express the two-staged (“before” and “after”) description of a concept, called the *Markov influence*. Besides extending this idea to our definition of an interaction, we hope to develop a more general mechanism to handle all the relevant temporal issues.

10.3.2 Changes of Representation in an Evolving KB

The contents of the KB changes as more information is gathered. These changes can either be in the amount of the information or in the certainty of the information. In the process of building the medical KB, we realized that a KB can be useful even when the information it contains is not complete nor certain. For example, in the beginning we only knew that PCP is most commonly associated with *AIDS* and that **pentamidine** is a **treatment** for PCP, without knowing the exact interactions among them. Such information, however, is sufficient to support very simple queries that might be posed by the planner in formulating a decision model. As the information in the KB improves in amount or certainty, the planner could pose more sophisticated queries, but the overall environment that supports such queries should be stable despite the changes.

Moreover, we believe the amount of *knowledge* in the KB, which could roughly be measured by the number and the complexity of the queries it could answer, changes as the amount and the certainty of the information vary. These changes in knowledge are possible only if the contents of the KB could be transformed, rearranged, summarized, or “learned” when necessary.

From the cognitive psychology viewpoint, this notion of an *evolving* or *adaptive* KB is supported by studies on cognitive conceptual development. These studies show that the acquisition of conceptual knowledge is a gradual but stable process; human beings could manipulate mental concepts before we actually understand what they mean [20].

To develop such an evolving medical KB, we envision having 1) a representation framework which would provide a stable environment for the various changes in the form and the amount of knowledge and 2) a KB-manager which has a learning component that would incorporate the changes when necessary. A lot of challenging issues will be involved in assessing the expressiveness adequacy of our current representation framework, and how it could be extended to support such a KB. The learning component of the KB-manager itself, on the other hand, will open up a whole new direction of research.

of the relevant information, and 3) can be classified hierarchically.

Based on our analysis and a study of related representations that might meet our requirements, we have developed a framework with the following features:

- A concept is an intensional description that reflects the causal/relational structure of the underlying phenomena.
- A spectrum of “definitiveness” for the interactions of the concepts is defined. This allows the behavioral effects of the concepts to be expressed in varying degrees of probabilistic and temporal certainty and precision.
- A concept can be involved in the descriptions of other concepts via interactions; the concept is then named a *property* of those other concepts. This extends the “slots” feature in frame-based languages to capture the “role-chaining” notion introduced in NIKL [25].
- The concepts are categorized, according to their descriptions, into hierarchies or networks. Clear semantics of the categorizations are provided.
- Representation of contextual effects is incorporated into the descriptions of the concepts, their interactions, and hence also their categorizations. This provides a general way to express context-dependent information, conforming to the desiderata for the notion of a “context” mentioned earlier.

All the above features are handled to a certain extent in the various formalisms that influenced our work, but they have not been integrated in a comprehensive way. We believe our framework has provided a basis for developing a semantically clear KB, capable of handling context-dependent representations; this KB is particularly suitable for supporting decision model formulation. To support this claim, we have shown how a small medical KB could be built; we have also shown how the medical KB could support formulating a decision model for an example case. Since most of the representation constructs are domain independent, we believe our framework is also useful in a more general way. The proof for this claim, however, needs to be explored in the future.

10.2 Limitations

As part of an ongoing project, the work reported in this thesis inevitably has a lot of limitations. Based only on the work done so far, some of the more important shortcomings are as follows:

- We have not been able to study more cases in our analysis. Consequently, it is difficult for us to generalize our results. At this stage, we could only hypothesize the usefulness of our formalism, both for different domains and for different

9.7 Formulating Other Decision Models

So far we have described how our representation framework supports the formulation of *one* decision model, in *one* particular domain. We believe, however, that our framework is applicable in a large class of clinical decision problems. We shall now discuss some justifications for this claim.

Most clinical decision problems involve selecting some optimal actions, usually a set of tests, treatments, or simply observations for one or more adverse conditions. The differences in the decision problems usually lie with the assumptions and the constraints being considered; the general domain relations among the diseases, tests, treatments, and other concepts are constant.

For instance, some different decision models for the example case we discussed might involve any or all of the following:

- In addition to the opportunistic PIs, we also consider the opportunistic neoplasms, such as Kaposi's Sarcoma, and other complications of AIDS.
- We do not assume that the opportunistic pneumonias are mutually independent; in other words, we now assume that the presence of one pneumonia, say PCP, will increase the probability that other pneumonias might occur.
- In addition to the tests, we would also like to compare the different treatments for the different pneumonias found to be present.

These decision models can be formulated in a similar manner as described earlier; all the different or additional medical information can still be derived from the medical KB with the same set of accessing queries. The adjustments are made in the information selection procedure of the planner, which is in turn guided by the process knowledge.

Our framework has captured some general domain relations and can handle both invariant and context-dependent effects on these relations. Therefore, we suspect that it is also applicable in many other medical domains, and perhaps even in some other clinical problem solving processes such as diagnosis, therapeutic planning, *etc.* Although we cannot fully explore these hypotheses in this work, a quick supporting argument is as follows: None of the representation structures in our framework is specific to the individual diseases, tests, treatments in the domain of PIDs with AIDS. The designs of all these structures are indeed motivated by the characteristics of the domain, but we have been dealing with general issues like manifestations of diseases, risk-factors of infections, results of tests, complications of treatments, *etc.* We do not claim, however, that our design is even complete for the problem we intended to solve. We will discuss more about the achievements and limitations of our work in the next chapter.

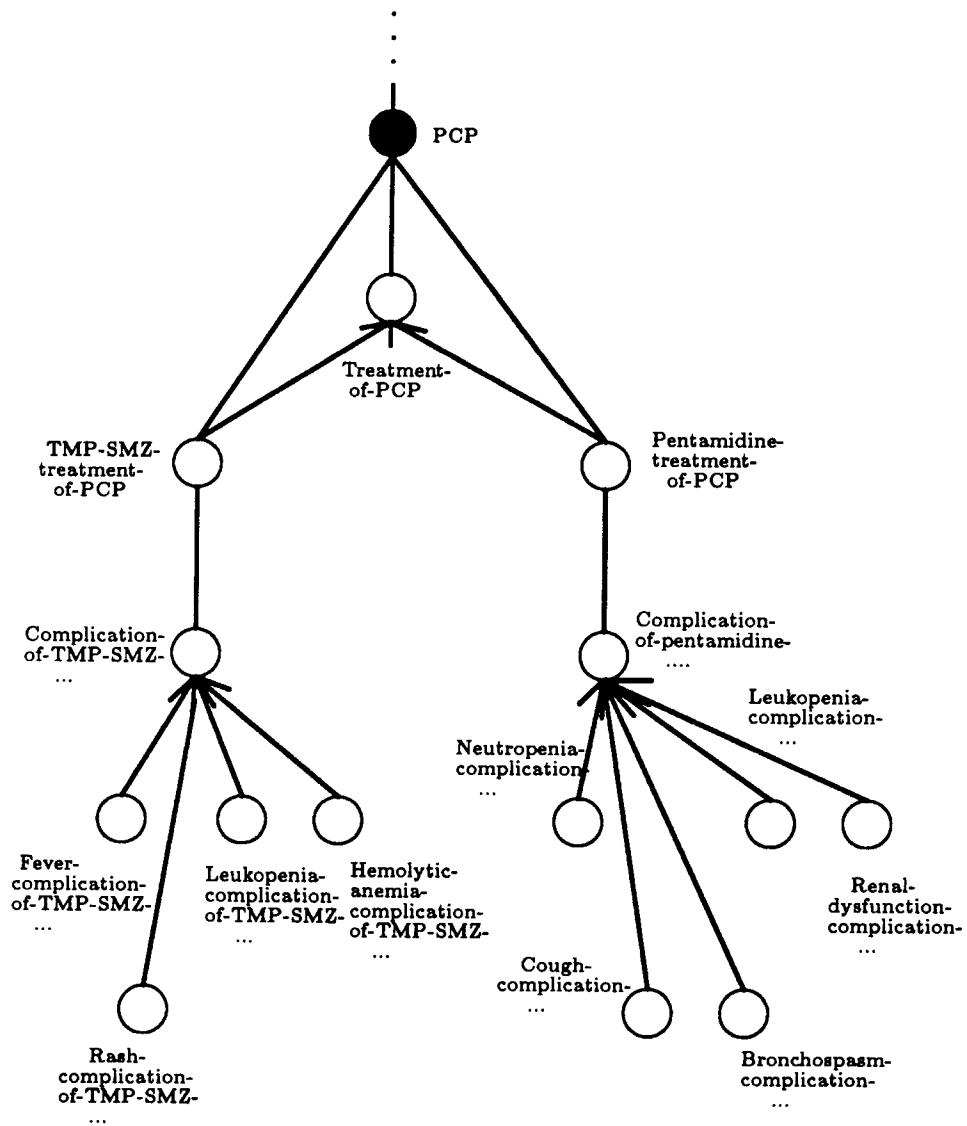


Figure 9-5: Fragment of the medical KB showing the treatments and their complications for PCP.

I6: What are the observable results of the diagnostic-tests for the PIs?

For each cause of the subclasses of opportunistic-pneumonia found in I3, the query Q1 is used to derive all possible results for each test given in the background information, *i.e.*, sputum-examination, gallium-scanning, BAL, and TBBx. In this case, the results of each test are either positive or negative for the corresponding opportunistic-pulmonary-infection. For example, Figure 9-4 shows part of the medical KB that supports the derivation of the query: What are the concepts that specialize the result-of-gallium-scanning-visualization-test-of-PC-infection?

I7: What are the complications of the diagnostic-tests for the PIs?

Similar to answering I6, the query Q1 is used to derived the complications for each test for each cause of the subclasses of opportunistic-pneumonia found in I3. The answers found are as listed in Table 3.2. Most of the complications found are inherent to the test performed. Very rarely does the test for a particular infection manifest specific complications. For example, Figure 9-4 also shows part of the medical KB that supports the derivation of the query: What are the concepts that specialize the complication-of-BAL-visualization-test-of-PC-infection? The complications found in this case are inherited from the generic BAL test.

I8: What are the treatments of the PIDs?

Again the query Q1 is used to find out the treatments for PCP and other PIDs. Figure 9-5 shows the part of the medical KB that supports answering the query: What are the concepts that specialize treatment-of-PCP? The answers found are: TMP-SMZ and pentamidine.

I9: What are the complications of the treatments of the PIDs?

For each treatment R of each PID, this question is answered by the Q1 query. For example, with reference to Figure 9-5, the complications of the treatments for PCP are found with the query: What are the concepts that specialize complication-of-R? The answers are again as listed in Table 3.2.

9.5 Supporting Decision Model Construction

In [15, Chapter 6], Holtzman suggests that there are five kinds of knowledge required to formulate and evaluate a decision model: *domain knowledge*, *preference knowledge*, *probabilistic knowledge*, *user data*, and *process knowledge*. The user data are

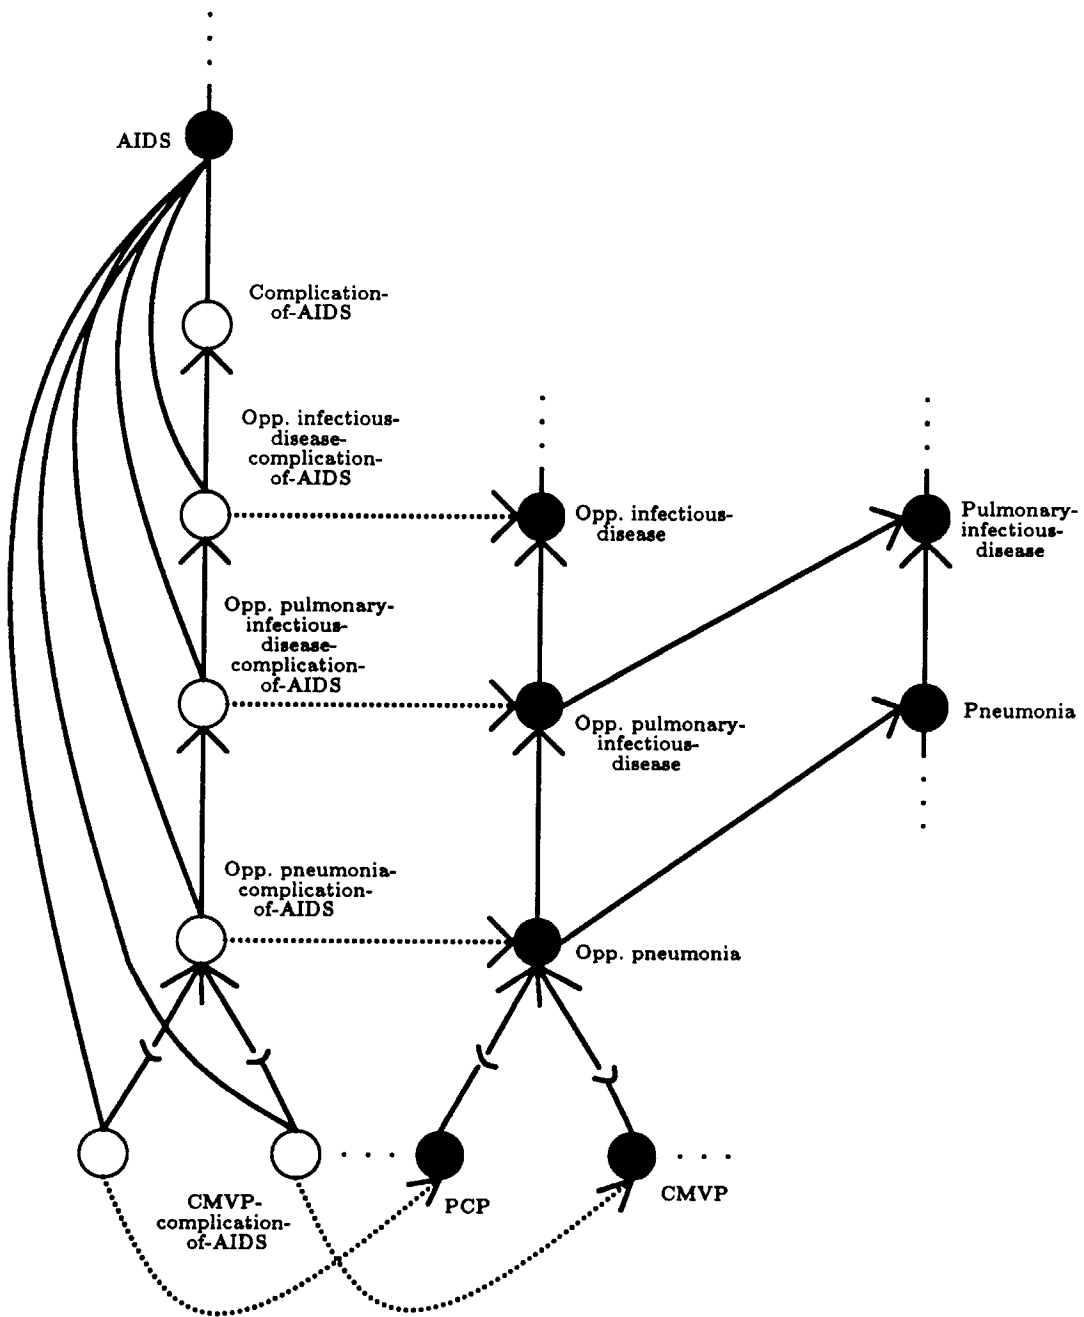


Figure 9-2: Fragment of the medical KB showing opportunistic pneumonias as complications of AIDS.

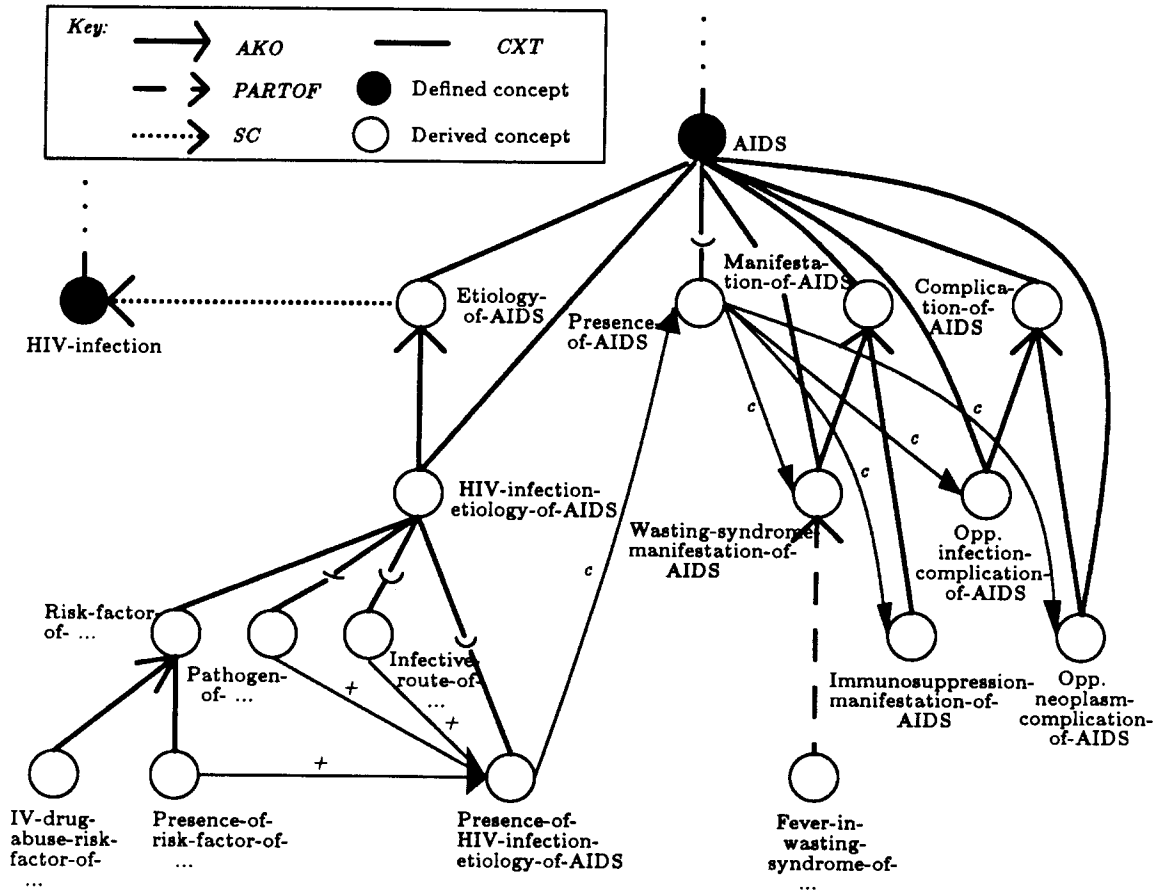


Figure 9-1: Fragment of the medical KB showing the etiology of AIDS.

In generally, query Q2 is adequate for categorizing all the input concepts as illustrated above. The search space can be limited by first asking if a concept is a kind of **state** or a kind of **process** because a history-finding, a sign or a symptom, a laboratory finding, a complication, and an outcome can either be a state or a process, while a disease, a test, or a treatment can only be a process.

After background characterization, the CSB contains the categorized concepts as shown in Table 3.1. Much future work could be saved if the KB-manager can somehow “remember” all the concepts identified and their positions in the medical KB. This could be implemented by marking the concepts in the CSB and tagging them with their positions in the medical KB. In addition, any relationship derived among the concepts could also be registered in the CSB for future reference.

9.3 Supporting Clinical Context Establishment

We assume the planner could derive, perhaps by following a set of criteria, the most specific and yet exhaustive clinical context from the **disease** category in the CSB. Given the characterized background information in the previous step, the planner could identify the clinical context by simply asking the question: What are the diseases suspected in input information? This can be formulated in Q1: What are the concepts (in the input information) that specialize **disease**? In addition, we assume the planner understands the clinical context, *i.e.*, pneumonia caused by PIDs with suspected AIDS, as follows: “Pneumonias *caused by* PIDs” indicates causal relations among the PIDs and the pneumonias, “*with suspected* AIDS” puts the PIDs in the context of AIDS and suggests the unconfirmed status of AIDS.

9.4 Supporting Decision Problem Formulation

The decision problem, in summary, is to decide whether empiric therapy for PCP should be given and how the non-invasive diagnostic tests compare with the invasive ones. A few assumptions are made:

1. The certainty of the presence of AIDS affects the certainty of the presence of the pneumonias.
2. The presence of the pneumonias are independent of each other.
3. Further testing is necessary only if the initial test-results are not conclusive; treatment outcomes do not affect the decision to perform further testing.

Ideally, the planner could formulate a decision problem with as few guidelines from the user as possible. We assume for now that the problem is entirely specified

3. Let $\mathcal{F}_\Omega = \{f_\omega | f_\omega \text{ is a function defined on } \omega, \forall \omega \in \Omega\} = \{ako, gen, partof, contain, eqv, sc\}$ as defined in Chapter 5.
4. Let $\mathcal{I} = \{association, precedence, positive-influence, negative-influence, cause, inhibitor\}$ = the set of all interaction types.
5. $\forall i \in \mathcal{I}$, let $\mathcal{F}_I = \{f_i | f_i \text{ is a function defined on } i\}$
6. $\forall f_i \in \mathcal{F}_I, i \in \mathcal{I}, a, b \in \mathcal{C}, f_i(a) = \{b | (a, b) \vee (b, a) \in i\}$.

Q1: What are the concepts related to A by <categoryzer>?

To find out the concepts related to a concept A in a categorization, let $\omega_0 \in \Omega$ be the categoryzer in question.

$$\text{Answer}_{Q1} = f_{\omega_0}(A).$$

An example of the Q1 query is: What are the concepts that are related to protozoal-infection by specialization? The answers are: Pneumocystis-carinii-infection, Toxoplasma-gondii-infection, etc.

Q2: Does A relate to B by <categoryzer>?

To find out if two concepts A and B are related in a categorization, again let $\omega_0 \in \Omega$ be the categoryzer in question.

$$\text{Answer}_{Q2} = \begin{cases} \text{yes} & \text{if } (A, B) \in \omega_0 \\ \text{no} & \text{otherwise.} \end{cases}$$

An example of the Q2 query is: Does pulmonary-infection relate to infection by specialization? The answer is: yes.

Q3: What are the concepts that directly <interact> (with) A?

To find out the concepts that directly interact with a concept A in an interaction, let $i_0 \in \mathcal{I}$ be the interaction in question.

$$\text{Answer}_{Q3} = f_{i_0}(A).$$

Two examples of the Q3 query are: What are the concepts that cause presence-of-AIDS? What are the concepts that presence-of-pneumonia causes? The answers are: presence-of-HIV-infection and fever, cough, hypoxemia-result-of-ABG, etc., respectively.

The actual organization of the medical KB is decided by the KB-manager. A possible construction process is as follows:

Whenever a concept is defined, the KB-manager derives all the concepts related to it and establish the appropriate relationships among them. For example, the concept PC-infection can be defined as:

```
(concept PC-INFECTION
  (: : Contexts
    (KEY all))
  (: : Organisms
    (SAND PROTONI-INFECTION ORGANISMI-INFECTION))
  (: : Properties
    (PATHEM PROPERTIES-CRITIC)
    (SERIALING-FACTO IMMUNOPROTECT))
  (: : Interactions
    ()))
```

The KB-manager will incorporate the knowledge into the KB, derive concepts such as pathogen-of-PC-infection, infection-route-of-PC-infection, etc. from the definition and the organization, derive more concepts from these concepts, and so forth. If more specific information is provided for the derived concepts, the KB-manager will integrate the information into the KB, making changes and compromises when necessary.




Figure 8-1: A concept as a context.

e.g., cough is a manifestation of bronchospasm.

The actual involvement of a physiological-state as part of the patient's medical history, a sign, a symptom, a laboratory finding, or a complication depends on the context in which these concepts are defined. For example, as mentioned in Chapter 5, the **complication-of-disease** is defined as a structural-copy of physiological-state or disease, in the context of another disease, which fits the description as shown in Figure 7-1.

8.1.2 Infections

An infection is a kind of process in which a harmful microorganism, called a pathogen in this case, invades the body via an infective-route. The infection may cause local inflammation and sometimes lead to more indirect effects, such as a full-blown infectious-disease with distinct signs and symptoms.

There are three other properties commonly associated with an infection: risk-factor, enabling-factor, and test. A risk-factor is a state or a process that will increase the chance of exposure to the pathogen, the infective-route, or other factors that may influence the presence of the infection. An enabling-factor is a state or a process that enables the infection to occur. A test is a diagnostic-test for confirming or ruling out the presence of the pathogen; the different classes of tests for infections will be described in Section 8.1.4.

The major subclasses of infections in the medical KB are: opportunistic infections, which are infections with immunosuppression as enabling-factor, viral infections, protozoal infections, bacterial infections, fungal infections, and pulmonary infections. These subclasses are all specializations of infection, along different specialization dimensions. The individual infection classes represented in the KB, which include HIV infection and all the infectious causes of the diseases listed in Table 2.1, are specializations of one or more of the above subclasses.

8.1.3 Diseases

A disease is a process with the following properties: etiology, severity, manifestation, complication, test, and treatment. All possible causes of the disease are called its etiology. The severity of a disease can usually be described as mild, moderate, or severe. The manifestations of a disease include all the observable signs and symptoms. The complications of a disease are the physiological-states and other diseases that are likely to occur because of the presence of the disease. The test of a disease, like that of an infection, is conducted to confirm or rule out the presence of the disease. The treatment usually reduces the severity of the disease.

The individual subclasses of disease represented are AIDS and the pulmonary infectious-diseases listed in Table 2.1.

concept is denoted as a pair: (genus specifier) in the ARD hierarchy where genus is the general class characterization of the concept, and specifier is the specialization dimension, e.g. (leg human). Derivative subclassification classifies a concept with respect to another concept according to the generality of both the genus and the specifier, e.g. AKO((leg human) (limb animal)). This is very similar to, but not as general as our context-dependent denotation and interpretation of a concept. Moreover, derivative subclassification was not generalized to other categorizations, e.g. facts like PARTOF((palmist # hand) (finger # human)), which can be expressed in our representation, are not handled in the OWL framework.

The context-dependent categorization specifications are orthogonal in different categorizations. For example, our current framework cannot infer the relationship between (bronchoscopy-in-BAL-test # pulmonary-infection) and (BAL # infection), where PARTOF((bronchoscopy-in-BAL-test # T) (BAL # T)) and AKO((pulmonary-infection # T) (infection # T)). The interrelations among the different categorizations and their implications on the expressiveness of the framework will be addressed in the future.

categorizing relations can be explicitly specified in the description of a concept. Otherwise, the position of a concept in a particular categorization can be deduced from the following definition:

Definition 7.2 (Context-dependent Categorization) *Let \mathcal{C} be the set of all concepts. Let Ω be the set of categorizers. Let and $\mathcal{O}_\omega \subseteq \mathcal{C}$ be the set of concepts in a categorization related by categorizer $\omega \in \Omega$. Let ID be the identity relation between two concepts¹ For all $a, b, x, y, x', y' \in \mathcal{C}$, where $a = (x\#y)$ and $b = (x'\#y')$:*

$$\begin{aligned} \omega(a, b) &\iff ID(x, x') \wedge ID(y, y') \\ &\text{or} \\ &ID(x, x') \wedge \omega(y, y') \\ &\text{or} \\ &\omega(x, x') \wedge ID(y, y') \\ &\text{or} \\ &\omega(x, x') \wedge \omega(y, y') \end{aligned}$$

For example, consider the following concepts in the AKO categorization:

- (complication # disease);
- (complication # AIDS);
- (PID-complication # disease); and
- (PID-complication # AIDS).

According to the table above, the following relations are valid:

- AKO((complication # AIDS),(complication # disease));
- AKO((PID-complication # disease),(complication # disease));
- AKO((PID-complication # AIDS),(PID-complication # disease));
- AKO((PID-complication # AIDS),(complication # AIDS)); and
- AKO((PID-complication # AIDS),(complication # disease)).

In the specialization hierarchy, the relations specified in Definition 7.2 are analogous to the idea of *derivative subclassification* in OWL[11, 33] In that system, a

¹This is different from the EQV categorizer.

Blocking Effect

A blocking effect, as opposed to an enabling effect, is a special case of negative synergy in which a second interaction will prevent the first one from taking place. The graphical notation of a blocking effect is shown in Figure 7-3.




Figure 7-3: Blocking effect example: The effect from concept B to concept C blocks the effect from concept A to concept C.

Again, sometimes we wish to derive the relationships among some concepts that are not directly involved in a contextual declaration. For example, given that *enabling-factor-of-opportunistic-infection* enables the causal effect on *opportunistic-infection*, we wish to know if the former has any effect on *opportunistic-pulmonary-infection*, which is a specialization of the latter.

The many-to-one contextual effects are inherited in the categorizations similar to the inheritance of interactions. Indirect contextual effects on the interaction can be derived using the *variable reduction* method in QPN [37, pages 280-283]. We will not go into the details of this method here because it involves instantiating a part of the knowledge-base and manipulating the network. For now we shall assume that a “demon” can derive the appropriate indirect effects when necessary.

7.3.2 One-To-Many Effects

If an concept A interacts in the same way with a few other concepts, say B, C, and D, we may sometimes need to differentiate the *relative* strengths of the interactions i_{ab} , i_{ac} , and i_{ad} . This is similar to the “conditioning” effect in probability, with A as the conditioning variable. The concepts being conditioned are usually sub-contexts of the conditioning concept.

Figure 7-4 shows an example of *relative conditioning*, a one-to-many contextual effect. The interactions are between AIDS and some, assumed exhaustive for now, of its common complications of *opportunistic-PID*. The contextual coverage is indicated by the hyper-edge. An arbitrary scale of 0 to 1, 0 being the weakest and

contextual effects: *synergistic*, *enabling*, and *blocking*. In the graphical notations as shown in Table 7.1, Figure 7-2, and Figure 7-3, the contextual effects are denoted by the hyper-edges. These effects do not affect associations of concepts; since causations and inhibitions are actually influences with additional temporal constraints, we shall only discuss the contextual effects on the positive- and the negative-influences.

Qualitative Synergy

Our definition of qualitative synergy is similar to that of QPN [37, pages 275-286]. Intuitively, two concepts A and B *positive-synergistically* influence a third concept C if their joint influence is greater than the separate, independent influences; the influence is *negative-synergistic* if the joint influence is smaller than the separate, independent influences.

Table 7.1 summarizes the different synergistic effects on the interactions i_{ac} and i_{bc} among three concepts A, B, and C. In the table, the descriptions of the “higher” and “lower” values for continuous concepts can be translated to the “true” and “false” values respectively for binary concepts. The results can easily be generalized to more than two concepts by considering a synergistic effect between every pair of them.

Qualitative synergy describes a rather general class of interactional patterns. Although we do not need to further refine the descriptions for now, there are two special cases of synergy that we need to differentiate: enabling and blocking effects.

Enabling Effect

An enabling effect is a special case of positive synergy, in which a second interaction must be present for the first one to take place. For example, **presence-of-PC-infection** has an enabling effect on the test outcome of **presence-of-BAL**, as shown in Figure 7-2.




Figure 7-2: Enabling effect example.

3. The *CXT* relation is irreflexive, asymmetric, and transitive.

In our framework, all concepts are in the context hierarchy because every concept is defined in terms of a particular context. In other words, each concept is also a context in which other concepts are defined. The descriptions of concepts that are valid in general are in the universal context, \top .

Unlike the multiply-connected categorizations, the context hierarchy is a singly-rooted tree. This is because each concept serves as the context of its derived-concepts, and there is no overlapping among the derived-concepts. All predecessors of a concept in the context hierarchy are called its *super-contexts*; all successors of a concept in the context hierarchy are called its *sub-contexts*.

Besides serving as a focusing mechanism for describing and organizing the concepts, the *CXT* relation keeps track of all the derived-concepts of a concept and their offspring. In other words, each node or concept in the context-hierarchy can be viewed as a *space* in a partitioned network [14]; each space contains all the sub-contexts and their descendants of the corresponding concept. By tracing the links in the context hierarchy, we could infer both the context-invariant and context-dependent interactions among different concepts.

7.2 Contextual Effects on Concept Descriptions

As mentioned in Section 5.3, if we derive a new concept **B-of-A** from the property **B** of a concept **A**, the properties of **B-of-A**, which has the basic-identity **B**, depends very much on **A**; moreover, the description of **B-of-A** is valid only in the context of the description of **A**.

Figure 7-1 shows a simplified example of a derived concept, **complication-of-disease**, from the property **complication** of the concept **disease**. A **complication-of-disease** is equivalent to a (different) **disease** or a **physiological-state** that is affected by a **disease**. In other words, in the context of the description of **disease**, all instances of (another) **disease** or **physiological-state** that fit the description in Figure 7-1, *i.e.*, affected by the **disease**, are **complication-of-disease**. The description is valid only in the defining context and in the descriptions of all concepts to be derived from the current one, *i.e.*, concepts defined in the context of the current description, unless otherwise specified.

Hence, the context relation actually allows different context-dependent descriptions of the same phenomenon. The two different descriptions are independent unless otherwise specified. For instance, the phenomenon represented by **disease** is described as a **complication** under a specific situation, *i.e.*, in the presence of another **disease** which it interacts with. The description of **disease** is referenced in the description of a **complication-of-disease**, as specified by the structural-copy relation in the latter; the description of a **disease**, however, does not include that of

Definition 8.8 (Structural-copy) For all $a, b \in C$, and for $SC \in \Omega$ where $SC \subseteq C \times C$:

$$1. SC \models (a, b) \text{ iff } \exists \alpha \in \Omega \text{ s.t. } \alpha \models a \text{ and } \alpha \models b$$

2. Let $sc: C \rightarrow \Omega$ be a function defined on SC :

$$sc(a) = \{b \mid (a, b) \in SC\}$$

Two major properties are observed for the SC categorizer:

$$1. a \in SC \iff \exists b \text{ s.t. } (a, b) \in SC \text{ or } (b, a) \in SC$$

2. The SC relation is reflexive, asymmetric, and transitive.

Intuitively, the SC relation provides a means for different concepts to share description under different constraints or situations, e.g., in the presence of another disease. These extra constraints or situations are usually captured in the context (CXT) relations of the involving concepts, to be described in Chapter 7.

8.2 Relationships Among Different Categorizations

In the definitions above, monotonic inheritance is assumed in each categorization of concepts. The properties and interactions are inherited independently from the different categorizations in concept description. We now further assume that monotonic inheritance is applicable across all the categorizations in the framework. In other words, all properties and interactions inherited and specified are consistent and complete with respect to the available information in the knowledge base, in each concept description. Consequently, the indirect interactions described in Section 2.4.2 are only relevant in the description of each concept.

6.3 Equivalence

The equivalence EQV relation is a set-equivalence relation between two concepts; the two concepts have an identical set of instances. The different concepts related through the EQV relation are different names or descriptions of the same concept.

Definition 6.5 (Equivalence) For all $a, b \in \mathcal{C}$, and for $EQV \in \Omega$ where $EQV \subseteq \mathcal{C} \times \mathcal{C}$:

1. $EQV \stackrel{def}{=} \{(a, b) | \forall \alpha \in a \iff \alpha \in b\}$.
2. Let $equiv : \mathcal{C} \rightarrow 2^{\mathcal{C}}$ be a function defined on EQV:
 $equiv(a) = \{b | (a, b) \in EQV\}$.

Two major properties are observed for the EQV categorizer:

1. $a \in \mathcal{O}_{EQV} \iff \exists b, (a, b) \in EQV$ or $(b, a) \in EQV$.
2. The EQV relation is reflexive, symmetric, and transitive.

The equivalence categorization is more naturally seen as an undirected network instead of a hierarchy. All concepts in the categorization have the same description and instances, only with different names. For instance, the concepts **therapy** and **treatment** are equivalent.

6.4 Structural-copy (SC)

Another useful categorizer is a variant of the EQV relation, called structural-copy (SC); the SC relation can be viewed as a unidirectional EQV relation. If the relation SC(A,B) is specified in the description of A, the description of B is *visible* in A. In other words, the properties and interactions of B may be used in the description of A, with the appropriate references. For example, SC(complication-of-disease,disease) is specified in the description **complication-of-disease**, which means that the latter may include all the properties and interactions of a **disease**³. A **disease**, however, cannot be described in terms of the properties and interactions of a **complication-of-disease**, because inheritance in an SC categorization is only one-way. We usually reason about a concept with its own properties and interactions; the visible properties and interactions are referenced only when necessary.

³The involvement of the concept **disease** itself in the referenced properties and interactions will now be replaced by **complication-of-disease**

6.2 Decomposition and Aggregation

6.2.1 The Decomposition Relation

The decomposition (**PARTOF**) relation is a set-membership relation between a concept and its *containing* parent concepts; all instances of the decomposed concept are part of some instances of its parent concepts.

Definition 6.3 (Decomposition) For all $a, b \in \mathcal{C}$, and for $PARTOF \in \Omega$ where $PARTOF \subseteq \mathcal{C} \times \mathcal{C}$:

1. $PARTOF \stackrel{def}{=} \{(a, b) | \forall \alpha, \exists \beta, \alpha \in a \implies \alpha \in \beta, \beta = \{\alpha_1, \alpha_2, \dots, \} \in b\}$.
2. Let $partof : \mathcal{C} \rightarrow 2^{\mathcal{C}}$ be a function defined on $PARTOF$:
 $partof(a) = \{b | (a, b) \in PARTOF\}$.

Two major properties are observed for the **PARTOF** categorizer:

1. $a \in \mathcal{O}_{PARTOF} \iff \exists b, (a, b) \in PARTOF$ or $(b, a) \in PARTOF$.
2. The $PARTOF$ relation is irreflexive, asymmetric, and transitive.

A concept can be part of one or more parent concepts and be decomposed along different dimensions. Therefore, the decomposition relation also induces a multiply-connected hierarchy of concepts. For example, **BAL** is a kind of **test** which decomposes into two other testing-processes: **bronchoscopy** and **bronchoalveolar-lavage**. In **BAL**, a bronchoscopy, which is to examine the bronchus by insertion of a tube, is performed before a bronchoalveolar-lavage, which is to wash the bronchoalveolar area with a fluid collected later for analysis.

The properties and interactions of the concepts are upward inheritable in the decomposition categorization. The properties and interactions of a concept include the conjunctive combination of the properties and interactions of its component concepts, *e.g.*, the properties for **BAL** include those of **bronchoscopy** and **bronchoalveolar-lavage**. Unless specified otherwise, however, the inherited descriptions from the component concepts may not exhaustively characterize a concept.

Again, we assume that the inheritance in the decomposition categorization is monotonic; the properties and interactions that a concept inherits from its decomposed components are additive and non-contradictory with respect to the component concepts.

In reality, we realize that monotonic inheritance of properties and interactions in a decomposition categorization rarely occurs. In cases where the same properties are

6.1 Specialization and Generalization

6.1.1 The Specialization Relation

The specialization (AKO)¹ relation is perhaps the most common and the most important relation defined in many representation frameworks. It can be defined as a proper set-inclusion relation between a concept and its subsuming parent concepts; all instances of the specialized concept are instances of its parent concepts.

Definition 6.1 (Specialization) For all $a, b \in \mathcal{C}$, and for $AKO \in \Omega$ where $AKO \subseteq \mathcal{C} \times \mathcal{C}$:

1. $AKO \stackrel{def}{=} \{(a, b) | a \subset b, \text{ i.e., } \forall \alpha, \alpha \in a \implies \alpha \in b\}$.
2. Let $ako : \mathcal{C} \longrightarrow 2^{\mathcal{C}}$ be a function defined on AKO :
 $ako(a) = \{b | (a, b) \in AKO\}$.

Two major properties are observed for the AKO categorizer:

1. $a \in \mathcal{O}_{AKO} \iff \exists b, (a, b) \in AKO \text{ or } (b, a) \in AKO$.
2. The AKO relation is irreflexive, asymmetric, and transitive.

A concept can specialize one or more parent concepts. On the other hand, a concept can be specialized along different dimensions. For instance, **PC-infection** specializes **protozoal-infection** and **opportunistic-infection**, both of which in turn specialize **infection**; **protozoal-infection** specializes **infection** in the type of pathogen, while **opportunistic-infection** specializes **infection** in the host's immunity status. Therefore, the specialization relation induces a multiply-connected hierarchy of concepts.

The properties and interactions of the concepts are downward inheritable in the specialization hierarchy. A specialized concept inherits all the properties and interactions of its parent concepts, unless otherwise specified. For example, an **opportunistic-pulmonary-infection** is both an **opportunistic-infection** and a **pulmonary-infection**, with all the properties and interactions of both parent concepts. The values of the inherited properties and the patterns of the inherited interactions can be modified in the description of the specialized concept.

¹The specialization relation is usually labeled as an **ISA** relation, which does not distinguish between the “subclass-of” and the “instance-of” implications. The specialization relation we are interested in is the “subclass-of” relation; we will reserve the **ISA** label for denoting the “instance-of” relation, to be worked out in future.

5.4.2. Indirect Interactions

The presence of an interaction in a model can have indirect effects on other interactions. For example, if a model contains an interaction between variables A and B, and another interaction between variables B and C, then the interaction between A and C is indirectly affected. This is because the effect of A on C is now dependent on the level of B. In this section, we will discuss how to identify and model these indirect interactions.

\otimes	a	p	+	-	c	i
a	a	a	a	a	a	a
p	a	p	a	a	p	p
+	a	a	+	-	+	-
-	a	a	-	+	-	+
c	a	p	+	-	c	i
i	a	p	-	+	i	c

\oplus	a	p	+	-	c	i
a	a	p	a	a	p	p
p	p	p	p	p	p	p
+	a	p	+	-	p	p
-	a	p	a	-	p	i
c	p	p	c	p	c	p
i	p	p	p	i	p	i

Table 5.1: Indirect Effects of Interactions: The \otimes operator for combining interaction chains and the \oplus operator for combining parallel interactions. The tables are indexed from interaction entries in "row" then "column", and the net interaction is read from their intersection. For example, $\otimes i = +$ and $\oplus c = i$.

The qualitative influences are defined in terms of an ordering criterion called *first-order stochastic dominance* (FSD). It is an ordering on the cumulative probability density functions (CDFs) F_c over the concept C . FSD holds for CDFs F_c and F'_c iff for any given value c_0 of C , the probability of obtaining c_0 or less is smaller for F_c than for F'_c [37].

Definition 5.3 (Influence, Wellman [37]) *Let \mathcal{C} be the set of all concepts in a decision model, $\forall a, b, x \in \mathcal{C}$, where $x \neq a$, and x relates to b via some direct or indirect interactions:*

- *For binary concepts a and b :*
 $a \xrightarrow{+} b \iff Pr(b|a, x) \geq Pr(b|\bar{a}, x)$, and
 $a \xrightarrow{-} b \iff Pr(b|a, x) \leq Pr(b|\bar{a}, x)$.

- *For continuous concepts a and b :*
 $a \xrightarrow{+} b \iff \forall a_1, a_2. a_1 \geq a_2 \Rightarrow F_b(\cdot|a_1, x) \text{ FSD } F_b(\cdot|a_2, x)$, and
 $a \xrightarrow{-} b \iff \forall a_1, a_2. a_1 \leq a_2 \Rightarrow F_b(\cdot|a_1, x) \text{ FSD } F_b(\cdot|a_2, x)$
where $F_b(\cdot|a_i, x)$, $i = 1, 2$ = cumulative probability function (CDF) over b conditioned by a_i .
FSD = First-order stochastic dominance relation.

The arrow-heads in the above notations indicate conditioning directions at the time of encoding. Again, there are no causal or temporal implications.

Some examples of influences are: “**presence-of-risk-factor-of-infection** positively-influences **presence-of-pathogen-of-infection**”, “**presence-of-treatment-of-disease** negatively-influences **severity-of-disease**”, *etc.*

Causal and Inhibitive Links

The *causal/inhibitive links* denote positive-/negative-influences from some concepts to others with known temporal precedence. An example of a cause is: “**presence-of-pulmonary-infection** causes **presence-of-pneumonia**.”

Definition 5.4 (Causality/Inhibition) *Let \mathcal{C} be the set of all concepts in a decision model, $\forall a, b \in \mathcal{C}$:*

$$a_i \xrightarrow{c} b \iff a_i \xrightarrow{+} b \text{ and } a \text{ temporally precedes } b,$$

and

$$a_i \xrightarrow{i} b \iff a_i \xrightarrow{-} b \text{ and } a \text{ temporally precedes } b.$$

Our current interpretation of the causal/inhibitive links do not distinguish between discrete and continuous effects, instantaneous and delayed effects, direct and potential

5.4.1 Interaction Types

A major task in this work is to formally define the types of interactions in a concept description. To balance between intuitive expressiveness and semantic precision, our definitions of interactions are based on an integration of temporal ordering and qualitative probabilistic interpretation.

Each interaction has two components: *temporal precedence*, with “known” or “unknown” as values, and *qualitative probabilistic influence*, with “positive,” “negative,” or “unknown” as values. The interactions can thus be expressed as four types of links in the network interpretation of our framework: *associational* links, which denote probabilistic correlation with an unknown type of influence and unknown temporal precedence; *precedence* links, which denote temporal precedence with unknown type of probabilistic influence; *influential* links, which denote conditional probabilistic dependency; and *causal/inhibitive* links, which denote known temporal precedence in addition to probabilistic dependency.

Our definition of the qualitative probabilistic interaction component is based on Wellman’s work on *qualitative probabilistic influences* [37]. In Wellman’s work, the qualitative probabilistic semantics is formally defined with respect to the QPNs [36], which correspond to instantiations of part of the knowledge base, or *sample spaces* in the probabilistic sense. In other words, all the events involved in a QPN are event instances. Since the concepts described in the knowledge base are concept types, the corresponding semantics for the interactions is not immediately clear.

While much work is still needed to develop a full formal semantics for the interactions in a knowledge base, for now we shall adopt an *operational* semantics for these interactions. The definitions of the interactions presented below describe how they are to be interpreted in the construction and manipulation of a decision model. Moreover, all interactions between two concepts A and B in the knowledge base should be read as: “All instances of A *affect* some instances of B”. This interpretation preserves the property of closure under transitivity [36], and ensures the proper behavior of indirect effects to be described in Section 5.4.2.

Associational Links

The *associational links* indicate probabilistic correlations with unknown type of influences; temporal precedence is also unknown in the associations. Although there may be direct or indirect involvement, an association is just a general claim without implication for the underlying mechanism.

An associational link corresponds to the *?-influence* in the QPN formalism [37], but does not include the *0-influence* correlation in both qualitative and numerical influence diagrams. We believe the 0-correlation is useful for decision model manipulation but unnecessary in a knowledge base. In other words, we impose a closed-world assumption in our definition: all associations are explicitly represented.

hierarchy. For example, the value of the `pathogen` property for `PC-infection` is `pneumocystis-carinii-pathogen`, a specialization of `protozoal-pathogen` which is in turn a specialization of `pathogen`. Thus the values of a property indicate more specific associations between the base-concept and the specializations of the property-concept; the property-concept itself can now be seen as an index to the more specific associations.

A concept may have a set of special properties which represent its attributes, *i.e.*, its inherent qualities and characteristics, *e.g.*, `existence` (of an infection), `severity` (of a disease), `location` (of a tumor), `duration` (of a drug-treatment), *etc.* These attributive properties are descriptions whose corresponding values can only be identified through their dependence on the base-concept. The values of the attributive property `severity of disease`, for example, are `mild`, `moderate`, and `severe`.

The concepts which represent attributes are called *attributive concepts*. An attributive concept is always associated with another concept, called the *target-concept*. The target-concept is the base-concept in which the attributive concept is specified as an attributive property. All other properties in an attributive concept are its valid values for the particular target-concept. For example, as shown in Figure 5-2, the concepts `mild`, `moderate`, and `severe` are the default valid values for describing `severity-of-disease`.

```
(defconcept SEVERITY-OF-DISEASE

  (;;; Contexts
    (#CXT disease))

  (;;; Categorizers
    ($AKO severity))

  (;;; Properties or values
    (!TARGET-CONCEPT disease)
    (!MILD)
    (!MODERATE)
    (!SEVERE))

  (;;; Interactions
    ()))
```

Figure 5-2: A simplified attributive concept example: Severity of disease

Properties as Derived-Concepts and Indices

To represent context-dependent information about the properties of a concept, we can derive a new concept from each property. For example, the

nary relation that specifies the properties and the interactions of, and hence also the categorizers on a concept in accordance with those of another concept. For example, `treatment-of-AIDS` is specified as “a kind of” `treatment-of-disease` because `treatment-of-AIDS` is defined in the context of `AIDS`, and `AIDS` is “a kind of” `disease`. All concepts are described in some contexts; the descriptions that are valid in general are in the *universal* context. The context (CXT) relation allows us to represent context-sensitive information. The partial-ordering imposed by this relation forms a *context-hierarchy*.

Figure 5-1 shows a simplified example of the concept `PC-infection` for `Pneumocystis-carinii-infection`³. Some of the properties and interactions shown in the figure may actually be inherited in the categorizations and may not need to be explicitly encoded.

```
(defconcept PC-INFECTION

  (;;; Contexts
    (#CXT nil))

  (;;; Categorizers
    ($AKO protozoal-infection)
    ($AKO opportunistic-infection))

  (;;; Properties
    (!EXISTENCE-STATUS)
    (!PATHOGEN pneumocystis-carinii)
    (!INFECTIVE-ROUTE)
    (!RISK-FACTOR)
    (!ENABLING-FACTOR immunosuppression)
    (!LOCATION lung)
    (!TEST visualization-test))

  (;;; Interactions
    (affects (presence ** risk-factor) (presence ** pneumocystis-carinii))
    (affects (presence ** risk-factor) (presence ** infective-route))
    (affects (presence ** risk-factor) (presence ** enabling-factor))
    (affects (presence ** risk-factor) (presence ** *))
    (affects (presence ** pneumocystis-carinii) (presence ** *))
    (affects (presence ** infective-route) (presence ** *))
    (affects (presence ** immunosuppression) (presence ** *)))
```

Figure 5-1: A simplified example of a concept: `PC-infection`. The “*” in the description of the interactions specifies the concept being defined.

In the following sections and the next two chapters, we shall look at the definition

³The full representations can be found in Appendix A. The Lisp-like syntax illustrated is an example of the data structures that could implement our formalism.

Proposition 3 The interaction and context-dependent interactions specified among the concepts adequately reflect the phenomena being modeled.

Based on the above propositions, we design a framework for describing the classes of concepts and the types of interactions among them in a context-sensitive manner. Our descriptions of the concept hierarchies are very much influenced by the assumptions and the approaches in NIKI [25] and OWL [11, 33], although our representation is very different. Our descriptions of the interactions among the concepts are based mainly on Wellman's work on the QPN formalism. In the current work, we shall only describe the intended interpretations of the relevant interactions; the interested reader should refer to [27] for the theoretical basis of the QPN framework. Our context-sensitive representation is inspired by the ideas in Hendrix's partitioned network [14].

In the following chapters, we will discuss the overall design of a concept and the different types of interactions among the concepts. We will also describe how to represent context-dependent information in this framework. Due to the large number of issues involved, the descriptions in the following chapters may indeed seem incomplete. We shall leave the intricate details and illustrations to the comprehensive example in Chapter 9.

View 2 *Concepts as relational/causal structures.*

In *The Society of Mind*, Minsky argues that the meaning of something depends on the meanings of all the other things connected to it [24, page 64]. Indeed, theoretical and empirical evidence in philosophy and cognitive psychology has shown that there are three kinds of concepts: *natural kinds*, *artifacts*, and *nominal kinds*, with increasing ease of definability (Figure 4-2) [20]. For instance, a nominal concept “triangle” can be defined as a conjunction of some finite properties, and we can claim that all triangles have these properties; but it is difficult to do so with an artifact concept “chair,” or, as noted earlier, a natural concept “infection” or “disease.” When we reason about an artifact or a natural concept, we usually associate it with some behavioral relationships, or interactions as defined in Chapter 3, with other concepts, *e.g.*, “complication” of a “disease,” “diagnostic-test” of an “infection.” These relationships, which may vary under different circumstances, are not part of the definitional properties of a concept. Therefore, the meaning of a concept includes not just its constituent features but also its context-dependent interactions with other concepts; different contexts have different structuring influences on the meaning [20].

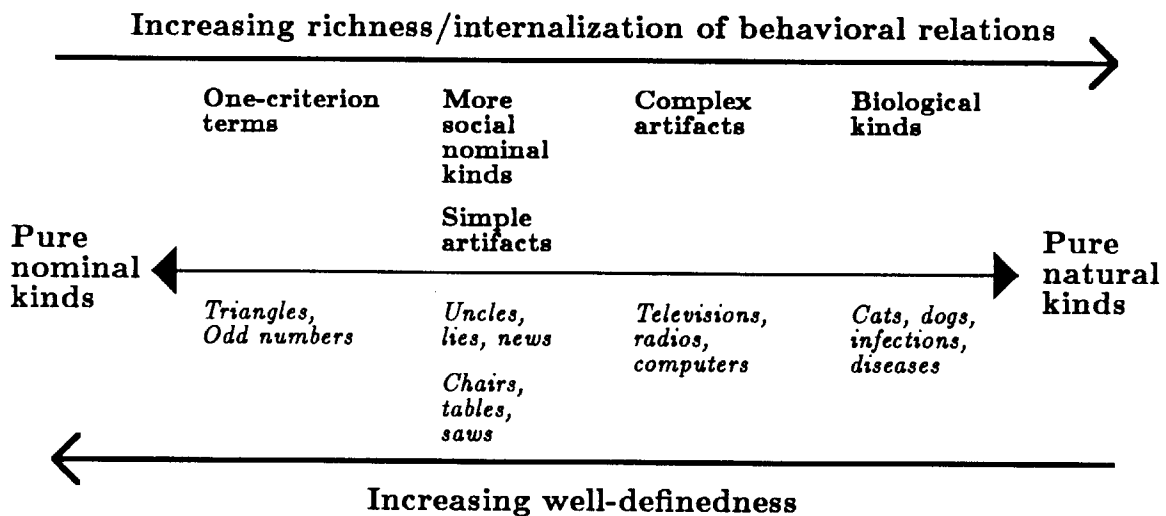


Figure 4-2: The Definability Spectrum

In summary, the basis of our representation design approach consists of the following propositions:

Proposition 1 *Concepts can be viewed as intensional descriptions of some “roughly distinguishable” phenomena. Such intensions may or may not have extensions.*

Proposition 2 *The meaning or structure of a concept may include both its internal characteristics and its interactions with other concepts.*

4.2 Representation Design Approach

Concluding from the brief survey above, we believe our representational requirements call for a hybrid framework that integrates a terminological component, an assertional component, and a network interpretation. The terminological component would capture the categorical knowledge in hierarchical representation of the domain concepts and explicit description of their properties or characteristics. The assertional component would allow expression of the uncertain interactions among the concepts. The network interpretation of the framework would facilitate expression of context-dependent effects on the terminological and assertional components.

We hypothesize that our desired framework could be accomplished by integrating a term subsumption language and a bayesian, or qualitative probabilistic, network formalism into a *partitioned network* [14] semantics. The integration, however, is not straightforward. In particular, the traditional way of distinctly separating the terminological and the assertional components in such a hybrid framework is inadequate for our requirements. This is mainly because term subsumption languages are not very expressive [6]; the rigid logical necessary and sufficient conditions render many important clinical concepts undefinable, *e.g.*, “infection,” “disease,” “treatable disease,” “untreatable disease,” *etc.* Moreover, the description of a concept may actually include some of the non-definitional facts expressed in the assertional component.

Before examining our own representation design, we shall first consider some assumptions in our approach. In the following discussions, all design decisions are primarily motivated by and aimed at representing clinical events or phenomena. In other words, we are not trying to design a *general* knowledge representation framework. Even though we believe some of the ideas presented below are applicable outside the clinical setting, we will not attempt to justify nor evaluate their generality.

We assume real world phenomena can be abstracted into the following types:

- *Objects*: Physical things that are perceivable by the senses, *e.g.*, “drug,” “needle,” “testing instrument,” *etc.*
- *States*: States of being; descriptions of part of the world at particular instances of time, including specific existence status and general perceivable conditions, *e.g.*, “presence of a disease,” “absence of an infection,” “jaundice,” “pallor,” “fever,” *etc.*
- *Processes*: Descriptions of part of the world over a period of time, usually involving a series of changes, *e.g.*, “infection,” “disease,” “diagnostic-test,” “treatment,” *etc.*
- *Attributes*: Inherent qualities or characteristics of the objects, states, processes, or even attributes themselves, *e.g.*, “duration (of a treatment),” “degree (of a fever),” *etc.*

The first order logic-like representations, such as those employed by Breese [2, 4], and Goldman and Charniak [10], have no explicit hierarchical dimensions. In Breese's framework, relations at different levels of detail may co-exist, but they are not treated distinctly [36]. Although Goldman and Charniak's system aims to construct parameterized classes of decision models, the domain entities involved are also not represented hierarchically; instead, the forward-chaining assertional rules are arranged to create multi-level models.

Breese's work captures contextual information in the dependency rules; these rules define the conditions in which the informational dependencies and the probabilities are valid. Since each dependency rule has a complete conditional probabilities matrix, this formalism assumes a finite set of pre-defined contexts. Goldman and Charniak's representation, on the other hand, allows contextual information to be arranged hierarchically to some extent. Instead of specifying a complete conditional probabilities matrix in each rule, some rules contain matrix patterns from which actual probabilities can be derived. This hierarchical arrangement of contexts, however, is again not explicit but induced by the activation of the rules.

Perhaps the approach closest to satisfying our representational requirements is the one adopted by Wellman's SUDO-PLANNER [36]. In this system, domain concepts such as diseases, tests, and treatments are defined in a terminological language, NIKL [25]. The effects of the concepts on each other are expressed in an assertional language based on the QPN formalism. Domain descriptions can be expressed in multiple levels of precision in this framework, thus facilitating decision-modeling in multiple levels of abstraction.

The expressiveness of Wellman's framework, however, is still limited with respect to our needs. In particular, the terminological component is subjected to the shortcomings of most *term subsumption languages*, as we shall discuss later. The purely probabilistic nature of the effects or influences, on the other hand, does not reflect the temporal precedence nor the deterministic causality among the concepts. To formulate a decision model, therefore, the planner has to consider all the concepts that affect or are affected by a particular concept, irrespective of their significance. Moreover, although some contextual effects on the influences are expressible in the *qualitative synergies* defined in QPN, there is no general mechanism for capturing contextual information in the whole framework.

4.1.2 Other Relevant Representation Frameworks

The formalisms most relevant to our work, in addition to the ones developed or used in existing knowledge-based decision systems, are those that incorporate an uncertainty model to a hierarchical representation framework. Most hierarchical representations, including early semantic networks and frame-based languages, are designed to support deductive reasoning based on truth-value interpretations. In other words, only absolute or categorical answers are derivable from these frameworks. Efforts in ac-

treatments, outcomes, and other entities related to the decision problem. Guided by the set of requirements, we shall evaluate some relevant representation formalisms in the next chapter. We shall also propose a representation design approach aimed at fulfilling the above requirements.

context establishment and deriving missing information in decision problem formulation. For example, we want to express facts such as: “PCP is a kind of pneumonia,” “pulmonary infection is a kind of infection,” etc. This type of knowledge should provide us with the power of *abstraction* and *inheritance*. For instance, knowing a general or aggregate class of concepts would allow us to derive the specialized or decomposed classes respectively, and vice versa. The inheritance capability would allow us to specify the generic description for a class of concepts at an appropriate level of abstraction.

3.6.2 Uncertain Knowledge

The uncertain knowledge captures the interactions, *i.e.*, the correlational, influential, or causal relations among the events. These relations are needed to support the derivation of missing information in problem formulation and model construction. For example, we want to express facts such as: “presence of treatment affects severity of disease,” “HIV-infection causes AIDS,” etc. This type of knowledge should allow us to express the varying degrees of probabilistic dependencies among the clinical concepts.

3.6.3 A Contextual Notion

In addition to the categorical and uncertain knowledge, a notion of “context” should be expressible in the knowledge base. This contextual notion has the following properties:

1. It sets a boundary on the relevant categorical and uncertain knowledge, and can be thought of as a focusing mechanism. This enables us to identify what should be considered at different situations. For instance, in the example case, the presence of AIDS would lead us to consider only certain PIDs.
2. It allows differentiation of the relational significance among a set of concepts; the more important information can thus be distinguished from the less important information in different situations. For example, the PIDs most frequently occurred in AIDS patients should be considered first in formulating the decision problem.
3. It is compositional and can be defined hierarchically. In other words, multiple, interacting contexts may coexist and a context can be defined within another context. For example, “AIDS” and “PIDs” combine to form the context of “PIDs with AIDS”; the latter, in turn, is a subcontext of “diseases associated with AIDS.”

In summary, we need a representation that captures the context-sensitive categorical and interactional relations among relevant clinical concepts: diseases, tests,

3.4 Decision Model Construction

In this step, a decision model is constructed from all the concepts identified in Section 3.3. Figure 3-1 shows a QPN, with unlabeled arcs, constructed from the information in Table 3.2.

To construct a decision model, we need to understand its structure, *e.g.*, nodes and links in an influence diagram, and its preference models, *e.g.*, evaluation criteria such as morbidity, mortality, and monetary costs associated with utilities. We also need to correlate the decision-analytic knowledge with the medical knowledge involved. For example, we have to express facts like “presence of a disease will lead to morbidity or mortality.” The temporal constraints on the decision model structure, *i.e.*, what concepts should be considered first and what their consequences are, should be inferable from the interactions of the underlying medical knowledge.

The way we think about the decision-analytic knowledge should not be very different from the medical knowledge. For example, we can think of the entities involved in the decision model structure and preference models as concepts. It may be, however, necessary to distinguish between the two types of knowledge for clarity and modularity.

3.5 Decision Model Evaluation

Upon completion, the decision model is evaluated by some procedure with respect to the evaluation criteria. Here, evaluation of a decision model refers to solving the model with procedures such as folding back of a decision tree, or graph reduction of an influence diagram or QPN. The evaluation criteria we assume are expected monetary cost and quality-adjusted life expectancy, *i.e.*, a measure of time remaining in a patient’s life, taking into account the inconveniences caused by the illness (morbidity). Given a well-formed decision model, only procedural knowledge is needed in this step.

3.6 Representation Requirements

The above analysis shows that to support automated clinical decision analysis, we need to express three types of knowledge: categorical knowledge, uncertain knowledge, and a notion of “context.”

3.6.1 Categorical Knowledge

The categorical knowledge captures the definitional/structural relations of the clinical concepts. These relations are needed to support background characterization,

Table 3.2 shows all the relevant concepts in the decision problem for the example case. All the concepts listed are assumed to be derivable from the input information, within the clinical context.

The concepts in Table 3.2 are derivable only if their relationships to the background information are expressible in the knowledge base. These concepts are only some of the concepts we expect the knowledge base to contain. For example, most of the PIDs listed are individual diseases; we would also expect more general classes of PIDs such as “pulmonary viral infectious disease,” “pulmonary fungal infectious disease,” *etc.*, to be included in the knowledge base as well. The concepts in the knowledge base relate to each other in two ways:

- *Definitional/structural relations:* In Section 3.1, we concluded that in addition to identifying the categorical nature of a concept, we want to infer the properties of a concept from the classes it belongs to. For example, we want to know that PCP occurs in the lungs because it is a kind of pneumonia, and that a pathogen is involved because it is a kind of infectious disease.

Categorically, a concept is usually described as either a specialization or a part of a more general class. For example, BAL is a kind of diagnostic test, while bronchoscopy, which is also a kind of diagnostic test, is part of the BAL procedure. Sometimes, a concept could be an equivalent or a *structural copy* of another concept. For instance, a complication of a disease is another disease or physiological state that occurs in its presence. A complication, when viewed as such, is not a specialization of a disease or physiological state. Instead, it is another name for a disease or a physiological state in the presence of a second disease. The properties of a disease or a physiological state may be used to describe the complication, however, only in the presence of a second disease. Such restricted equivalence relationships are called structural copies. There may be other useful categorical relations defined similarly. These categorical relations should induce multiply-connected hierarchies among all the concepts in the knowledge base, allowing their structural descriptions or definitions to be inferred from each other.

- *Correlational/influential/causal relations:* We use the term *interactions* to denote the correlational, (probabilistic) influential, or causal relations among the concepts in different hierarchies, *e.g.*, presence of pentamidine treatment affects severity of PCP. The interactions may be undirected, unidirectional or mutual, may be of different relational strengths, and may involve two or more concepts at a time. Moreover, the patterns of interactions among certain concepts may change in the presence of some other concepts. For instance, different administrative modes of a treatment for a particular disease may have different efficacies and different side effects: IV pentamidine therapy of PCP is more effective but has more serious complications than aerosol pentamidine.

- *Diseases*: Abnormalities identified or hypotheses to be tested.
- *Alternatives*: The list of available actions.
- *Complications*: The possible interactions or complications among the concepts and actions.
- *Outcomes*: The possible outcomes of the diseases, actions, and complications.

Table 3.1 shows the characterized background information of the example case. The information given is insufficient for formulating a decision model. For example, the different PIDs being considered are not explicitly stated, the empiric therapy for PCP is not defined, and the evaluation criteria are not mentioned. The missing information, which may be related to the medical domain or the decision-analytic methodology, must be derived when necessary.

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

Table 3.1: Characterized Background Information

We can think of each event in Table 3.1 as a *concept*. A concept is 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. Some concepts listed in Table 3.1 could be interpreted as “presence of <concept>,” where *concept* is the event listed, *e.g.*, BAL. In general, when the intended meanings are clear, such intuitive short-hand for describing a property of the concept will be used throughout this report. Other examples of concepts include “severity of a disease,” “dosage of a drug,” “complication of a test,” *etc.* We will define the notion of a concept more formally in later chapters.

To support the characterization, different *categories* or *classes* of concepts should be distinguishable. We should be able to identify the different concepts as discrete entities with distinct properties; concepts with similar properties belong to the same class. For example, sputum examination belongs to the class of non-invasive diagnostic procedures; the latter, in turn, belongs to the class of all actions.

reference to the above case and propose a knowledge representation framework

Diagnosis	Treatment	Essential Complications
PCP	Trimethoprim-Sulfamethoxazole	Rash, fever, leukopenia, hemolytic anemia, Neutropenia, leukopenia, renal dysfunction
Pul. toxoplasmosis	Intravenous Pentamidine	Cough, bronchospasm, Rash, fever, neutropenia, leucopenia, hepatitis
Pul. TB	Aerosol Pentamidine (prophylactic)	Leucopenia, hepatitis, Allergic reactions, renal dysfunction
MAL complex	Pyrimethamine-Sulfadiazine	Gastrointestinal upset
Pyo. bac. pneumonia	Antibiotics regimen	Fever, chills, renal dysfunction, liver dysfunction
Regionellosis	Antibiotics regimen	Fever, chills, renal dysfunction
Pul. cryptococcosis	Erythromycin	
	Amphotericin B-Flucytosine	
Oral fungal infections	Amphotericin B	
CMVP	No standard treatment	
HSP	Acyclovir	

Table 2.3. Treatments for AIDS-related Pulmonary Infections

lating rapid development and evaluation of safer and more cost-effective management strategies [31]. The diversity and uncertainty involved in diagnosis and therapy make PIDs with suspected AIDS a challenging domain for decision analysis. The richness of the underlying knowledge is ideal for examining the representational requirements for supporting automated decision making.

2.2 An Example Case

An example case in the domain of PIDs with suspected AIDS, a simplified version of a case presented in [31], 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. He has no known drug allergies and has never been tested for HIV infection. The initial impression was pneumonia possibly due to opportunistic infection in this patient with suspected AIDS. The problem is to investigate whether or not to employ empiric therapy for PCP, and how non-invasive diagnostic tests such as sputum examination and gallium scanning compare with invasive procedures such as BAL and TBx.

In the following chapter, we will examine the decision-making process with rel-

Diagnosis	Symbol	Relative Frequency
<i>Infectious Causes</i>		
<i>Pneumocystis carinii</i> pneumonia	PCP	85%
Pulmonary toxoplasmosis		<1%
Pulmonary tuberculosis	Pul. TB	4-20%
<i>Mycobacterium avium-intracellulare</i> complex	MAI complex	17-27%
Pyogenic bacterial pneumonia		2-4%
Legionellosis		4%
Pulmonary cryptococcosis		2%
Other fungal infectious diseases		2%
Cytomegalovirus pneumonia	CMVP	17%
Herpes simplex pneumonia	HSP	2%
Miscellaneous		<1%

Table 2.1: Differential Diagnosis of AIDS-related Pulmonary Infectious Diseases

qualitative estimates of sensitivities and specificities, possible complications, and estimated costs [31].

Test	Sensitivity	Specificity	Complications	Mortality	Cost
Sputum examination	Low	High	None	None	Low
Gallium scanning	High	Low	None	None	Moderate
Bronchoscopic bronchoalveolar lavage	High	High	Fever, worsening oxygenation	Low	Moderate
Bronchoscopic transbronchial biopsy	High	High	Pneumothorax, hemorrhage	Moderate	Moderate

Table 2.2: Diagnostic Tests for Pulmonary Infections

Diagnostic tests are done to permit specific treatments for the diseases. Most treatments for the PIDs, even if effective, have serious complications. Given the seriousness of the PIDs and the duration of the testing plans, sometimes empirical treatments, *i.e.*, treatments administered without a precise diagnosis, are necessary; however, since most treatments are very toxic, such decisions have to be carefully weighed. Table 2.3 shows a list of common treatments for each PID and their corresponding complications [21, 31].

The spread of the AIDS epidemic and the severity of the related PIDs are precipi-




Figure 1-1: A knowledge-based clinical decision system.

1.4 Overview of Thesis

This introductory chapter has briefly discussed the nature of knowledge-based decision systems, how they might assist clinical decision making, and how we plan to analyze and formalize the representation requirements. The remainder of this thesis is organized as follows. Chapter 2 describes the medical domain addressed in this work, namely, pulmonary tuberculosis disease with AIDS, and explains why it is interesting from the decision-analytic perspective. An example case is also included for illustration. Chapter 3 details the analysis of the clinical decision analysis process and characterizes the different types of medical knowledge involved. Chapter 4 compares the knowledge requirements found in the analysis with some existing representations. The motivations behind our own representation formalism are then presented. Our representation design is described in Chapter 5 and Chapter 6. Chapter 5 describes the representation of concepts; the interactions among these concepts are discussed in greater detail in Chapter 6. Chapter 7 illustrates the contents of the medical KB. Chapter 8 shows a comprehensive example of how the medical KB supports clinical decision-modeling and postulates how it could be used in general. Finally, Chapter 9 lists the accomplishments, limitations, and future directions of this work.

probabilistic networks (QPN) [37]. Unlike previous efforts, instead of concentrating on the structural components of the decision *model* such as nodes, conditional probabilities, and influences, we focus on the structural features of the decision *problem* such as clinical contexts, classes of evidence and hypotheses, classes of diagnostic procedures, classes of therapeutic interventions, causality, and probabilistic and contextual dependency. By gaining insights into the nature of clinical decisions, this exercise serves as a step toward realizing a uniform representation language for knowledge-based decision systems in medicine [34].

Based on the analysis, the effectiveness of some representation formalisms for supporting automated clinical decision analysis is evaluated. This evaluation motivates a representation design for building knowledge bases with the features mentioned in Section 1.1.3.

1.3 A Hypothetical System

A few assumptions are made about the general system architecture of this work:

1. Our proposed knowledge-based decision system consists of two major components: a *planner* or *decision-maker* which decides what information to process, and a *knowledge base* (KB) which contains all the facts involved in the decision making.
2. The planner constructs a decision model by processing the relevant information and determining the different options for consideration. The relevant information is either available from the initial input or derivable from the KB.
3. The KB can be divided into two parts: a medical KB and a decision-analytic KB. Information from the two KBs are integrated by a *knowledge-base manager* (KB-manager).
4. The KB-manager, besides organizing and integrating the information, serves as an interface from the KB to the planner. The planner can access the KB via some general queries, without knowing the exact structure of the KB and the corresponding inferences.

Figure 1-1 shows the hypothetical system on which we base our discussions.

Our current effort concentrates on analyzing and representing the contents of the medical KB; issues related to other parts of the system will be mentioned without further analysis. Moreover, in this work, the only decision models focused on are the QPNs. Since QPNs are the qualitative variants of influence diagrams, and since each influence diagram can be transformed into a decision tree, we expect our results to be easily generalizable.

et al. [12] and Holtzman's RACHEL [15], the decision-analytic knowledge and the domain knowledge are combined into various *influence diagrams* [30] or their variants. Decision model construction then simply involves selecting, combining, and adapting some of these influence diagrams or "templates" for a specific problem.

However, several recent efforts have advocated that decision models should be regarded as target representations which, in response to a problem, can be dynamically constructed from a knowledge base [36]. Some of these systems, notably Breese's ALTERID and more recent work [3, 4], and Goldman and Charniak's FRAIL3 [10], use first order logic-like languages to encode domain relations and decision-analytic rules. Wellman's SUDO-PLANNER [36], on the other hand, uses a terminological (definitional) language and a simple assertional (factual) language to represent domain relations and some decision-analytic knowledge.

1.1.2 Static versus Constructive Decision-Modeling

We call systems that treat the knowledge bases as decision models "static decision modelers." This approach has two major advantages:

1. The knowledge base usually has a clear and precise semantics, identical or similar to that of the target decision models.
2. Decision models are easy to construct because of the structural and semantical homogeneity of the knowledge base and the target decision models.

On the other hand, the constructive decision-modeling approach has three major advantages over the static approach [4, 36]:

1. *Scalability*: Since no pre-enumeration of anticipated problems is necessary, the knowledge base can support formulating decision models across a wide range of sizes and complexities.
2. *Relevance*: The knowledge in the knowledge base is not committed to be used in any pre-determined manner. Only the relevant knowledge with respect to a specific problem is included in the decision model. This avoids irrelevant information that usually comes along in a fixed model.
3. *Context-sensitivity*: Since most decision problems are context-sensitive, *e.g.*, patient-specific in the clinical setting, a fixed model may not be applicable in all situations. The constructive approach allows both general and specific patterns to be derived from the knowledge base as needed.

Hence, the decision models resulting from the constructive approach are easier to assess, more accurate, and able to address a larger variety of problems. All these features, however, are at the cost of a more complicated model-construction process.

List of Figures

8	4-1	A knowledge-based clinical decision system
17	3-1	A QPN for the Example Case
28	4-1	The Meaning Triangles
28	4-2	The Definability Spectrum
30	2-1	A simplified example of a concept: PC-infection
32	2-2	A simplified attribute concept example: Severity of disease
37	2-3	Part of the interactive model of PC-infection
49	7-1	Complication of disease
50	7-2	Enabling effect example
52	7-3	Inhibiting effect example
53	7-4	An example of one-to-many contextual effects
60	8-1	A concept as a context
68	9-1	Fragment of the medical KB showing the etiology of AIDS
70	9-2	Fragment of the medical KB showing opportunistic pneumonias as complications of AIDS
71	9-3	Fragment of the medical KB showing the outcomes of the pneumonias
73	9-4	Fragment of the medical KB showing the results and the complications of BAL for Pneumocystis carinii infection
74	9-5	Fragment of the medical KB showing the treatments and their complications for PCP

List of Tables

8	2.1 Differential Diagnosis of AIDS-related Pulmonary Infections Diseases	8
8	2.2 Diagnostic Tests for Pulmonary Infections	8
9	2.3 Treatments for AIDS-related Pulmonary Infections	9
12	3.1 Characterized Interpersonal Interaction	12
15	3.2 Concepts Involved in Decision Problem	15
38	5.1 Indirect Effects of Interactions	38
41	7.1 Syncretic effects on interaction	41

B. Glossary Contextual Effects on Concept Categorizations 79

8 The Medical Knowledge Base 87

8.1 The Contents 87

8.1.1 Physiological States 87

8.1.2 Infections 88

8.1.3 Diseases 88

8.1.4 Tests 89

8.1.5 Treatments 89

8.2 The Structure 89

8.3 Constructing The Medical Knowledge Base 89

9 Supporting Decision Model Formulation 93

9.1 General Query Format 93

9.2 Supporting Background Characterization 95

9.3 Supporting Clinical Context Establishment 98

9.4 Supporting Decision Problem Formulation 99

9.5 Supporting Decision Model Construction 75

9.6 Supporting Decision Model Evaluation 75

9.7 Formulating Other Decision Models 76

10 Conclusions 77

10.1 Achievements 77

10.2 Limitations 78

10.3 Future Directions 79

10.3.1 Temporal Representation 79

10.3.2 Changes of Representation in an Evolving KB 80

10.3.3 Integrating Decision Analytic Knowledge 81

10.4 Summary 81

A The Medical Knowledge Base 83

4	Representation Design Approach	21
4.1	Related Work	21
4.1.1	Representations in Existing Knowledge-Based Decision Systems	21
4.1.2	Other Relevant Representation Frameworks	22
4.2	Representation Design Approach	24
5	Representation of Concepts	29
5.1	Overview	29
5.2	Denotation of A Concept	31
5.3	Properties of A Concept	31
5.4	Interactions of Concepts	33
5.4.1	Interaction Types	34
5.4.2	Indirect Interactions	37
6	Categorization of Concepts	39
6.1	Specialization and Generalization	40
6.1.1	The Specialization Relation	40
6.1.2	The Generalization Relation	41
6.2	Decomposition and Aggregation	42
6.2.1	The Decomposition Relation	42
6.2.2	The Aggregation Relation	43
6.3	Equivalence	44
6.4	Structural-copy (SC)	44
6.5	Relationships Among Different Categorizations	45
7	Context-dependent Representation of Concepts	47
7.1	The Context Hierarchy	47
7.2	Contextual Effects on Concept Descriptions	48
7.3	Contextual Effects On Concept Interactions	49
7.3.1	Many-To-One Effects	49
7.3.2	One-To-Many Effects	52

Acknowledgments

I would like to thank:

- Peter Szolovits, my thesis supervisor, for his guidance, encouragement, and support; his understanding and patience during some very difficult times are especially appreciated;
- Mike Wellman, whose excellent work inspired this project, for his advice, encouragement, and reading through numerous drafts of this thesis; his friendship and enthusiasm have helped me throughout the project;
- Frank Sonnenberg, for suggesting and teaching me about the medical domain, and for providing the example case;
- Jon Doyle and Alex Yeh, for commenting on early drafts of this thesis;
- Tom Russ and Tom Wu, for teaching me about various \LaTeX hackery;
- All members of the Clinical Decision Making Group, past and present, for their friendships, and for providing a pleasant and exciting working environment;
- Yang Meng Tan, for his love and support; and
- My parents, for their love and undaunted belief in me.