

# **New types of probabilistic graphical models**

## **Applications to medicine**

Iñigo Bermejo Delgado

Madrid, 2015



Universidad Nacional de Educación a Distancia (UNED)  
Escuela Técnica Superior de Ingeniería Informática  
Departamento de Inteligencia Artificial



Departamento de Inteligencia Artificial  
Escuela Técnica Superior de Ingeniería Informática  
Universidad Nacional de Educación a Distancia

PhD Thesis Dissertation:  
**New types**  
**of probabilistic graphical models.**  
**Applications to medicine**

by

**Iñigo Bermejo Delgado**  
M. Sc. in Artificial Intelligence

Supervised by  
**Prof. Dr. Francisco Javier Díez Vegas**

Madrid, May 2015



# Contents

<b>Abstract</b>	<b>1</b>
<b>Acknowledgments</b>	<b>3</b>
<b>Nomenclature</b>	<b>5</b>
<b>List of Figures</b>	<b>8</b>
<b>List of Tables</b>	<b>9</b>
<b>I. FUNDAMENTALS</b>	<b>11</b>
<b>1. Introduction</b>	<b>13</b>
1.1. Motivation . . . . .	13
1.2. Objectives . . . . .	14
1.3. Methodology . . . . .	15
1.4. A brief note on notation . . . . .	17
1.5. Organization of the thesis . . . . .	18
<b>2. State of the art</b>	<b>19</b>
2.1. An introduction to probabilistic graphical models . . . . .	19
2.2. Asymmetric decision problems . . . . .	27
2.3. Cost-effectiveness analysis . . . . .	30
<b>II. NEW TYPES OF PGM'S</b>	<b>37</b>
<b>3. Decision analysis networks</b>	<b>39</b>
3.1. Basic properties of DANs . . . . .	39
3.2. Equivalent decision tree . . . . .	46
3.3. Recursive decomposition into symmetric DANs . . . . .	50
3.4. Empirical comparison of evaluation algorithms . . . . .	54
<b>4. Markov processes with atemporal decisions</b>	<b>57</b>
4.1. Introduction . . . . .	57
4.2. Definition . . . . .	59

4.3.	Evaluation of MPADs . . . . .	62
4.4.	Overcoming the Markov property . . . . .	66
4.5.	Probabilistic sensitivity analysis with MPADs . . . . .	68
4.6.	Comparison of MPADs with other tools . . . . .	70
4.7.	Conclusions . . . . .	74
<b>5.</b>	<b>Tuning networks</b>	<b>77</b>
5.1.	Basic properties . . . . .	77
5.2.	The tuning model . . . . .	78
5.3.	Construction of a tuning network . . . . .	82
5.4.	Inference: finding a near-optimal strategy . . . . .	84
5.5.	Discussion . . . . .	85
<b>III.</b>	<b>APPLICATIONS TO MEDICINE</b>	<b>87</b>
<b>6.</b>	<b>Cost-effectiveness analysis with MPADs</b>	<b>89</b>
6.1.	Cost-effectiveness of PleurX® vs. talc . . . . .	89
6.2.	Cost-effectiveness of mammography by risk factors for breast cancer .	103
<b>7.</b>	<b>Cochlear implant programming</b>	<b>113</b>
7.1.	Fundamentals . . . . .	113
7.2.	Cochlear implant programming . . . . .	117
7.3.	The Opti-FOX project . . . . .	122
7.4.	Hearing Minds . . . . .	125
7.5.	Conclusions and future work . . . . .	128
<b>IV.</b>	<b>CONCLUSION</b>	<b>129</b>
<b>8.</b>	<b>Conclusions</b>	<b>131</b>
8.1.	Main contributions . . . . .	131
8.2.	Future work . . . . .	134
<b>V.</b>	<b>APPENDICES</b>	<b>137</b>
<b>A.</b>	<b>Functionality implemented in OpenMarkov</b>	<b>139</b>
A.1.	Reorganization in Maven projects . . . . .	139
A.2.	Build and test automatization . . . . .	139
A.3.	Interactive learning of Bayesian networks . . . . .	140
A.4.	The tuning model . . . . .	140
A.5.	Object-oriented probabilistic networks . . . . .	140
A.6.	Likelihood weighting . . . . .	140

A.7. Cost-effectiveness analysis with MPADs . . . . . 141  
A.8. Inference on DTs . . . . . 141  
A.9. Conversion of an ID into an equivalent DT . . . . . 141  
A.10. Conversion of a DAN into an equivalent DT . . . . . 141  
A.11. Evaluation of DANs: recursive decomposition . . . . . 142

**Bibliography** **143**





# Abstract

Probabilistic graphical models (PGMs) play a major role in much of the modern research in reasoning with uncertainty, decision analysis, planning, pattern recognition, and many other areas. Several types of PGMs have been proposed in the last two decades. However, there are some problems for which none of these types are appropriate.

For example, none of the types of PGMs proposed has been widely adopted for representing solving asymmetric decision problems. Decision analysis networks (DANs) have been recently proposed by our research group but they needed efficient evaluation algorithms in order to be applicable to real-world problems. In this thesis we have proposed a new algorithm that evaluates DANs by recursively decomposing them into a set of symmetric DANs, which can then be evaluated with standard algorithms, such as variable elimination or arc reversal. The efficiency of this algorithm matches or even improves that of existing algorithms for other asymmetric representations.

Similarly, existing types of PGMs were not apt as dynamic modeling methods for cost-effectiveness analysis (CEA). The existing dynamic PGMs are burdened by the complexity of their evaluation and can only solve unicriterion problems. Only Markov processes with atemporal decisions (MPADs), a more restricted type of dynamic PGMs also proposed by our research group, are suitable to build complex dynamic models to perform CEA. I have developed new types of potentials and new sensitivity analysis algorithms, with which I have been able to replicate as MPADs several models proposed in the literature and to build two new models for CEA: one for malignant pleural effusion and another one for mammography screening.

Finally, with the help of an expert, we have built a decision-support system for cochlear implant programming (i.e., parameter tuning) based on PGMs. In this thesis we also describe tuning networks, a new type of PGM we developed because existing PGMs were not suitable to model the behavior of systems with a high number of tunable parameters. This decision-support system is now routinely used at a hearing clinic in Antwerp (Belgium) to assist audiologists in the programming of cochlear implants.

All the contributions to PGMs described in this thesis have been implemented in `OpenMarkov`, an open-source software tool developed at the UNED, and are publicly available.



# Acknowledgments

I would like to thank everyone who has helped me throughout my doctoral research.

First and foremost, I want to thank my supervisor, Prof. Francisco Javier Díez, for accepting to supervise my thesis and for his help and guidance ever since, for his availability despite his busy agenda, and for his patience despite my stubbornness. I am very grateful for all the opportunities he has offered me during this time to progress on my research.

I am also very grateful to Dr. Paul Govaerts for his trust and continuous support, for all the motivation, and for pushing me to do my best. I owe you a beer.

I would like to thank the rest of members of CISIAD, for the enlightening discussions, for their help, and for the collaboration in the development of **OpenMarkov**. In particular, I would like to thank Prof. Manuel Luque, with whom I have shared many hours of work and discussion, not only on probabilistic graphical models, but also on football and politics.

I want to thank the colleagues working at the Eargroup and Otoconsult in Antwerp, for their kindness and for taking the extra work to keep testing the prototypes of our decision-support system for fitting cochlear implants.

I am also grateful to Prof. Manfred Jaeger for his guidance during my stay at the Machine Intelligence Group during the spring of 2012. Likewise, I would like to thank Prof. Mark Sculpher, Claire McKenna, and Marta Soares for their collaboration and insightful comments, and to the rest of the members of the Centre for Health Economics of the University of York for the good times during my stay during the spring of 2013.

On a more personal level, I would like to thank my parents and my brother, for their constant love and support. Finally, I am grateful to Clara for encouraging me to embark on this trip and for supporting me through thin and thick.



# Nomenclature

BN	Bayesian network
CEA	Cost-effectiveness analysis
CEAC	Cost-effectiveness acceptability curve
CI	Cochlear implant
DAG	Directed acyclic graph
DAN	Decision analysis network
DCIS	Ductal carcinoma in situ
DES	Discrete event simulation
DT	Decision tree
ECAP	Electrically evoked compound action potential
EVPI	Expected value of perfect information
ICER	Incremental cost-effectiveness ratio
ICF	Intensity-coding function
ICI	Independence of causal influence
ID	Influence diagram
LIMID	Limited-memory influence diagram
MDP	Markov decision process
MPAD	Markov Process with Atemporal Decisions
MPE	Malignant pleural effusions
NMB	Net monetary benefit
PGM	Probabilistic graphical model

POMDP Partially observable MDP

PSA Probabilistic sensitivity analysis

QALY Quality-adjusted life year

SEER Surveillance, Epidemiology and End Results Program

STM State transition model

# List of Figures

2.1.	A Bayesian network for the differential diagnosis of two hypothetical diseases. . . . .	20
2.2.	Influence diagram representing the decision over the test. . . . .	21
2.3.	A family consisting of a child node $Y$ and its parents . . . . .	24
2.4.	Auxiliary structure for the derivation of a noisy ICI model. . . . .	25
2.5.	Cost-effectiveness plane of treatments for GORD relative to maintenance H <sub>2</sub> RA . . . . .	31
2.6.	Cost-effectiveness acceptability curve for different types of hip prostheses Briggs et al. (2004) . . . . .	32
3.1.	A DAN for the diabetes problem. . . . .	40
3.2.	A DAN for the reactor problem. . . . .	41
3.3.	Compatibility table for the link <i>Result of test</i> $\rightarrow$ <i>Build decision</i> . . . . .	41
3.4.	Probabilistic potential of <i>Blood test result</i> , including total restrictions. . . . .	42
3.5.	Decomposition of a DAN representing the diabetes problem during the execution of the conversion into a DT. . . . .	49
3.6.	A decision tree equivalent to the diabetes DAN . . . . .	50
3.7.	Decomposition of the diabetes DAN into symmetric DANs. . . . .	51
3.8.	Computational time for the $n$ -test problem. . . . .	56
4.1.	Transition diagram for Chancellor's model . . . . .	59
4.2.	The first two cycles of an MPAD for the HIV example . . . . .	60
4.3.	Temporal evolution of the variable <i>State</i> for combination therapy . . . . .	65
4.4.	Aggregate value of the variable <i>Life years</i> over time. . . . .	66
4.5.	Transition diagram for Chancellor's model with tunnel states. . . . .	67
4.6.	Transition diagram for Chancellor's model with relapses. . . . .	69
4.7.	Assigning probability distributions to transition probabilities . . . . .	70
4.8.	Cost-effectiveness plane for the HPV vaccine model, evaluated using PSA. . . . .	70
4.9.	Expected value of perfect information for the hip replacement model Briggs et al. (2004). . . . .	71
5.1.	Conditioned interaction in a tuning network. . . . .	84
6.1.	State transition diagram for patients with MPE . . . . .	90
6.2.	MPAD for the CEA of PleurX® vs. talc . . . . .	91

6.3.	Weekly mortality rate and cumulative survival following a Weibull distribution. . . . .	93
6.4.	Conditional probability distribution of <i>Patient state</i> [1]. . . . .	94
6.5.	Utility function for the cost of PleurX supplies, assuming we need two vacuum bottles during the first four weeks and only one thereafter. . . . .	96
6.6.	Temporal evolution of the cohort for both treatments. . . . .	97
6.7.	Evolution of the ICER for patients with an expected survival ranging from 1 to 52 weeks. . . . .	98
6.8.	Evolution of the ICER of PleurX® vs. talc for different effusion resolution rates with talc pleurodesis. . . . .	99
6.9.	Evolution of the ICER of PleurX® vs. talc for different quality of life scores for patients undergoing PleurX® treatment. . . . .	99
6.10.	Evolution of the ICER with, without and 50% of nurse visits. . . . .	100
6.11.	Evolution of the ICER with different drainage frequencies. . . . .	101
6.12.	MPAD for breast cancer screening through mammography. . . . .	105
6.13.	State transition diagram from the cancer diagram. . . . .	106
6.14.	Probability potential for incidence of invasive cancer. . . . .	108
7.1.	Tonotopic distribution of the cochlea. . . . .	114
7.2.	Main parts of a cochlear implant. . . . .	115
7.3.	The result of a loudness scaling test in Audiqueen©. . . . .	120
7.4.	Prototype of the Opti-FOX network. . . . .	123
7.5.	Loudness scaling of patient before and after applying the MAP proposed by the Opti-FOX prototype. . . . .	124
7.6.	Speech audiometry test patient before and after applying the MAP proposed by the Opti-FOX prototype. . . . .	124
7.7.	Screenshot of ICCI showing the effect of a 20 point increase in the T parameter. . . . .	126



# List of Tables

- 3.1. Main features of several methods for representing decision problems. . . 44
- 3.2. Computational time for the different algorithms (ms). . . . . 55
  
- 5.1. Conditional probability table for link  $X_i \rightarrow Y$  in the tuning model. . . 79
- 5.2. Conditional probability table for a link  $X_i \rightarrow Y$  of the *direct* class. . . 80
- 5.3. Conditional probability table for the auxiliary variable  $A1$ . . . . . 84
  
- 6.1. Model parameters . . . . . 93
- 6.2. Results of the cost-effectiveness analysis for the base case. . . . . 97
- 6.3. Cancer stage distributions according to age, mammographic frequency,  
and breast density. . . . . 109
- 6.4. ICER for different mammography frequencies according to age range,  
breast density, family history, and history of biopsy. . . . . 111
  
- 7.1. Features of the ICF and parameters across manufacturers . . . . . 116



**Part I.**  
**FUNDAMENTALS**



# 1. Introduction

*All models are wrong, but some are useful.*

George E. P. Box

## 1.1. Motivation

Probabilistic graphical models (PGMs) have had a significant impact on almost all aspects of artificial intelligence. Soon after their inception in the early 1980s (Howard and Matheson, 1984; Pearl, 1988), they rose to prominence and they now lead much of the modern work in reasoning with uncertainty, decision analysis, planning, pattern recognition, and many other research areas. PGMs have proved that intuitive, qualitative models capturing the probabilistic and informational dependence can help in the construction, understanding, and solution of complex quantitative inference, learning, and decision problems.

Several types of PGMs have been proposed in the literature: Bayesian networks (BNs) (Pearl, 1988), influence diagrams (IDs) (Howard and Matheson, 1984), dynamic BNs (Murphy, 2002), dynamic IDs (Hazen, 2004), limited-memory IDs (LIM-IDs) (Lauritzen and Nilsson, 2001), dynamic LIMIDs (Díez and van Gerven, 2011), factored Markov decision processes (MDPs) (Boutilier et al., 1995) and partially observable MDPs (POMDPs) (Åström, 1965). However, there are some problems for which none of these types of models are appropriate, especially in the field of medicine.

For example, no type of PGM has received wide acceptance to solve asymmetric decision problems, i.e., problems where a total order in the sequence of information is not established or where some variables only make sense only in certain scenarios. Decision analysis networks (DANs) (Díez et al., 2012) had been recently proposed, but they needed efficient evaluation algorithms in order to be accepted by the community.

On the other hand, dynamic PGMs are not used to build dynamic models to perform cost-effectiveness analysis (CEA), because the complexity of most of the algorithms for MDPs, POMDPs, and DLIMIDs grows exponentially with the number of cycles and because they are limited to solve unicriterion problems. For this reason we proposed a new type of PGM called Markov Processes with Atemporal Decisions

(MPAD), a more restricted type of dynamic PGM that are suitable to build complex dynamic models to perform CEA.

Finally, during the construction of a decision-support system for parameter tuning (also called programming or fitting) of cochlear implants, we realized that general purpose decision analysis PGMs were not suitable to model the behavior of systems with a high number of adjustable parameters. The enormous number of possible adjustments makes IDs unsuitable for this task. For this reason, we developed a new type of PGM, called *tuning networks*.

This research has followed the usual approach at the Research Center for Intelligent Decision-Support Systems (CISIAD) of the UNED: the main contributions of this thesis have been motivated by concrete medical problems, but the models, the algorithms, and the software tools we have developed can be applied to many other medical problems and even to other fields. These contributions have been implemented and are now publicly available in the open-source program **OpenMarkov**<sup>1</sup>, a software for the construction and evaluation of PGMs developed at the CISIAD.

## 1.2. Objectives

The general objective of this thesis was to contribute to the development of new types of PGMs and to apply them to solve several medical problems. I have divided this general objective into three more specific objectives:

- Our first objective was to implement the conversion of DANs into decision trees, proposed by Díez et al. (2012). Decision trees have several disadvantages, mainly that their size grows exponentially and that they cannot represent conditional independencies. However, they are still the most usual decision analysis tool and can be very useful to explain the decision process in small problems (Lacave et al., 2007). Nevertheless, this conversion is extremely inefficient. For this reason, a more relevant objective was to develop efficient algorithms for DANs, implement them in **OpenMarkov**, and evaluate them.
- The second objective was to perform cost-effectiveness analyses for several medical problems using MPADs. A necessary sub-goal was to complete and extend the functionalities to perform CEA with MPADs in **OpenMarkov**. Another sub-goal was to reproduce with MPADs some of the models published in the literature. The objective of building new MPADs for CEA was concreted in two particular problems: analyzing the cost-effectiveness of different frequencies for mamography and that of two treatments for malignant pleural effusion.
- The third goal was to build a decision-support system for fitting cochlear implants. This objective led to the emergence of new

---

<sup>1</sup>[www.openmarkov.com](http://www.openmarkov.com)

## 1.3. Methodology

### 1.3.1. Algorithms for DANs

The development of DANs had begun at the CISIAD, before the start of my doctoral research. As a matter of fact, I started working on them after they were first presented by Díez et al. (2012) at the Sixth Workshop on Probabilistic Graphical Models, in Granada (Spain). I was first assigned the task of designing and implementing the conversion of DANs into equivalent decision trees in `OpenMarkov`, from the instructions in (Díez et al., 2012). This conversion was meant to serve as a definition of the syntax of DANs rather than as an evaluation algorithm, for the well known inefficiency of the decision tree evaluation algorithm (Raiffa, 1968). As the first step, I got acquainted with the literature on the representation of asymmetric decision problems. During the implementation of the conversion algorithm, we found a few loose ends in the definition of DANs. After discussing these with the authors of the original proposal, some of the features of DANs changed and so did the conversion algorithm.

The main objective was to develop an algorithm that would match the efficiency of the algorithms proposed in the literature for other asymmetric decision representations. After analyzing these other algorithms, I designed a new one for DANs. The new algorithm, which consisted in a recursive decomposition of the original DAN into a set of symmetric DANs, was evaluated on a wide range of problems described in the literature—such as the  $n$ -test problem for different values of  $n$ . We also compared its efficiency with that of another algorithm for DANs proposed and implemented by Prof. Manuel Luque, as explained in Section 3.4.

### 1.3.2. Cost-effectiveness analysis with MPADs

MPADs had also been under development within the CISIAD previous to the beginning my doctoral research, and a model had been built to model the cost-effectiveness of the human papilloma virus (HPV) vaccine. I took over the development of MPADs in `OpenMarkov` and completed the implementation of algorithms for cost-effectiveness analysis and reimplemented sensitivity analysis almost from scratch. During my three-month stay at the Centre for Health Economics of the University of York in 2013, funded by a PhD scholarship awarded by UNED, I implemented as MPADs several models published in the literature (Chancellor et al., 1997; Briggs et al., 2004; Ryan et al., 2008) based on their `Excel` and `R` implementations. Those models required new types of potentials, such as the Weibull hazard potential, which were not yet available in `OpenMarkov`, as well as the extension of some of the existing algorithms, for example to be able to do inference with numeric variables in certain scenarios. I collaborated with other members of the CISIAD and of the University of York on a paper that describes the MPADs and will be sent the *Medical Decision Making* journal.

I then built MPADs to perform cost-effectiveness analyses of two medical problems in collaboration with two master students of the *Modular Program on Tools for Health Research and Management* of the UNED. Dr. María José Roca, a medical specialist in thoracic surgery at the Hospital Universitario Virgen de la Arrixaca in Murcia and associate professor at the Department of Thoracic Surgery at the University of Murcia Medical School, was interested in analyzing the cost-effectiveness of PleurX® compared to talc pleurodesis for the treatment of malignant pleural effusion. Given that the life expectancy of those patients is short, we initially intended to build an atemporal model; however, we soon realized that it would be more appropriate to build a Markov model with a short cycle length (one week).

We first conducted a literature review aimed at covering all the studies published on the cost and effectiveness of the alternative treatments for malignant pleural effusion. I then built the MPAD with the assistance of the expert: we identified the variables that needed to be in the model and drew causal links amongst them. Finally I encoded the conditional probability potentials and the utility functions. We took most of the model parameters (probabilities and utilities) from the literature, but others were estimated by the expert. I then performed the cost-effectiveness analysis with *OpenMarkov*, including one-way sensitivity analysis for the parameters having a higher variability in the literature. Finally, we wrote a paper describing the model and the results of the CEA, which we intend to submit to the *Medical Decision Making* journal after the paper on MPADs. The collaboration with the expert was carried out with regular meetings remotely (through Skype and TeamViewer).

Similarly, Inés Vázquez wanted to conduct an analysis on the cost-effectiveness of mammography in Spain, considering different frequencies, age ranges, and risk factors. She conducted a literature review. We started to build the MPAD based on the model built by Schousboe et al. (2011). That model was based on data from the Surveillance, Epidemiology, and End Results (SEER) database, which contains epidemiological information about cancer in the USA. Even though that model is described with great detail in the supplement of their paper, I had to obtain some transition probabilities from the SEER database using their statistical software, *SEER\*Stat*. We also contacted Dr. Schousboe to discuss with him some counter-intuitive results of his model and ours.

We then tried to replace the parameters calculated from the SEER database with others calculated from Spanish data. The problem is that the model of Schousboe et al., which we had reproduced approximately, was not causal but observational, and therefore we could not use the available Spanish data, as they were not of the same nature as those in the SEER database.

### 1.3.3. Decision-support system for cochlear implant fitting

When CISIAD joined the Opti-FOX project, funded by EU 7th Framework Programme, I became involved in the development of a decision-support system for



cochlear implant fitting in collaboration with experts at the Eargroup, a hearing clinic in Antwerp (Belgium). I first conducted a literature review on cochlear implants and previous applications of artificial intelligence to the problem of cochlear implant fitting, which is summarized in Section 7.2.2. We then proceeded to the construction of a probabilistic graphical model to model the behavior of cochlear implants with assistance of the experts. We had regular meetings with them, most of the times remotely (through Skype) but we also traveled to Antwerp regularly for face to face meetings. In one of them we felt the need for a new type of canonical model (Díez and Druzdel, 2006), which we called the *tuning model*, and a new type of PGM, which we called *tuning networks*. Due to the size of the resulting network and its high degree of connectivity, exact inference algorithms were inapplicable because they soon ran out of memory. For this reason, I implemented a well-known stochastic algorithm for Bayesian networks, called *likelihood weighting* (Fung and Chang, 1990; Shachter and Peot, 1990), and extended it to PGMs with utility nodes. During my three month research stay at the Machine Intelligence group of the University of Aalborg in 2012, I designed and implemented object-oriented probabilistic networks as an extension of object-oriented Bayesian networks (Koller and Pfeffer, 1997; Bangsø and Willemin, 2000) to models with utility and decision nodes. We evaluated the model mainly testing it with cases stored in the clinical database. The expert evaluated the quality of the output of the model and we compared it with that of a previous approach.

The collaboration with the Eargroup continued after the end of the Opti-FOX project with a new grant from the Marie Curie Action “Industry-Academia Partnerships and Pathways” program (FP7-PEOPLE-2012-IAPP) funded by the European Union. Since July 2013, I have been collaborating with the ENT clinician Dr. Paul Govaerts in the further development of the decision-support system. This time we changed our approach to combine the physical model published in Vaerenberg et al. (2014b) with probabilistic graphical models. The new system has been evaluated in several patients, achieving generally good results.

## 1.4. A brief note on notation

The following notation rules are observed throughout this dissertation:

- Variables are represented by an italic capital letter, such as  $X$ , sometimes followed by a subscript, e.g.,  $X_1$ .
- Values taken by variables are represented by an italic lower case letter. For example,  $x_i$  is a value of  $X_i$ .
- Sets of variables  $\{X_1, \dots, X_n\}$  are represented by a bold capital letter, e.g.,  $\mathbf{X}$ .
- Configurations of  $\mathbf{X}$ , where each variable  $X_i$  takes on a value  $x_i$ , are represented by a lower case bold letter, such as  $\mathbf{x}$ .

## 1.5. Organization of the thesis

The first part of this thesis, Fundamentals, is completed with Chapter 2, which describes the state of the art of PGMs and cost-effectiveness analysis. Part II describes three new types of PGMs: Chapter 3 introduces decision analysis networks (DANs) and the algorithms for evaluating them, Chapter 4 presents Markov processes with atemporal decisions (MPADs) and their application to cost-effectiveness analysis in medicine, and Chapter 5 defines tuning networks. Part III describes the applications of these new types of PGMs to particular medical problems: Chapter 6 includes two cost-effectiveness analyses performed using MPADs, one for malignant pleural effusion and the other for mammography, and Chapter 7 describes the decision-support system for fitting cochlear implants. Finally, Chapter 8 summarizes the conclusions and proposes some lines of future work.

## 2. State of the art

### 2.1. An introduction to probabilistic graphical models

Probabilistic graphical models (PGMs) are models that use a graph-based representation for compactly encoding a complex joint probability distribution over a set of variables.

#### 2.1.1. Basic definitions

A *graph*  $G = (\mathbf{V}, \mathbf{E})$  consists of a set of nodes  $\mathbf{V}$  and a set of edges  $\mathbf{E}$ . An edge is a binary relation between two nodes  $(X, Y)$ , where  $X, Y \in \mathbf{V}$  and  $X \neq Y$ . An edge is *directed* if  $X$  and  $Y$  are ordered in  $(X, Y)$  and usually represented as  $X \rightarrow Y$ ; if an edge is not directed it is said to be *undirected*. We will refer to directed edges as *arcs*. If every edge in  $\mathbf{E}$  is directed, then  $G$  is a *directed graph*, whereas if every edge in  $\mathbf{E}$  is undirected, then  $G$  is an *undirected graph*.

A *path* from node  $X$  to node  $Y$  in a graph  $G = (\mathbf{V}, \mathbf{E})$  is a sequence  $\{X_0, X_1, \dots, X_n\}$  of distinct nodes such that  $(X_{i-1}, X_i)$  is an edge in  $\mathbf{E}$  for each  $i$  such that  $1 \leq i \leq n$ ,  $X = X_0$ , and  $Y = X_n$ . The path is *directed* if all edges are directed. A graph is said to be a *tree* if each pair of distinct nodes is connected by exactly one path.


A *cycle* is a path where  $X_0 = X_n$ , and a *directed cycle* is a directed path with  $X_0 = X_n$ . A directed graph with no directed cycles is said to be a *directed acyclic graph* (DAG).

Given an edge  $(X, Y)$  from  $X$  to  $Y$ , the node  $X$  is said to be a *parent* of  $Y$ ,  $Y$  a *child* of  $X$ . The set of parents of a node  $Y$  is denoted by  $Pa(\mathbf{Y})$ , and the set of children for a node  $X$  is denoted by  $Ch(\mathbf{X})$ . The nodes from which there exists a directed path to  $X$  are named the *ancestors* of  $X$ . Similarly, the nodes to which there exists a directed path from  $X$  are the *descendants* of  $X$ .

#### 2.1.2. Bayesian networks

A *Bayesian network* (BN) (Pearl, 1988)  $BN = (G, P)$  consists of a DAG  $G = (\mathbf{V}, \mathbf{E})$ , in which each node  $X \in \mathbf{V}$  represents a *chance variable*, and a probability distribution over  $\mathbf{V}$ ,  $P(\mathbf{V})$ , which can be factored as:

$$P(\mathbf{v}) = \prod_{X \in \mathbf{V}} P(x \mid pa(X)), \quad (2.1)$$



**Figure 2.1.:** A Bayesian network for the differential diagnosis of two hypothetical diseases.

where  $pa(X)$  denotes a configuration of the parents of  $X$ .


Since there is a one to one relation between variables and nodes in a BN, we will use the terms node and variable indifferently.

We will assume in this thesis that all the variables in a BN are *discrete*, which means each variable  $X \in \mathbf{V}$  can only take a finite set of mutually exclusive values; this set is called the *domain* of a variable  $X$  and is denoted by  $dom(X) = (x_1, x_2, \dots, x_n)$ .

The quantitative information of a BN is given by assigning a conditional probability distribution  $P(X | Pa(X))$  to each node  $X \in \mathbf{V}$ . A conditional probability distribution defines a conditional probability for each value  $x \in dom(X)$  and each configuration  $pa(X)$  of its parents. A *potential* is a real-valued function defined over a set of variables. The conditional probability distribution  $\phi = P(X | Pa(X))$  is a particular case of potential whose domain is  $dom(\phi) = \{X\} \cup Pa(X)$ . Figure 2.1 shows an example of a Bayesian network.

### 2.1.3. Influence diagrams

An *influence diagram* (ID) (Howard and Matheson, 1984) is a generalization of a BN used to model and solve decision problems. An ID consists of a DAG  $G = (\mathbf{V}, \mathbf{E})$ , where the set  $\mathbf{V}$  has three types of nodes: *chance nodes*  $\mathbf{V}_C$ , *decision nodes*  $\mathbf{V}_D$ , and *utility nodes*  $\mathbf{V}_U$ . Chance nodes, drawn as circles, represent events which are not under the direct control of the decision maker. Decision nodes, drawn as rectangles, correspond to actions under the direct control of the decision maker. Utility nodes, drawn as diamonds, represent the expected benefit or loss, or more generally, the values of the decision maker. Utility nodes can not be parents of chance or decision



**Figure 2.2.:** Influence diagram representing the decision over the test. For a disease whose prevalence is 8% there exists a test with a sensitivity of 75% and a specificity of 96%. The test has a cost of 0.2. There is a therapy for the disease, with an effectiveness of 8 if the disease is present and 8 if it is absent. The effectiveness of not applying the therapy is 3 if the disease is present but 10 if it is absent. Is it worth doing the test? Put another way, the test will give us useful information for deciding when to apply the therapy, but the question is whether this benefit outweighs the cost of the test. In what cases should the therapy be applied?

nodes, but in an extension of IDs proposed by (Tatman and Shachter, 1990) they can be parents of other utility nodes. Such nodes are called *super value nodes* (SVNs).

There are three types of arcs in an ID, depending on the type of node they point at. Arcs pointing at chance nodes represent probabilistic dependency. Arcs pointing at decision nodes, named *informational arcs*, represent availability of information; i.e., if there is an arc from a node  $X$  to a decision node  $D$  then the state of  $X$  is known before decision  $D$  is made. Finally, arcs pointing at utility nodes represent functional dependency: arcs into an ordinary utility node indicate the domain of the associated *utility function*. A utility function is a potential that defines the utility of a node given its parents. Arcs into a SVN  $U$  indicate that the associated utility function is a combination (generally a sum or a product) of the utility functions of the parents of  $U$ .

Standard IDs require a directed path that connects all the decision nodes indicating the order in which the decisions are made. In an ID with  $n$  decisions  $\{D_1, \dots, D_n\}$ , this induces a partition of  $\mathbf{V}_C$  into  $n + 1$  disjoint subsets,  $\mathbf{C}_0, \mathbf{C}_1, \dots, \mathbf{C}_n$ , where  $\mathbf{C}_i$  contains every chance variable  $C$  with an arc  $C \rightarrow D_i$  but no arc  $C \rightarrow D_j$  to a previous decision,  $j < i$ ; i.e.,  $\mathbf{C}_i$  is the subset of chance variables known for  $D_i$  but unknown for previous decisions. This induces a *partial order*  $\prec$  in  $\mathbf{V}_C \cup \mathbf{V}_D$ :

$$\mathbf{C}_0 \prec \{D_0\} \prec \mathbf{C}_1 \prec \dots \prec \{D_n\} \prec \mathbf{C}_n . \quad (2.2)$$

The variables known when making decision  $D_j$  are known as the *informational predecessors* of  $D_j$  and denoted  $iPred(D_j)$ . By assuming the *no-forgetting* hypothesis, which states that the decision maker remembers all previous decisions and observations, we have  $iPred(D_i) \subseteq iPred(D_j)$  (for  $i \leq j$ ). In particular,  $iPred(D_j)$  is the set of chance variables that occurs before  $D_j$  under  $\prec$ , i.e.,  $iPred(D_j) = \mathbf{C}_0 \cup \{D_0\} \cup \mathbf{C}_1 \cup \dots \cup \{D_{i-1}\} \cup \mathbf{C}_i$ .

The quantitative information defining an ID consists of a conditional probability distribution  $p(C \mid Pa(C))$  for each chance node  $C$ , and a utility function or potential  $\psi_U(pa(U))$  for each utility node. The potential of an ordinary utility node maps each configuration of its parents onto a real number, whereas the potential of a SVN defines a utility-combination function.

For each configuration  $\mathbf{v}_D$  of the decision variables in  $\mathbf{V}_D$  we have a joint probability distribution defined over the set of random variables  $\mathbf{V}_C$ :

$$P(\mathbf{v}_C \mid \mathbf{v}_D) = \prod_{X \in \mathbf{V}_C} P(x \mid pa(X)), \quad (2.3)$$

This equation represents the probability of configuration  $\mathbf{v}_C$  when the decision variables are set to the values  $\mathbf{v}_D$ .

### 2.1.3.1. Policies and strategies

A *stochastic policy*  $P_D$  for a decision  $D$  is a probability distribution defined over  $D$  and conditioned on the set of its informational predecessors:  $P_D(d \mid iPred(D))$ . If  $P_D$  is degenerate (consisting only of ones and zeros) then we say that the policy is *deterministic*. A deterministic policy  $\delta_{D_i}$  assigns each of the configurations of the informative predecessors with a value  $d$  of  $D$ .

A *strategy*  $\Delta$  is the set of policies for all the decisions in  $\mathbf{V}_D$ ,  $\{P_D \mid D \in \mathbf{V}_D\}$ . If all the policies in the strategy  $\Delta$  are deterministic, then  $\Delta$  is said to be *deterministic*; otherwise it is *stochastic*. The literature about IDs generally assumes that the strategies are deterministic and so will we in this dissertation.

A strategy  $\Delta$  *induces* a joint probability distribution over  $\mathbf{V}_C \cup \mathbf{V}_D$  defined as:

$$\begin{aligned} P_\Delta(\mathbf{v}_C, \mathbf{v}_D) &= P(\mathbf{v}_C \mid \mathbf{v}_D) \prod_{D \in \mathbf{V}_D} P_D(d \mid iPred(D)) \\ &= \prod_{C \in \mathbf{V}_C} P(c \mid pa(C)) \prod_{D \in \mathbf{V}_D} P_D(d \mid iPred(D)). \end{aligned} \quad (2.4)$$

Let  $I$  be an ID,  $\Delta$  a strategy for  $I$  and  $\mathbf{r}$  a configuration defined over a set of variables  $\mathbf{R} \subseteq \mathbf{V}_C \cup \mathbf{V}_D$  such that  $P_\Delta(\mathbf{r}) \neq 0$ . The conditional probability distribution *induced by the strategy  $\Delta$  given the configuration  $\mathbf{r}$* , defined over  $\mathbf{R}' = (\mathbf{V}_C \cup \mathbf{V}_D) \setminus \mathbf{R}$ , is

$$P_\Delta(\mathbf{r}' \mid \mathbf{r}) = \frac{P_\Delta(\mathbf{r}, \mathbf{r}')}{P_\Delta(\mathbf{r})}. \quad (2.5)$$

Using this distribution  $P_{\Delta}(\mathbf{r}' \mid \mathbf{r})$  we can compute the *expected utility of node  $U$  under the strategy  $\Delta$  given the configuration  $\mathbf{r}$*  as:

$$EU_U(\Delta, \mathbf{r}) = \sum_{\mathbf{r}'} P_{\Delta}(\mathbf{r}' \mid \mathbf{r}) \psi_U(\mathbf{r}, \mathbf{r}'). \quad (2.6)$$

We define the *expected utility of  $U$  under the strategy  $\Delta$*  as  $EU_U(\Delta) = EU_U(\Delta, \blacklozenge)$ , where  $\blacklozenge$  is the empty configuration. We have that

$$EU_U(\Delta) = \sum_{\mathbf{v}_C} \sum_{\mathbf{v}_D} P_{\Delta}(\mathbf{v}_C, \mathbf{v}_D) \psi_U(\mathbf{v}_C, \mathbf{v}_D). \quad (2.7)$$

The *expected utility of the strategy  $\Delta$* , denoted by  $EU(\Delta)$ , is the sum of the expected utilities under the strategy  $\Delta$  of all the utility nodes in  $I$  having no descendants.

An *optimal strategy* is a strategy  $\Delta_{opt}$  that maximizes the expected utility:

$$\Delta_{opt} = \arg \max_{\Delta \in \Delta^*} EU(\Delta), \quad (2.8)$$

where  $\Delta^*$  is the set of all the strategies for  $I$ .

The policies in an optimal strategy are said to be *optimal policies*. An *optimal policy*  $\delta_{D_i}$  is a function that maps each configuration of the variables in  $iPred(D_{i-1})$ , onto the value  $d_i$  of  $D_i$  that maximizes the expression at the right of  $D_i$  (in the case of a tie, any maximizing value can be chosen arbitrarily):

$$\delta_{D_i}(iPred(D_i)) = \arg \max_{d_i \in D_i} \sum_{\mathbf{c}_i} \max_{d_{i+1}} \dots \sum_{\mathbf{c}_{n-1}} \max_{d_n} \sum_{\mathbf{c}_n} P(\mathbf{v}_C : \mathbf{v}_D) \psi_{U_0}(\mathbf{v}_C, \mathbf{v}_D). \quad (2.9)$$

The *maximum expected utility (MEU)* is the expected utility of the optimal strategy.

$$MEU = EU(\Delta_{opt}) = \max_{\Delta \in \Delta^*} EU(\Delta). \quad (2.10)$$

The evaluation of an ID consists in finding the *MEU* and an optimal strategy, composed by an optimal policy for each decision. It can be proved (Cowell et al., 1999) that

$$MEU = \sum_{\mathbf{c}_0} \max_{d_0} \dots \sum_{\mathbf{c}_{n-1}} \max_{d_{n-1}} \sum_{\mathbf{c}_n} P(\mathbf{v}_C : \mathbf{v}_D) \psi_{U_0}(\mathbf{v}_C, \mathbf{v}_D). \quad (2.11)$$

We can always find a deterministic optimal strategy for an ID and, in the case of ties, there may even be more than one.

### 2.1.4. Canonical models

The construction of probabilistic graphical models, such as Bayesian networks and influence diagrams, requires the specification of many conditional probability distributions of the form  $P(y \mid \mathbf{x})$ , where  $\mathbf{X} = \{X_1, \dots, X_n\}$  is the set of parents of a node  $Y$  in the network—see Figure 2.3.


The set  $\{Y\} \cup \mathbf{X}$  is called a *family*. When all the variables in a family are discrete,  $P(y \mid \mathbf{x})$  can be expressed in the form of a conditional probability table (CPT), with as many numerical parameters as the number of different combinations of the values of the variables in the family. Therefore, the number of numerical parameters grows exponentially with the number of parents. Usually, the numerical parameters are estimated from databases or assessed by human experts and, for this reason, it is usually difficult to build a CPT for a family having more than three or four parents.

One way of reducing the complexity of the elicitation of the numerical probabilities is to rely on the so-called *canonical models*. The term “canonical” is used because such models are elementary units used in the construction of more complicated models (Pearl, 1988). Canonical models usually allow the definition of complex probability distributions from a small number of parameters. Different canonical models may coexist in any PGM. For instance, in causal Bayesian networks modeling real-world domains, it is frequent to find a significant number of families that interact through OR/MAX-models and a few who do it through AND-models, while the rest of the families use regular CPTs.

Canonical models not only simplify the construction of probabilistic models (knowledge engineering), but also save storage space and computation time (Díez and Druzdzel, 2006) and they correspond to causal patterns that can be exploited to generate user explanations (Lacave and Díez, 2002).

#### 2.1.4.1. ICI models

A common type of canonical models are ICI models, which are based on the *independence of causal influence* (ICI) assumption. This assumption implies that there are no interactions among the causal mechanisms by which the parents  $X_i$  affect the



**Figure 2.3.:** A family consisting of a child node  $Y$  and its parents,  $\mathbf{X} = \{X_1, \dots, X_n\}$ .



value of the child  $Y$ . *Noisy ICI models* are defined by introducing  $n$  auxiliary variables  $\{Z_1, \dots, Z_n\}$ , as shown in Figure 2.4, such that  $Y$  is a deterministic function of the  $Z_i$ s,  $y = f(\mathbf{z})$ , and the value of each  $Z_i$  depends probabilistically only on  $X_i$ , as captured by the CPT  $P(z_i | x_i)$ . The conditional probability  $P(y | \mathbf{x})$  is obtained by marginalizing out the  $Z_i$ s:


$$P(y | \mathbf{x}) = \sum_{\mathbf{z}} P(y | \mathbf{z}) \cdot P(\mathbf{z} | \mathbf{x}), \quad (2.12)$$

where

$$P(y | \mathbf{z}) = \begin{cases} 1 & \text{if } y = f(\mathbf{z}) \\ 0 & \text{otherwise.} \end{cases} \quad (2.13)$$

Therefore,

$$P(y | \mathbf{x}) = \sum_{\mathbf{z} | f(\mathbf{z})=y} P(\mathbf{z} | \mathbf{x}). \quad (2.14)$$



**Figure 2.4.:** Auxiliary structure for the derivation of a noisy ICI model.

Given the graph in Figure 2.4, the ICI assumption implies the absence of links  $X_i \rightarrow Z_j$  and  $Z_i \rightarrow Z_j$  for all  $i \neq j$ , which means that

$$P(\mathbf{z} | \mathbf{x}) = \prod_i P(z_i | x_i) \quad (2.15)$$

and, consequently,

$$P(y | \mathbf{x}) = \sum_{\mathbf{z} | f(\mathbf{z})=y} \prod_i P(z_i | x_i). \quad (2.16)$$

Each parameter  $P(z_i | x_i)$  of an ICI model is associated with a particular link  $X_i \rightarrow Y$ , while each parameter  $P(y | \mathbf{x})$  in a CPT corresponds to a certain configuration  $\mathbf{x}$  of all the parents of  $Y$ , and cannot be associated to any particular link. This property, stemming from the ICI assumption, entails two advantages from the point of view of knowledge engineering. The first is a significant reduction in the number of parameters required to specify a model, from  $O(\exp(n))$  in a general model to  $O(n)$  in an ICI model. This can amount to a substantial reduction of the elicitation effort. For example, a binary node with 10 binary parents will have a CPT consisting of  $2^{11} = 2,048$  numerical parameters. Adding one more parent doubles this number to  $2^{12} = 4,096$  parameters. In contrast, a noisy OR model would require only 10 and 11 parameters, respectively. The second advantage is that the parameters in canonical models lend themselves to fairly intuitive interpretations, which facilitates the task of eliciting them from human experts.

It is possible to define leaky ICI models (Díez and Druzdzel, 2006), which only differ from their noisy counterparts in the addition of another auxiliary variable,  $Z_L$ , which accounts for the effect of variables not explicitly represented in the model.

## 2.2. Asymmetric decision problems

### 2.2.1. Basic definitions

The term *asymmetry* in decision analysis stems from decision trees (DTs) (Raiffa and Schlaifer, 1961), the first and probably the most widely used formalism in the field, and it comprises two rather different features. In a DT, a path from the root to a leaf is called a *scenario*. A decision problem is said to be *symmetric* if all its possible DT representations are symmetric. For a DT to be symmetric, (1) the number of scenarios must be equal to the Cartesian product of the state spaces of all chance and decision variables, and (2) a decision tree representation of the problem must exist such that the sequence of variables is the same in all scenarios (Bielza et al., 2011). A decision problem is said to be *asymmetric* if it is not symmetric. In order to distinguish between the two types of asymmetry, we say there is *structural asymmetry* when the value taken on by a variable restricts the domain of other variables, and that there is *order asymmetry* when several orderings of the decisions and observations are possible (Jensen et al., 2006).

The  $n$ -test problem is a good example containing both order and structural asymmetry. It can be stated as follows: a patient goes to the doctor showing some symptoms and the doctor wants to know whether the patient has a certain disease or not. The doctor has to decide which of the  $n$  available tests (each one with an associated cost, specificity, and sensitivity) to perform and in which order. Each test can be performed at most once and its result will be known immediately. The diabetes problem Demirer and Shenoy (2006) is an instance of the  $n$ -test problem with  $n = 2$ . It contains structural asymmetry because the information available (the test results) depends on the decision made for each test, we will know its result only if we decide to do a certain test. It also contains order asymmetry because different orderings of the decisions lead to different information being available when making the decisions.

### 2.2.2. Shortcomings of traditional representations

DTs have the advantage of almost absolute flexibility, but also have three drawbacks: their size grows exponentially with the number of variables, they cannot represent conditional independencies, and they require in general a preprocessing of the probabilities (Howard and Matheson, 1984; Bielza et al., 2011); for example, medical diagnosis problems are usually stated in terms of direct probabilities, namely the prevalence of the diseases and the sensitivity and specificity of the tests, while DTs are built with inverse probabilities, i.e., the positive and negative predictive values of the tests. Even in cases with only a few chance variables, this preprocessing of probabilities is a difficult task.

After DTs, IDs (Section 2.1.3) are the most widely used formalism and have the advantages of being very compact, easily representing conditional independence, and using direct probabilities, but they are only appropriate to model symmetric decision problems. Problems having only structural asymmetry can be symmetrized by using two modeling tricks (Smith et al., 1993). The situation where a chance variable  $X$  (e.g. the test result) has no sense for a value of parent variable  $Y$  (e.g. the decision not to make a test), can be modeled by adding a dummy state to  $X$ . For example, the test problem can be represented with an ID if the test result has three values: *positive*, *negative*, and (the dummy state) *not-performed*. The drawback of this trick is that it complicates the edition of the probability tables. It also makes the evaluation less efficient, due to the enlarged probability and utility tables, and complicates the interpretation of the results, because some policies contain configurations that can never occur, such as (*do not test, positive*), (*do not test, negative*), and (*do test, not-performed*). If a value  $d$  of decision  $D$  cannot be made when parent node  $X$  takes value  $x$ , a dummy utility node  $U$  can be added such that  $u(x, d) = -\infty$  and  $u(x', d') = 0$  for the other configurations of  $X$  and  $D$ .

Order asymmetry poses a much more serious difficulty to IDs. For example, if we try to model the  $n$ -test problem with an ID, we need a  $n$  decision nodes for the tests,  $\{T_1, \dots, T_n\}$ , each having  $n + 1$  options; the extra option is “do no test”. We then need  $n(n - 1)/2$  dummy utility nodes to represent the restriction that tests cannot be repeated, and that if the  $i$ -th decision is not to test, then all subsequent test decisions are not to test. We also need  $n$  chance variables,  $\{R_1, \dots, R_n\}$ , for the results of the tests. The meaning of  $R_i$  depends which test has been done in the  $i$ -th place. Each  $R_i$  has  $m + 1$  possible states, where  $m$  is the maximum of the number of outcomes of the tests; the extra state is “not performed”. If a test has fewer than  $m$  outcomes, some of the states of the  $R$ 's are meaningless. Combining continuous and discrete tests in the same model is impossible. Specifying the conditional probability tables for the  $R$ 's is cumbersome. Additionally, the same table is repeated for every  $R_i$ , which complicates the maintenance of the model and makes sensitivity analysis virtually impossible. If a new test is added to the model, the domains of the all the  $T$ 's and the  $R$ 's must be revised, as well as the conditional probability tables of the  $R$ 's and the tables of the dummy utility nodes; the effort is comparable to building the new model from scratch.

### 2.2.3. Representations for asymmetric decision analysis

In practice, virtually all real-world problems are asymmetric, in particular all those that involve the possibility of getting additional information at a cost. Several formalisms have been proposed for representing and solving asymmetric decision problems.

Smith, Holtzman, and Matheson (Smith et al., 1993) proposed an asymmetric representation of IDs (SHM IDs) where the structure of the ID remained unchanged

but the conditional distribution was represented by distribution trees, thus avoiding the need for dummy states.

Asymmetric influence diagrams (AIDs) (Nielsen and Jensen, 2000) were presented as IDs where the requirement for a total ordering of the decisions was relaxed and where the structural asymmetry was encoded at the qualitative level. In practice, the total order was required except in the cases where the structure of the decision problem rendered the decisions conditionally independent.

Sequential valuation networks (SVNs) (Demirer and Shenoy, 2006; Shenoy, 2000) are a representation based on valuation networks (Shenoy, 1992) and sequential decision diagrams (Covaliu and Oliver, 1995). Structural asymmetry is modeled in SVNs by specifying conditions in labeled arcs but it is more difficult to model order asymmetry, since all admissible decision and observation sequences of variables must be established with arcs.

Unconstrained influence diagrams (UIDs) (Jensen and Vomlelová, 2002) were the first extension of IDs where order asymmetry was efficiently represented: the order of decisions and observations need only be partially specified, without the need for information links. Unfortunately, they were as inappropriate to represent structural asymmetry as conventional IDs.

Sequential influence diagrams (SIDs) (Jensen et al., 2006), combine the representation of structural asymmetry with labeled arcs of SVNs with the order asymmetry representation of UIDs.

In summary, SHM IDs, AIDs, and SVNs were designed to represent structural asymmetries, while UIDs were designed only for order asymmetry. SIDs were designed for both; they consist of two overlapping graphs, one for the probability and utility relations, and the other for modeling the sequence of decisions and observations, which makes them difficult to build, to evaluate, to communicate to experts, and to maintain—see also the comparisons in (Bielza et al., 2011; Bielza and Shenoy, 1999).

## 2.3. Cost-effectiveness analysis

### 2.3.1. Basic concepts

Cost-effectiveness analysis (CEA) is the economic analysis that compares the relative costs and outcomes (effects) derived from two or more possible course of actions. In medicine, these actions are usually medical treatments, also called interventions. The origin of CEA is in constrained optimization, as a means to maximize the health-related function subject to a budget constraint, such as in a public health care system. As such, CEA is a particular case of multicriteria decision making with two objective functions: the health-related outcome (effectiveness), and the economic cost, measured in monetary units. The impact of a treatment on health is often measured in quality-adjusted life years (QALY) (Drummond et al., 2005), which seeks to combine both length of life and health-related quality of life in a single measure.

For each pair of interventions  $A$  and  $B$ , being  $C_A$  and  $C_B$  their respective costs, and  $E_A$  and  $E_B$  their respective effectiveness: if  $C_A < C_B$  and  $E_A > E_B$ , then we say  $A$  dominates  $B$ , as it is less costly and more effective; if on the other hand  $C_B > C_A$  and  $E_B > E_A$ , we define the *incremental cost-effectiveness ratio* (ICER) of  $A$  and  $B$  as:

$$ICER(A, B) = \frac{(C_B - C_A)}{(E_B - E_A)}.$$

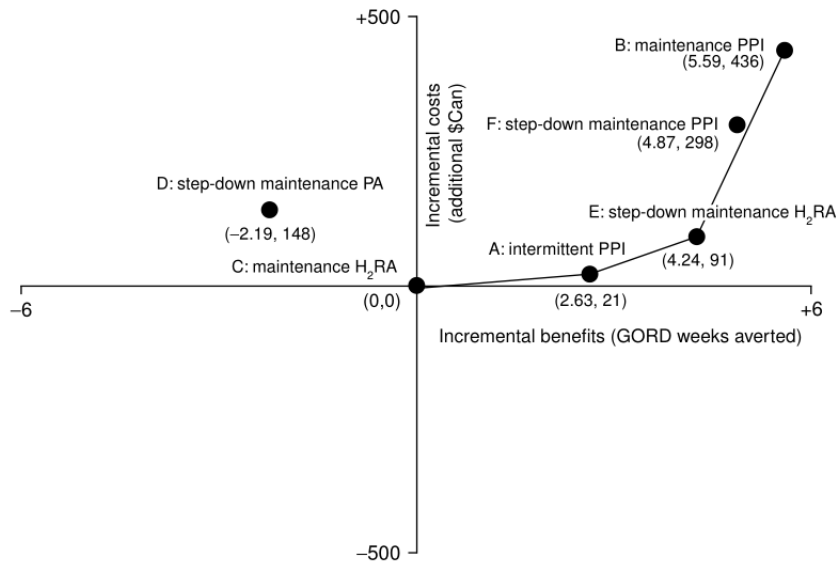
An intervention  $A$  might be *dominated by extended dominance* if there are interventions  $B$  and  $C$  such that  $ICER(A, C) > ICER(B, C)$ .

The results of cost effectiveness analysis can be presented in a graph where the horizontal axis represents effectiveness and the vertical axis represents cost. In this graph, called the *cost-effectiveness plane*, each intervention is represented by a point whose position is determined by its cost and effectiveness. The frontier of efficiency is formed by those interventions that are not dominated. An illustrative example is shown in Figure 2.5.

When the ICER for a treatment is lower than a threshold  $\lambda$ , usually called *willingness to pay* or *cost-effectiveness threshold*, we say it is cost-effective. This threshold varies wildly across different countries and has been and still is subject to passionate debate Claxton et al. (2015). The (NMB) of an intervention  $I_A$  as a function of  $\lambda$  can be defined as

$$NMB_{I_A}(\lambda) = \lambda \cdot E_A - C_A.$$

Therefore, if the value of  $\lambda$  is known, we are able to compute the NMB of each intervention. The intervention that maximizes the NMB is the most cost-effective.

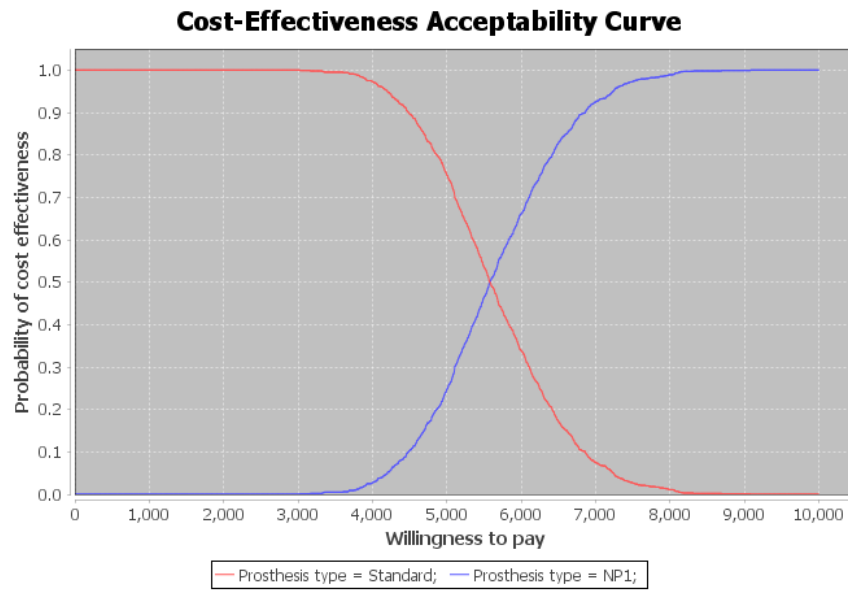


**Figure 2.5.:** Cost-effectiveness plane of treatments for gastro-oesophageal reflux disease (GORD) relative to maintenance H<sub>2</sub>RA Goeree et al. (1999). The line joining strategies C, A, E and B is the efficiency frontier. Treatment F is dominated by extended dominance by treatment B.

### 2.3.2. Modeling for CEA

Decision analytic models use mathematical relationships to define a series of consequences derived from a set of alternative options. Based on the inputs of the model, the likelihood of each consequence is expressed in terms of probabilities. Each consequence has an estimated cost and an outcome. It is therefore possible to calculate the expected cost and expected outcome of each option being evaluated. A key purpose of decision modeling is to allow for the variability and uncertainty associated with decisions. The use of models for economic evaluation in health care has been controversial, but the popularity of decision modeling has increased due to the required features of any economic evaluation, such as using all relevant evidence, considering all relevant options, using an appropriate time horizon (often as long as a lifetime), and the ability to translate uncertainty in the evidence into decision uncertainty. Randomized controlled trials are important to generate evidence for the evaluation of health care interventions, but replacing models with trials for economic evaluation is usually not feasible, as it is not possible to meet the requirements we just listed.

The parameters of a model are usually estimated—either from data or by human experts—and are therefore subject to uncertainty. The purpose of *sensitivity analysis* is to reflect the uncertainty in the input parameters and analyze how it translates into uncertainty over the outputs of interest (cost, effects, and derived measures, such as ICER) and therefore into decision uncertainty (Briggs et al., 2006). The simplest



**Figure 2.6.:** Cost-effectiveness acceptability curve for different types of hip prostheses Briggs et al. (2004)

form of sensitivity analysis, called *one-way sensitivity analysis* is to simply vary one value in the model by a given amount, and examine the impact that the change has on the model's results. *Probabilistic sensitivity analysis* (PSA), is a form of sensitivity analysis where some parameters are defined by probability distributions (also called second order probabilities) rather than point estimates (Claxton et al., 2005). PSA has the advantage, in comparison with one-way sensitivity analysis, of accounting for the uncertainty in all parameters at the same time. The choice of distribution to reflect uncertainty is not arbitrary and must be done according to the type of the parameter and its method of estimation.

Commonly Monte Carlo simulation is used to propagate uncertainty in the model. A high number of simulations is run by drawing samples from the probability distributions. Each simulation results in a estimation of expected costs and efficiency. The variability in these results reflects the uncertainty in the expected cost and efficiency of interventions.

One way to represent the results of PSA is the *cost-effectiveness acceptability curve* (CEAC) (Briggs, 1999). It presents the probability that a given intervention is more cost-effective than the alternatives for a range of willingness to pay thresholds. For example, Figure 2.6 shows the CEAC for different prostheses for hip replacement as described in Briggs et al. (2004).

Decision uncertainty implies that even if we make the best decision based on available evidence, there is a chance that the wrong decision will be made. The expected cost of uncertainty is determined by the probability of making the wrong decision and the



cost of making the wrong decision in terms of health benefit and resources forgone. The expected costs of uncertainty can be interpreted as the expected value of perfect information (EVPI), as perfect information eliminates the possibility to make the wrong decision.

There is a wide range of modeling techniques available, and some problems are more naturally represented in certain modeling techniques than in others. The rest of this section describes the most relevant modeling frameworks, their strengths and weaknesses.

### 2.3.2.1. Atemporal models

For relatively simple models or decision problems with very short time horizons, atemporal models may be appropriate. The common feature of these models is that the elapse of time is not explicit. Such models include decision trees and influence diagrams (Nielsen et al., 2007; Arias and Díez, 2015). Atemporal models are not suitable to model decision problems involving a risk that is continuous over time, in which events may occur more than once, and when the utility of an outcome depends on when it occurs. Most analytic problems involve at least one of these considerations. Hence, modeling such problems with atemporal models may require unrealistic or unjustified simplifying assumptions.

### 2.3.2.2. State transition models

State transition models (STM) are probably the most widespread kind of models for health economic modeling. STMs conceptualize the treatment of a patient as its evolution over time through a finite set of mutually exclusive and exhaustive health states: it is considered that a patient is in one of the health states at any time. All clinically important events are modeled as transitions from one state to another. Transitions have an assigned probability whereas each state has a cost and an outcome. The structure of such models is usually represented using state transition diagrams, which define the set of states (with circles) and the allowed transitions amongst them (with arrows), and Markov cycle trees (Hollenberg, 1984), which consist of a decision tree whose leaves are state transition diagrams.

STMs do not capture interactions among individuals; they model a single (closed) cohort, and allow transitions to occur only at specified time intervals. In fact, time is discretized into a finite number of cycles of a fixed length, and transitions only happen at a certain moment in the cycle. STMs are useful when the model requires time-dependent parameters (e.g., probability of recurrence for a patient with AIDS), time to an event (e.g., disease-free survival), or repeated events (e.g., second heart attack). STMs include Markov model cohort simulation (Becker and Geiger, 1994; Sonnenberg and Beck, 1993) as well as individual-level STMs (Siebert et al., 2012).

Markov models, also called cohort models, are named after their assumption of the Markov property. According to the Markov property, transition probabilities do not depend on history, i.e. neither on past states nor on the time spent in the current state. A Markov model is evaluated calculating the proportion of the cohort in each state in each cycle based on transition probabilities. Markov models can overcome the Markov property by creating states that contain information about an individual's history, such as tunnel states (Sonnenberg and Beck, 1993), which represent the number of cycles spent in a certain state. However, this can greatly increase the number of states, resulting in very large models that are difficult to manage. Their main advantages over more complex formalisms are that they are easy to develop, debug, and analyze and that in general are not computationally expensive. Their principal disadvantage over individual-level STMs is their inability to handle history without it leading to state explosion.

Individual-level STMs (also called patient-level simulation or microsimulations) are not limited by the Markov property, and unlike cohort models, they keep track of each individual's history using tracker variables. These models simulate one individual at a time, using first-order Monte Carlo. The main disadvantages of individual-level STMs over Markov models are that they are usually more demanding of data, computationally expensive (often requiring millions of simulations to obtain stable outcomes), and that they are more difficult to build and debug. It is worth noting that an individual-level STM without tracker variables is actually a Markov model that is evaluated using a stochastic inference algorithm instead of exact inference.

### **2.3.2.3. Discrete event simulation**

Discrete event simulation (DES) is a modeling method characterized by the ability to represent complex behavior within populations as well as interactions between individuals with their environments (Pidd, 2004). The term “discrete” refers to the fact that DES moves forward in time at discrete intervals, but time management is more flexible and efficient than restricting occurrences to fixed time intervals, because the model jumps from the time of one event to the time of the next. Instead of a cohort, DES models define a set of entities, which are objects that have attributes, experience events, consume resources and enter queues, over time. DES is especially useful in applications to analyze systems with constrained resources or interactions between individuals (Karnon et al., 2012). In fact, experts agree (Karnon, 2003; Roberts et al., 2012) that simpler models should be used unless the characteristics of the problem justify the extra complexity.

### **2.3.2.4. Dynamic transmission models**

Dynamic transmission models (Pitman et al., 2012), also called system dynamics, are capable of reproducing the direct and indirect effects that may arise from a communicable disease control program. Contrary to other (static) models, which assume

a constant risk of infection, dynamic transmission models define risk of infection as a function of the number of infectious individuals. Dynamic models are important when an intervention affects a pathogen's ecology or when the intervention affects disease transmission. Most dynamic transmission modeling has been performed using system dynamics, where transitions between compartments are governed by differential equations.



**Part II.**  
**NEW TYPES OF PGM'S**



## 3. Decision analysis networks

*You take the blue pill, the story ends. You wake up in your bed and believe whatever you want to believe. You take the red pill, you stay in wonderland, and I show you how deep the rabbit hole goes.*

Morpheus

### 3.1. Basic properties of DANs


Pushed by the severe complications encountered when trying to model complex medical decision problems with existing formalisms, (Díez et al., 2012) proposed a new formalism called *decision analysis networks* (DANs). The objective was to create a formalism that would help represent real-world asymmetric decision problems more naturally.

**Example** We will illustrate the definition and evaluation of DANs with a running example, which is very similar to the diabetes problem described in (Demirer and Shenoy, 2006). Consider a physician who is trying to diagnose whether or not a patient is suffering from diabetes. The patient might or might not show a visible symptom of the illness. The physician has the chance to perform two different tests, a blood test and a urine test. Each test has an assigned cost and different specificity and sensitivities. Ultimately, the physician has to decide whether to treat the patient for diabetes or not. The DAN representing this problem is shown in Figure 3.1.

#### 3.1.1. Definition

The set of variables of a DAN,  $\mathbf{V}$ , is partitioned, as in influence diagrams, into three disjoint subsets: chance variables  $\mathbf{C}$ , decisions  $\mathbf{D}$ , and utility variables  $\mathbf{U}$ . A DAN also has an directed acyclic graph such that each node represents a variable.

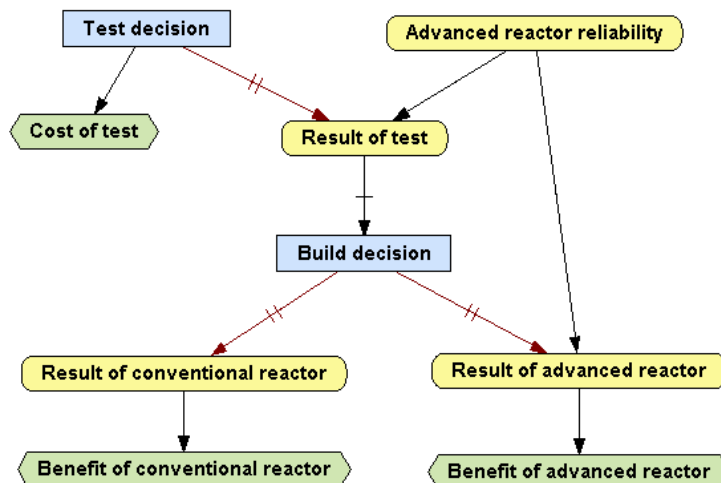
**Restrictions** A *restriction* associated to a link  $X \rightarrow Y$ , such that  $X$  and  $Y$  are chance or decision variables, is a pair  $(x, y)$ , where  $x$  is a value of  $X$  and  $y$  is a value of  $Y$ . It means that variable  $Y$  cannot take the value  $y$  when  $X$  takes the value



**Figure 3.1.:** A DAN for the diabetes problem. The decision about a test may depend on the presence of the symptom and on the result of the other test. There is no constraint on the order of the tests.

$x$ . When all the values of  $Y$  are incompatible with  $x$ , a particular value of  $X$ , we say that there is a *total restriction* and denote it by  $(x, Y)$ . It implies that in some scenarios the variable  $Y$  does not exist. When there is a restriction  $(x, y)$  but every value of  $X$  is compatible with at least one value of  $Y$ , we say that there is a *partial restriction*. Thus, the DAN for the reactor problem (Covaliu and Oliver, 1995; Bielza and Shenoy, 1999) (Figure 3.2) has a partial restriction for the link *Result of test*  $\rightarrow$  *Build decision*, because a bad test result prevents the construction of an advanced reactor, as shown in Figure 3.3, but every value of *Result of test* is compatible with at least one value of *Build decision*.





**Figure 3.2.:** A DAN for the reactor problem. The main decision is which type of reactor to build, if any.

If  $Y$  is a chance variable, the restriction  $(x, y)$ —associated to the link  $X \rightarrow Y$ , which implies that  $X \in Pa(Y)$ —means that  $P(y | pa(Y)) = 0$  for all the configurations of  $Pa(Y)$  in which  $X = x$ . Therefore, when  $Y$  is a chance variable, there is not a significant difference between attaching a partial restriction to a link and setting to 0 the corresponding cells in the conditional probability table.

Result of test	bad	good	excellent
build none	1	1	1
build conven...	1	1	1
build advanced	0	1	1

**Figure 3.3.:** Compatibility table for the link  $Result\ of\ test \rightarrow Build\ decision$  in Figure 3.2. It means that a bad result prevents the construction of an advanced reactor.

In contrast, if  $D$  is a decision, the restriction  $(x, d)$  means that when  $X = x$  the decision maker cannot choose the option  $d$ —see again the example in Figure 3.3.

**Potentials** Each chance node  $Y$  has an associated conditional probability potential, denoted by  $\psi(y | pa(Y))$ . When there is a restriction  $(x, y)$  and  $X = x$  in the configuration  $pa(Y)$ , then  $\psi(y | pa(Y)) = 0$ , because  $y$  is incompatible with  $x$ . If at least one value of  $Y$  is compatible with all the values of its parents in the configuration  $pa(Y)$ , then  $\psi(y | pa(Y))$  is a conditional probability distribution for  $Y$  given that configuration, as shown in Figure 3.4. Similarly, each utility node  $U$  has an associated potential,  $u(pa(U))$ .

Dec: Blood Test	no	no	yes	yes
Diabetes	absent	present	absent	present
positive	0	0	0.02	0.96
negative	0	0	0.98	0.04

**Figure 3.4.:** Probabilistic potential associated to the chance node *Blood test result*.

The first two columns are 0 due to the total restriction (*Dec: Blood Test = no, Blood test result*). The third column defines a conditional probability distribution for *Blood test result* (0.98 is the specificity of the test). The fourth column defines another conditional probability distribution for the same variable (0.96 is the sensitivity of the test).

### 3.1.2. Representing the availability of information

In DANs there are two ways to indicate when a variable becomes observed: *always-observed variables* and *revelation arcs*. In both cases, we rely on the *no-forgetting hypothesis* (see Section 2.1.3).

**Always observed variables** The value of a chance variable declared as *always-observed* is known before making any decision. In Figure 3.1, the node *Symptom* is marked with a red thick border to indicate that it is always observed.

**Revelation links** Given a link  $X \rightarrow Y$ , such that  $Y$  is a chance variable, we can declare that certain values of  $X$  *reveal* the value of  $Y$ ; we then say that  $X \rightarrow Y$  is a *revelation link*. In general,  $X$  is a decision node, but it might also be a chance node; in this case, it means that  $Y$  is known only if a fortuitous event  $X$  occurs and the value of  $X$  is observed. The *revelation conditions* are the set of values of  $X$  that reveal  $Y$ . For the link *Dec: Blood test*  $\rightarrow$  *Blood test result* for the DAN of the diabetes problem, the value *yes* (i.e., the decision to do the test) reveals the result of the test, but the value *no* (not to do the test) does not. It is contradictory to declare  $x$  as a revelation condition for  $Y$  when there is a total restriction  $(x, Y)$ .

### 3.1.3. The meaning of links

In DANs a link  $X \rightarrow Y$  may have five meanings:

1. *Causal influence:* For example, the links *Diabetes*  $\rightarrow$  *Symptom* and *Diabetes*  $\rightarrow$  *Blood test result* mean that the presence of the disease affects the probability of having the symptom and the outcomes of this test, respectively.  $Y$  must be a chance node.
2. *Functional dependence:* For example, the four links pointing at *Cost of blood test*, *Cost of urine test*, and *Quality of life* denote which variables affect directly the decision maker's preferences.  $Y$  must be a utility node.

3. *Temporal order*: For example, the link  $Dec: Blood\ Test \rightarrow Therapy$  indicates which decision is made first.  $Y$  must be a decision.
4. *Revelation*: In general  $X$  is a decision, but it may also be a chance node, as explained above;  $Y$  must be a chance node.
5. *Restriction*:  $X$  must be a chance node or a decision, because utility nodes do not have children;  $Y$  must be a decision or a chance node.

The first three meanings are the same as in IDs. Revelation links in DANs replace information links in IDs. Restrictions are a novelty of DANs with respect to IDs.

### 3.1.4. Symmetric DANs

A DT is *symmetric* if (1) every path from the root to a leaf has the same variables and in the same order, and (2) every node representing the variable  $X$  has an outgoing branch for each value  $x$ . A DAN is *symmetric* if (1) it has no restrictions, (2) if a value of a variable  $X$  reveals  $Y$ , then all the other values of  $X$  reveal  $Y$ , and if (3) a directed path connects all the decisions.

The third condition induces a total ordering of the decisions in the DAN:  $D_1 < \dots < D_n$ . The first and second conditions imply that a symmetric DAN can be converted into an equivalent symmetric DT because when expanding the DT, the set of candidate chance nodes is the same for every branch of the tree. Therefore the first condition for a symmetric DT is always met. The second condition of symmetric DTs is guaranteed by the first condition of symmetric DANs.

The properties of symmetric DANs are very similar to those of IDs (see Section 2.1.3). A DAN with  $n$  decisions  $\{D_1, \dots, D_n\}$  induces a partition of  $\mathbf{C}$ , the set of chance variables, into  $n + 1$  disjoint subsets,  $\{\mathbf{C}_0, \dots, \mathbf{C}_n\}$  such that a chance variable  $X$  belongs to  $\mathbf{C}_i$ , with  $0 \leq i < n$ , if and only if the decision  $D_i$  reveals  $X$  and no other decision  $D_j$  with  $j < i$  reveals it while with IDs the condition is that there is a link  $C \rightarrow D_{i+1}$  but no link  $X \rightarrow D_j$  to a previous decision  $D_j$  ( $j \leq i$ ). Both in a DAN and in an ID,  $\mathbf{C}_0$  is the set of variables observed before making the first decision,  $\mathbf{C}_i$  is the set of variables observed after decision  $D_i$  and before  $D_{i+1}$ , and  $\mathbf{C}_n$  is the set of unobservable variables. Therefore the transformation of a symmetric DAN into an ID is straightforward: for every chance variable  $X$ , if  $X$  is always-observed, i.e., if  $X \in \mathbf{C}_0$ , draw a link  $X \rightarrow D_1$ ; if  $X$  is revealed by decision  $D_i$  (and not revealed by any previous decision), i.e., if  $X \in \mathbf{C}_i$ , draw a link  $X \rightarrow D_{i+1}$ . The conversion of an ID into a symmetric DAN is analogous.

As a consequence of this close relation between the two types of networks, any algorithm for IDs can be used to evaluate symmetric DANs. In general, it is not even necessary to convert the DAN into an ID; it is enough to infer the partition  $\{\mathbf{C}_0, \dots, \mathbf{C}_n\}$  from the graph of the DAN. For example, the variable elimination algorithm (Jensen and Nielsen, 2007; Luque and Díez, 2010), which consists essentially of applying Equations 2.9 and 2.11, is identical in both cases. The arc reversal

**Table 3.1.:** Main features of several methods for representing decision problems.

	dummy states	total order	inform. links
IDs (Howard and Matheson, 1984)	yes	yes	yes
SHM IDs (Smith et al., 1993)	no	yes	yes
AIDs (Nielsen and Jensen, 2000)	no	yes	yes
SVNs (Shenoy, 2000)	no	yes	yes
UIDs (Jensen and Vomlelová, 2002)	yes	no	no
SIDs (Jensen et al., 2006)	no	no	yes
DANs (Díez et al., 2012)	no	no	no

algorithm (Olmsted, 1983; Shachter, 1986) can also be applied to DANs: the only difference is that in an ID the first nodes to be removed (i.e.,  $\mathbf{C}_n$ ) are those that have no outgoing information links, while in a DAN they are those that have no incoming revelation links.

### 3.1.5. DANs vs. existing formalisms

In this section we examine six formalisms for decision analysis with DANs. First we compare DANs with IDs, which are the standard framework for decision analysis, and then we analyze six formalisms for asymmetric decision problems.

Table 3.1 summarizes the main features of each method. The first column indicates whether a formalism needs dummy states to “symmetrize” the problems containing structural asymmetries, as explained above. IDs and UIDs suffer from this drawback. The second column specifies which formalisms require a total ordering of the decisions. Obviously, these formalisms are inadequate for modeling problems having order asymmetry. Finally, the third column shows which methods need information links. These links, which indicate whether a variable is known for a decision, complicate the graph when there are many sequences of decisions that can reveal a variable.

**DANs vs. IDs** We have shown that symmetric DANs are equivalent to IDs. Therefore, they are equally suited for representing symmetric problems. As discussed earlier, IDs need dummy states or dummy utility nodes to handle structural asymmetry and order asymmetry poses great complications that grow exponentially with the number of decisions. In contrast, DANs do not need dummy states, nor dummy utility nodes, the size of the model is proportional to the number of decisions, and adding new decisions to the model is straightforward.

**DANs vs. other formalisms for asymmetric decision problems** UIDs and DANs share the advantage of not needing information links, but differ from each other in

that UIDs are inappropriate for problems having structural asymmetry. First, they can not represent restrictions. Second, in a UID an observable variable becomes observed when all its ancestral decisions have been made, independently of the option chosen for each decision; in a DAN a revelation link  $X \rightarrow Y$  indicates which values of  $X$  reveal  $Y$ . For example, the decision to do a test reveals its result, while the decision not to do it does not reveal anything. This property cannot be modeled directly with a UID: it would be necessary to add a dummy state to the result-of-test variable, as explained in Section 2.2.

The five frameworks proposed previously for asymmetric problems have been illustrated only with toy problems, which in several cases were designed just to illustrate the strengths of the new formalism. None of these formalisms has tried to solve all the asymmetric problems proposed previously in the literature. Additionally, no software tool, either commercial or open-source, has implemented them any of these formalisms.<sup>1</sup> These limitations may explain why none of these frameworks, proposed as an alternative to IDs, has been used to build any real-world application.

In contrast, DANs were developed with a complex medical problem in mind (the mediastinal staging of non-small cell lung cancer) based on several tests, which is very similar to the  $n$ -test problem. Later another DAN was built for deciding when to implant a knee prosthesis (León, 2011) These networks are available at [www.ProbModelXML.org/networks](http://www.ProbModelXML.org/networks), together with other DANs for the problems proposed previously in the literature on asymmetric decision formalisms: the used car buyer problem (Howard, 1984), the reactor problem (Covaliu and Oliver, 1995), the diabetes problem (Demirer and Shenoy, 2006), the dating problem (Nielsen and Jensen, 2000), and the king's problem (Jensen and Vomlelová, 2002); all these networks were built with **OpenMarkov**, an open source software tool that supports the edition and evaluation of DANs. Finally, DANs can also solve other real-world problems; for example, troubleshooting different types of devices (Heckerman et al., 1995), which is very similar to the  $n$ -test problem.

Decision analysis networks (DANs) are a new formalism to model and solve asymmetric decision problems, as described in Section 3.1. The author of this thesis began working on DANs after the formalism had been defined, and had the objective to develop inference algorithms to evaluate DANs. The first method proposed for the evaluation of DANs was the conversion to a decision tree (Section 3.2) that could be evaluated using the roll-back algorithm (Raiffa, 1968). We were aware at the moment of proposing such conversion that it was not the most efficient approach, but we decided it had the added value of defining the semantics of DANs for those familiar with decision trees. The need for more efficient algorithms led us to the development of the algorithm described in Section 3.3.

---

<sup>1</sup>UIDs are a partial exception to this assertion: they were implemented in Elvira (Jensen and Vomlelová, 2002; Ahlmann-Olsen et al., 2009; Elvira Consortium, 2002; Luque et al., 2008) with the goal of measuring the efficiency of some inference algorithms; but Elvira's graphical user interface for building UIDs is not mature, and no document explains how to build UIDs in Elvira.

## 3.2. Equivalent decision tree

In this section we describe the conversion algorithm from a DAN to a decision tree (DT). In the first phase, the algorithm builds the structure of the tree and then in the second it assigns probability and utility values. Then the DT can be evaluated with the standard roll-back algorithm (Raiffa, 1968).

We denote each node in the tree by the variable that it represents — for instance  $X$  — and each branch departing from this node by its associated value,  $x$ .

### 3.2.1. First phase: structure of the tree

**Initialization:** The tree is initialized by placing the original DAN as the only node in the tree.

**Invariant:** The leaves of the tree represent DANs or utility values. The rest of nodes represent either chance or decision variables, or a meta-decision (described below). The branches represent either the assignment of a value to the variable whose node is the source of the branch or an assignment to a meta-decision.

**Termination condition:** All the leaves in the tree are utility nodes (values).

#### 3.2.1.1. Recursion

There is a base case and four possible recursive cases. They are considered in order and if one holds, the next cases are ignored.

**Base case: The DAN only has utility nodes** Replace the DAN with a utility node whose value is the sum of the utility nodes remaining in the DAN.

The algorithm always terminates: in cases 1, 2, and 4 the expansion of a DAN with  $n$  nodes generates a finite number of DANs with  $n - 1$  nodes; in case 3, the new DANs still have  $n$  nodes, but every one is in case 2, which ensures the elimination of a decision node when expanding each new DAN; therefore, case 3 can only occur a finite number of times.

**Recursive case 1. The DAN has always-observed nodes** Select always-observed node  $X$  that is not a descendant of any other always-observed node. Replace the DAN in the tree with  $X$ . For each value  $x$  of  $X$ , create a copy of the DAN and:

1. delete the node  $X$  from the new DAN;
2. if  $x$  reveals the value of other variables, mark them as always-observed in the new DAN;

3. if there is a total restriction  $(x, Y)$ , remove  $Y$  and all the descendants of  $Y$  that are not descendants of another decision;
4. if there is a restriction  $(x, y)$ , remove  $y$  from the domain of  $Y$ ; and
5. put the new DAN in a branch  $x$  outgoing from node  $X$ .

**Recursive case 2. The DAN has a decision  $D$  that is an ancestor of all the other decisions** Put  $D$  as a node in the tree, create a copy of the DAN for each value  $d$  of  $D$ , and execute the same five steps as in the previous case, with one exception: in step 2, if  $d$  reveals a variable  $Y$  that is a descendant of other decisions  $\{D'_1, \dots, D'_n\}$ , draw a link  $D'_i \rightarrow Y$  for  $i \in \{1, \dots, n\}$  and declare that all values  $d'_i$  of  $D'_i$  reveal  $Y$ .<sup>2</sup> If  $d$  reveals a variable  $Y$  that is not a descendant of any other decision, then declare  $Y$  as always observed, as in case 1.

**Recursive case 3. There are  $n$  decisions  $\{D_1, \dots, D_n\}$  such that there is no directed path between any two of them** In the tree, replace the DAN with a meta-decision node indicating which decision will be made first. Create  $n$  branches outgoing from the meta-decision node, each branch  $i$  labeled with  $D_i$ . In the node at the end of each branch  $i$ , put a copy of the DAN where  $n - 1$  links  $D_i \rightarrow D_j$  ( $j \neq i$ ) have been drawn to indicate that  $D_i$  will be the first decision.

**Recursive case 4. The DAN has chance nodes (that are not always-observed)** Select one chance node randomly, say  $X$ , and replace the DAN with a node  $X$  in the tree. For each value  $x$  of  $X$ , create a copy of the DAN, remove  $X$ , and put the new DAN at a branch  $x$ .<sup>3</sup>

**Example** Consider the conversion of the DAN for the diabetes problem (3.1) into a DT. The tree is initialized by putting the original DAN in the root of the tree. In the first iteration, the DAN has one always-observed node, *Symptom* (case 1); then this node replaces the original DAN as the root of the tree (represented by  $S$  in 3.6); two branches are added, one for *present* ( $+s$ ) and one for *absent* ( $-s$ ). Each leaf is a DAN in which the node *Symptom* has been removed, and has two decisions (*Dec: Blood Test* and *Dec: Urine Test*) with no directed path between them (case 3). Therefore we put in the tree a meta-decision node *OD* (order of the decisions) with two branches. The DAN at the *bt* branch has a link *Dec: Blood Test*  $\rightarrow$  *Dec: Urine Test*, which indicates which is the first decision (case 2). This DAN is later replaced with a node *BT* (*Dec: Blood Test*) having two branches:  $+bt$  (do test) and  $-bt$

<sup>2</sup>The reason for this exception is that if  $Y$  is a descendant of  $D'_i$  then this decision can be interpreted as a cause of  $Y$  and therefore  $Y$  cannot be known before making  $D'_i$ .

<sup>3</sup>In this case, the restrictions and revelation conditions are ignored. In fact, if  $X$  can never be observed, then no link  $X \rightarrow Y$  should have total restrictions or revelation conditions.


(do not test). In the DAN at the  $\neg bt$  branch the node *Dec: Blood Test* has been removed when putting this variable in the tree; *Blood test result* has been removed because of the total restriction (*Dec: Blood Test = no, Blood test result*). In the DAN at the  $+bt$  branch the node *Dec: Blood Test* has been also removed, but the node *Blood test result* remains; it is now marked as always observed because of the revelation condition the link *Dec: Blood test  $\rightarrow$  Blood test result*. This DAN gives rise to the branches  $+b$  (positive test result) and  $\neg b$  (negative) in Figure 3.6. Each of the three *UT* nodes in this figure is the result of expanding a DAN in which *Dec: Urine Test* was the first decision to make (case 2). When the decision is  $+ut$  (do test), the variable *U* is marked as always-observed and in the subsequent expansion of the DAN it is put in the corresponding branch of the tree. When the decision is  $\neg ut$  (do not test), *U* is removed. The node *Th* (*Therapy*) is the result of expanding a DAN in which this was the only decision (again case 3). The node *D* (*Diabetes*) results from a DAN having one chance node and three utility nodes (case 4). The expansion of the DAN that was at this node, *D*, generates the two utility nodes shown in this figure (base case).

The resulting tree could be further pruned of suboptimal branches by introducing a priority in cases 2 and 3 for decisions revealing a set of variables that is a proper superset of the sets of variables revealed by other decisions. The formal justification is the same as in the case of unconstrained influence diagrams (Jensen and Vomlelová, 2002). For example, in Figure 3.6, the decision *Therapy* does not reveal any variable. Therefore, even if the links *Dec: Blood Test  $\rightarrow$  Therapy* and *Dec: Urine Test  $\rightarrow$  Therapy* were not present, the node *OD* in Figure 3.1 should have only two outgoing branches, *bt* and *ut*. Making the *Therapy* decision first would never result in a higher expected utility, as additional information can only improve the expected utility.

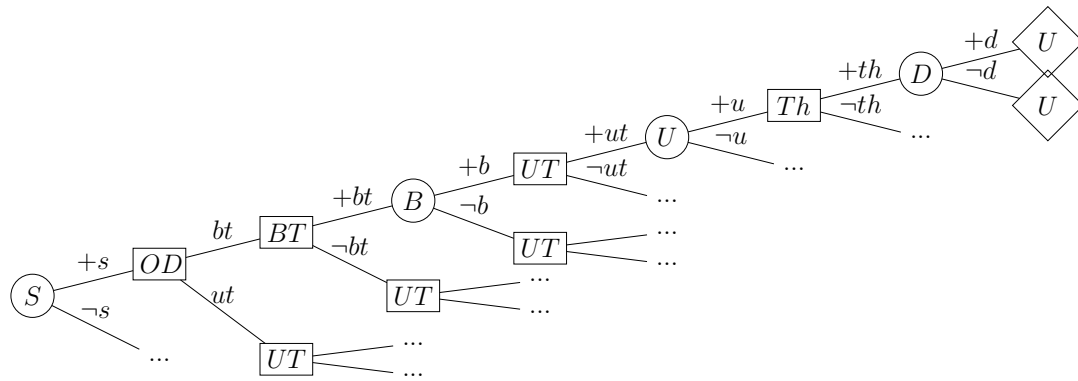
### 3.2.2. Second phase: assignment of utilities and conditional probabilities

The second phase of the algorithm consists of assigning the numerical parameters of the DT. Each path from the root node to a leaf node defines a *scenario*, i.e., a configuration of the variables in that path. The utility of a leaf is the sum of the utility functions in that scenario; for example, the utility for the upper utility node in 3.6 is  $u_1(+bt) + u_2(+ut) + u_3(+th,+d)$ , where  $u_1$ ,  $u_2$ , and  $u_3$  are the utility functions for the nodes *Cost of blood test*, *Cost of urine test*, and *Cost of therapy*, respectively. The probability of a scenario is the product of the conditional probabilities involved in it; for example, the probability of the upper scenario is  $P(+s \mid +d) \cdot P(+b \mid +d, +bt) \cdot P(+u \mid +d, +ut) \cdot P(+d)$ . Initially we compute the probability of each branch as the sum of the probabilities of the scenarios containing





**Figure 3.5.:** Decomposition of a DAN representing the diabetes problem during the execution of the conversion into a DT.



**Figure 3.6.:** A decision tree for the diabetes problem, equivalent to the DAN in Figure. 3.1.

it; then, we normalize the probabilities of the branches going out from each chance node. This way we obtain an equivalent DT for the DAN.


The above algorithm may generate different DTs for one DAN because in cases 1 and 4 there may be several nodes to select. However, these trees only differ in the order of the chance variables placed between two consecutive decisions. These DTs are equivalent in the sense that there is a one-to-one correspondence between the decision nodes of each pair of trees, and the optimal policies and the maximum expected utility are the same for all these trees.

### 3.2.3. Evaluation of the tree

This tree can be evaluated with the standard roll-back algorithm (Raiffa, 1968), which proceeds from the leaves to the root: the utility of a decision node is the maximum of the utilities of the nodes at its branches, and the utility of a chance node is the average of the utilities of the nodes at its branches, weighted by their probabilities.

## 3.3. Recursive decomposition into symmetric DANs

The algorithm in 3.2 converts a DAN into an equivalent DT recursively decomposing the original DAN, until it is completely decomposed and all the nodes in the tree represent single variables. However, the recursion can stop once the DANs in the leaves are symmetric, and therefore evaluable using any algorithm for IDs, which are much more efficient in general. For example, when expanding the tree in Figure 3.6 the DANs at the branches  $+ut$  and  $-ut$  are symmetric.



**Figure 3.7.:** Decomposition of the diabetes DAN into symmetric DANs.

Besides, we don't need to expand  $S$ , since it induces no asymmetry and therefore DANs containing  $S$  are still symmetric. Therefore, we may represent it as a tree whose leaves are symmetric DANs, as shown in Figure 3.7.

The recursive algorithm we propose here is similar to the expansion of a DT, but has two important differences: recursion stops when DANs are symmetric and only variables inducing asymmetry multiply the number of calls to the recursive function.

Let  $\mathbf{V}$  be the set of observed variables. We denote with  $\mathbf{V}_A$  the subset of observed variables that induce asymmetry by imposing restrictions on other variables or conditionally revealing other variables.

The recursive function receives a DAN as input parameter and returns a pair  $\{U(\mathbf{V}), P(\mathbf{V})\}$ , where  $U(\mathbf{V})$  is the expected utility potential for the set of variables  $\mathbf{V}$ , and  $P(\mathbf{V})$  its joint probability.

**Base case. The DAN is symmetric.** In this case, and both its expected utility  $U(\mathbf{V})$ , and the joint probability  $P(\mathbf{V})$ <sup>4</sup> can be calculated using any inference algorithm for IDs. There are two exhaustive<sup>5</sup> recursive cases that are considered in order in the recursive function.

**Recursive case 1.  $\mathbf{V}_A$ , the set of observed chance variables that impose restrictions or conditionally reveal other variables, is not empty** In this case, we select a variable  $X \in \mathbf{V}_A$  that is not a descendant of any other variable in  $\mathbf{V}_A$

---

<sup>4</sup> $P(\mathbf{V})$  can be calculated efficiently on a Bayesian network containing  $\mathbf{V}$  and their ancestors and by marginalizing out the unobserved variables.

<sup>5</sup>If we are not in the base case, we must be in one of these cases

and we call the recursive function for each value  $x$  of  $X$  passing it a DAN where  $X$  has been instantiated w.r.t  $x$  as input parameter. By instantiating  $X$  w.r.t  $x$ , we obtain a reduced DAN where  $X$  has been removed and in accordance with the restrictions of the form  $(x, y)$  domain of the  $Y$ 's have been reduced; if there was a total restriction  $(x, Y)$ ,  $Y$  and all its the descendants that are not descendants of another decision are removed; and the variables revealed by  $x$ ,  $\mathbf{V}_x$ , have been marked as always observed. The recursive calls return  $U_x(\mathbf{V}')$  and  $P_x(\mathbf{V}')$  for each value  $x$  of  $X$ , where  $\mathbf{V}' = \{\mathbf{V} \setminus X \cup \mathbf{V}_x\}$  is the set of observed variables in the reduced DAN. The expected utility for each  $x$  of  $X$ ,  $U_x(\mathbf{V})$ , is calculated multiplying  $P_x(\mathbf{V}_x | \mathbf{V})$  and  $U_x(\mathbf{V}')$  and marginalizing out  $\mathbf{V}_x$ .

The new  $P(\mathbf{V})$  is calculated marginalizing out  $\mathbf{V}_x$  from  $P_x(\mathbf{V}')$  for each value of  $x$ .

$$U_x(\mathbf{V}) = \sum_{\mathbf{V}_x} P_x(\mathbf{V}_x | \mathbf{V}) U_x(\mathbf{V}')$$

**Recursive case 2.**  $\mathbf{D} = \{D_1, \dots, D_n\}$ , **the set of initial decisions is not empty:**  $D_i \in \mathbf{D}$  **if and only if there exists no direct path from any other decision to  $D_i$**  For each  $D$  in  $\mathbf{D}$ , we do the following: If  $D$  imposes restrictions or conditionally reveals other variables, we call the recursive function for each value  $d$  of  $D$  passing it a DAN where  $D$  has been instantiated w.r.t  $d$  as input parameter. In contrast with the instantiation of chance variables, variable  $V_d$  conditionally revealed by  $d$  is only truly revealed (and therefore marked as always observed) if it has no ancestral decisions. Revelation links are added from all its ancestral decisions to  $V_d$  otherwise. The recursive calls return  $U_d(\mathbf{V}')$  and  $P_d(\mathbf{V}')$  for each value  $d$  of  $D$ , where  $\mathbf{V}' = \{\mathbf{V} \cup \mathbf{V}_d\}$ <sup>6</sup> is the set of observed variables in the reduced DAN. The expected utility for each  $d$  of  $D$ ,  $U_d(\mathbf{V})$ , is calculated multiplying the probability of  $P_d(\mathbf{V}_d | \mathbf{V})$ , and  $U_d(\mathbf{V}')$  and marginalizing out  $\mathbf{V}_d$ .

$$U_d(\mathbf{V}) = \sum_{\mathbf{V}_d} P_d(\mathbf{V}_d | \mathbf{V}) U_d(\mathbf{V}')$$

The utility for a decision node  $U_D(\mathbf{V})$  is the value  $U_d(\mathbf{V})$  that maximizes utility for each configuration of  $\mathbf{V}$ .

$$U_D(\mathbf{V}) = \max_{d \in D} U_d(\mathbf{V})$$

If on the contrary,  $D$  neither imposes restrictions nor conditionally reveals any variable, we call the recursive function only once passing it a DAN as input parameter

---

<sup>6</sup>Note that in this case,  $\mathbf{V}_d$  is the set of variables truly revealed by  $d$

where:  $D$  has been marked as always observed and for each variable  $R_D$  revealed by  $D$ , if it does not have ancestral decisions it has been marked as always observed and otherwise revelation links have been added from all its ancestral decisions to  $R_D$ . The recursive call returns in this case the expected utility  $U(\mathbf{V}', D)$ .  $U_D(\mathbf{V})$  is again computed by choosing the value that maximizes the utility for each configuration of  $\mathbf{V}$ .

---

**Algorithm 1:** Asymmetric decomposition.

---

**Function** Evaluate( $DAN$ ):

```

Result:  $\{U(\mathbf{V}), P(\mathbf{V})\}$ 
if  $DAN$  is symmetric then
     $\{U(\mathbf{V}), P(\mathbf{V})\} = \text{IDAlgorithm}(DAN)$ ;
else if  $\mathbf{V}_A \neq \emptyset$  then
    Pick  $X \in \mathbf{V}_A$ , where  $\neg \exists Y \in \mathbf{V}_A, Y \rightarrow X$ ;
    foreach  $x$  of  $X$  do
         $DAN_x = \text{Instantiate}(DAN, x, X)$ ;
         $\{U_x(\mathbf{V}'), P_x(\mathbf{V}')\} = \text{Evaluate}(DAN_x)$ ;
         $U_x(\mathbf{V}) = \sum_{\mathbf{V}_x} P_x(\mathbf{V}_x | \mathbf{V}) U_x(\mathbf{V}')$ ;
         $P_x(\mathbf{V}) = \sum_{\mathbf{V}_x} P_x(\mathbf{V}')$ ;
    else if  $\mathbf{D} \neq \emptyset$  then
        foreach  $D$  in  $\mathbf{D}$  do
            if  $D$  induces asymmetry then
                foreach  $d$  in  $D$  do
                     $DAN_d = \text{Instantiate}(DAN, d, D)$ ;
                     $\{U_d(\mathbf{V}'), P_d(\mathbf{V}')\} = \text{Evaluate}(DAN_d)$ ;
                     $U_d(\mathbf{V}) = \sum_{\mathbf{V}_d} P_d(\mathbf{V}_d | \mathbf{V}) U_d(\mathbf{V}')$ ;
                     $P(\mathbf{V}) = \sum_{\mathbf{V}_d} P_d(\mathbf{V}')$ ;
                 $U_D(\mathbf{V}) = \max_{d \in D} U_d(\mathbf{V})$ ;
            else
                 $DAN_D = DAN$  where  $D$  and  $\mathbf{V}_D$  have been revealed;
                 $\{U_D(\mathbf{V}'), P_D(\mathbf{V}')\} = \text{Evaluate}(DAN_D)$ ;
                 $U_D(\mathbf{V}) = \sum_{\mathbf{V}_D} P_D(\mathbf{V}_D | \mathbf{V}) U_D(\mathbf{V}')$ ;
                 $P(\mathbf{V}) = \sum_{\mathbf{V}_D} P_D(\mathbf{V}')$ ;
             $U(\mathbf{V}) = \max_{D \in \mathbf{D}} U_D(\mathbf{V})$ ;
    return  $\{U(\mathbf{V}), P(\mathbf{V})\}$ 

```

---

Having calculated  $U_D(\mathbf{V})$  for each  $D$  in  $\mathbf{D}$  implies that we have calculated every possible ordering of the decisions in  $\mathbf{D}$ . We can now calculate  $U_{\mathbf{D}}(\mathbf{V})$  maximizing the utilities of  $U_D(\mathbf{V})$  for each configuration of  $\mathbf{V}$ .

For example, at the upper node  $UT$  in Figure 3.7  $\mathbf{V} = \{S, BR\}$ ; the evaluation of  $DAN_1$  and  $DAN_2$  returns the utility functions  $U_1(s, br)$  and  $U_2(s, br)$ , respectively. The expected utility for that node is  $U_{UT(+bt)}(s, br) = \max(U_1(s, br), U_2(s, br))$ . The only variable observed at the node  $UT$  in the branch  $(bt, -bt)$  is  $S$ ; therefore, the utility at this node is  $U_{BT(-bt)}(s) = \max(U_3(s), U_4(s))$ . The utility at the root is  $U_{OD}(s) = \max(U_{BT}(s), U_{UT}(s))$ .

### 3.3.1. Complexity

The upper bound for the number of recursive calls is  $O(s^m \cdot n!)$ , where  $m$  is the number of variables that either impose restrictions or conditionally reveal variables,  $s$  is the highest cardinality of those variables, and  $n$  is the number of unordered decisions (assuming that there is only one level of unordered decisions). In the  $n$ -test problem, there are  $n$  unordered decisions, each of which imposes total restrictions on the corresponding chance variable; therefore, the complexity is  $O(2^n \cdot n!)$ , which makes the problem intractable even for relatively small values of  $n$ .

### 3.3.2. Reducing complexity reusing intermediate results

However, the number of recursive calls can be drastically reduced exploiting the fact that the order in which past decisions have been taken is not relevant for the expected utility of future decisions. For example, the reduced DAN resulting from taking the decision  $BT = +bt$  first and then  $UT = +ut$  is identical to the one resulting from taking the decisions in the reverse order. In fact, all DANs in 3.7 are repeated twice:  $DAN_1 = DAN_5$ ,  $DAN_2 = DAN_7$ ,  $DAN_3 = DAN_6$ , and  $DAN_4 = DAN_8$ . By storing the expected utilities of the reduced DANs in a cache we can bring the number of recursive calls from  $O(s^m \cdot n!)$  down to  $O(s^m)$ . The number of intermediate results to store is  $O(s^m)$ .

The advantage of this algorithm is that it profits from the reduction of the problem when restrictions are applied. This is also true for the size of the cache. For example, in the  $n$ -test problem the cache contains a single table with  $2^n$  entries; the rest of tables are smaller because the variables corresponding to the results of tests that are not performed are not present.

## 3.4. Empirical comparison of evaluation algorithms

### 3.4.1. Experiments

In our experiments we have used several DANs that model the some of the problems proposed in the literature: the used car buyer problem (Howard, 1984), the reactor

**Table 3.2.:** Computational time for the different algorithms (ms).

	DT	S-DAG	RD
Reactor	64	21	50
Used car buyer	102	26	55
Dating	122	83	46
King	3944	117	183
Mediastinet	84720	353	370
Diabetes (2 tests)	128	33	40
3-tests	825	73	89
4-tests	3457	240	187
5-tests	55799	633	419

problem (Covaliu and Oliver, 1995), the dating problem (Nielsen and Jensen, 2000), the king problem (Jensen and Vomlelová, 2002), the diabetes problem (Demirer and Shenoy, 2006), and Mediastinet. The latter is a DAN for the mediastinal staging of non-small cell lung cancer, which is an adaptation of the ID Mediastinet (Luque et al., 2009). The main difference is that in the DAN there is not a total ordering among the 5 possible tests; there is only a partial ordering, due to some technical requirements. Another difference is that all the dummy states have been removed, except for the result a test that affects the sensitivity and specificity of the three subsequent tests. These networks are available at [www.ProbModelXML.org/networks](http://www.ProbModelXML.org/networks).

We have also tested the  $n$ -test problem for different values of  $n$ . This problem consists in deciding how to treat a patient that may suffer from a certain disease; after an initial examination of the symptoms, the doctor may order some of the  $n$  available tests, each one having a cost; each test can be performed at most once and its result will be known immediately. In our experiments we have used the simplest version of the problem, in which there is only one symptom and all the variables are binary, i.e., the disease and the symptom are either present or absent, and the result of each test is either positive or negative. The diabetes problem is a particular case in which there are only two tests ( $n = 2$ ).

In the experiments we have compared the following algorithms:

1. equivalent decision tree (DT);
2. recursive decomposition into symmetric DANs (RD); the symmetric DANs in the base case were evaluated using variable elimination;
3. variable elimination on an S-DAG. This algorithm is the result of applying S-DAGs (Jensen and Vomlelová, 2002) for the evaluation of DANs.

The algorithms have been implemented in `OpenMarkov`. The experiments were conducted on a PC with an Intel Core i7-2670QM CPU @ 2.20 GHz using Java's default heap size of 1GB.

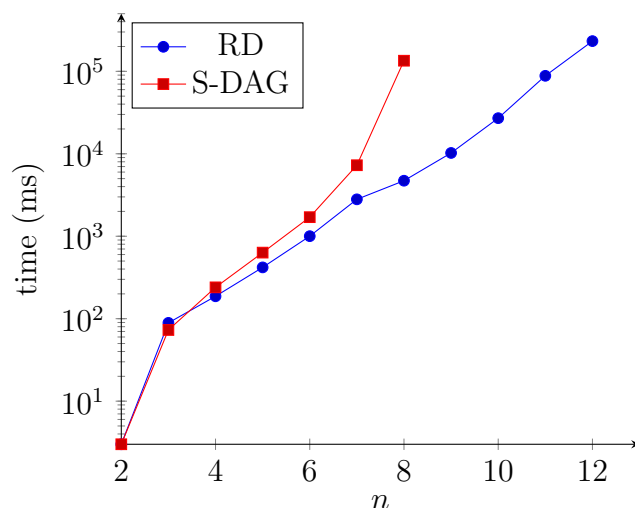


Figure 3.8.: Computational time for the  $n$ -test problem.

### 3.4.2. Results

The time necessary to evaluate the five toy examples proposed in the literature (the first five networks in Table 3.4.1) is small for all the algorithms. Therefore we can use any of them in practice without noticing much difference. However, for the networks that have 5 or more tests, including the real-world application Mediastinet, the algorithms proposed in this paper are much faster than the expansion of the DT. S-DAG is more efficient than RD except for the dating problem, which has strong structural asymmetry: one half of the network models the decision over where to go if Emily accepts the invitation and another one for deciding what to do if she does not; once we know her answer, the RD algorithm prunes almost half of the nodes, and for this reason it is more efficient than S-DAG. In turn, RD is more efficient than S-DAG because the coalescence compensates for the compacted S-DAG and the algorithm takes profit of problem reduction thanks to asymmetry.

The space complexity of these algorithms is of the same order of their time complexity. In practice, the space complexity may be more relevant than the time, especially when the DAN is used off-line to find out the optimal policy. In the case of the  $n$ -test problem, the DT algorithm runs out of memory at  $n = 6$ , S-DAG at  $n = 9$  and RD at  $n = 13$ . The advantage in space requirements of RD over S-DAG is that the former removes some variables as a consequence of the restrictions in the DAN, which reduces the size of the probability tables, while the latter treats the restrictions by adding dummy states that increase the size of the tables.



## 4. Markov processes with atemporal decisions

*History is a cyclic poem written by time upon the memories of man.*

Percy Bysshe Shelley

### 4.1. Introduction

Ever since the influence diagrams were proposed by Howard and Matheson (1984), probabilistic graphical models have been used for decision making under uncertainty in a wide range of areas, including medicine Boutilier (2005). Unfortunately, PGMs are still relatively unknown in the field of economic evaluation in medicine, a field that has long embraced the need for decision modeling Weinstein et al. (2003).

Typically, decision trees Raiffa (1968) were used to build and analyze economic models in medicine. Decision trees however, suffer severe limitations to model medical problems; mainly, they do not explicitly consider time and their size grows exponentially with the number of variables in the model, as discussed in Section 2.3.2.1. Markov models Beck and Pauker (1983) (and more generally state transition models) added a temporal dimension, thus overcoming some of the limitations of decision trees. They gained popularity rapidly, becoming the most popular instrument for health economic evaluation. In the last decade, some have advocated for the adoption of discrete event simulation (DES) modeling (see Section 2.3.2.3), but Markov models remain, in spite of their shortcomings, the first option for modelers.

Given that current models are computer-based, the approach and software used to implement them is of vital importance. According to a review of the use of software for decision modeling as part of the National Institute for Health and Care Excellence (NICE) of England and Wales, technology appraisal process Tosh and Wailoo (2008), state transition models are most commonly implemented in spreadsheet applications, such as **Microsoft Excel**. The structure of the model, as well as its numerical parameters (for example, the transition probabilities) and its evaluation (for example, the calculation of utilities) are encoded as formulas in the spreadsheet. Sensitivity analysis is usually implemented by a set of macros in a high-level programming language, such as **Visual Basic for Microsoft Excel**.

After spreadsheets, **TreeAge Pro** is the most used software package to build state transition models. It offers the possibility of building and evaluating decision trees

and Markov cycle trees using its GUI. However, the exponential growth of the size of trees and the restrictions imposed by the software make complex models unwieldy. Script-based statistical packages such as R, BUGS or STATA offer higher flexibility and a myriad of statistical functions, but require higher expertise and lack the clarity of graphical representations. DES models are usually built and evaluated using proprietary software (Arena® and SIMUL8 are the most popular), but remain marginal, due probably to lack of familiarity, the higher level of expertise required, and a widespread belief that its disadvantages outweigh its advantages in most cases.

Probabilistic graphical models, thanks to their graphical structure, allow for the design of intuitive qualitative models, capturing the probabilistic and informational dependencies, thus helping the construction, inference, and understanding of complex decision problems. Influence diagrams, the first and most popular PGM for decision analysis, are very useful for a variety of decision problems, as explained in Section 2.1.3, but are not appropriate to model decision problems involving a continuous risk over time, or in which events may occur more than once, or when the utility of an outcome depends on when it occurs.

In this chapter we propose a new formalism, called Markov Processes with Atemporal Decisions (MPADs), as a particular type of PGM. MPADs are able to encode state-transition models and are designed for cost-effectiveness analysis.

Support for building and evaluating MPADs has been implemented in **OpenMarkov**, an open-source software package for PGMs. It offers the possibility to perform cost-effectiveness analysis and several types of sensitivity analysis with MPADs using its graphical user interface. Several examples of MPADs replicating cost-effectiveness models in the literature are available online<sup>1</sup>.

**Running example: HIV therapy** In this chapter we use the evaluation of two therapies for HIV infection as a running example to illustrate the main properties of MPADs. The conceptual model was designed by Chancellor et al. (1997) to compare the cost-effectiveness of lamivudine plus zidovudine (combination therapy) with that of zidovudine alone (monotherapy). It is worth noting lamivudine was meant to be given for only two years. In this model the health state of the patients is classified into four possible states: *State A*, the least severe disease state, means that the CD4 count is between 200 and 500 cells/mm<sup>3</sup> whereas *State B* means less than 200 cells/mm<sup>3</sup>, *State C* represents AIDS, and *State D* represents death. The possible transitions are depicted as arrows in Figure 4.1. The cycle length is one year. The effectiveness was measured as the duration of life, without considering the quality of life.

---

<sup>1</sup>At <http://www.probmodelxml.org/networks>




Figure 4.1.: Transition diagram for Chancellor’s model

## 4.2. Definition

An MPAD consists of a graph, which defines the causal relations between the variables, and a set of potentials, namely probability distributions and utility functions. The numerical parameters of MPADs define the potentials.

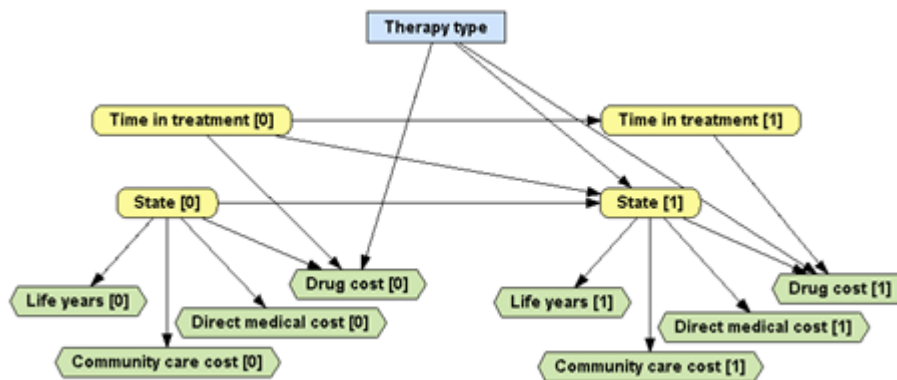
### 4.2.1. Structure

As mentioned in Section 2.1.1, a graph is an ordered pair  $G = (V, E)$  comprising a set of vertices or nodes  $V$  and a set of edges or links  $E$ . As in other PGMs, each node in an MPAD represents a variable and therefore both terms will be used interchangeably henceforth. The set of nodes in MPADs is divided into three types of nodes: decision, chance and utility<sup>2</sup>. Nodes representing properties that evolve over time have indexes, shown between square brackets, as shown in Figure 4.2. The subgraph that contains all the nodes  $X^t$  having the same temporal index  $t$  and the links between them is known as the  $t$ -th time slice; therefore there is a time slice for each cycle.

Decision nodes, drawn as rectangles, represent choices available to the decision maker. In Figure 4.2, *Therapy* represents the decision over which one of the two therapies under evaluation to apply.

Chance nodes, drawn as rounded rectangles, represent all the variables that are not under the decision maker’s direct control. In medical PGMs, they usually represent the patient’s features, such as sex, age and the health state. Figure 4.2 contains four chance variables. The variable *State [0]* represents the health state of the patient in cycle 0 (start of therapy) and can take four values, one for each health state: A, B, C and D. *State [1]* represents the state of the patient in cycle 1 (one year later) and has the same domain as *State [0]*. *Time in treatment [0]* and *Time in treatment [1]*

<sup>2</sup>“Utility” is the term used in the PGM literature, while “value” is more common in other fields, such as health economics.



**Figure 4.2.:** The first two cycles of an MPAD for the HIV example

are numeric variables that represent, for each cycle, for how long has the patient been under treatment.

Utility nodes represent the decision maker's payoffs, such as costs and health outcomes. They are drawn as hexagons. In Figure 4.2, *Life years [0]* and *Life years [1]* represent effectiveness, while the other utility nodes represent economic costs.

Chance and value nodes can be either temporal or atemporal; for example, gender is atemporal, whereas the health state and the costs incurred in each cycle are temporal. In contrast, decisions in MPADs are always atemporal, as they are made once and forever, generally at the start of the time represented by the model. This is a key difference between MPADs and dynamic influence diagrams Tatman and Shachter (1990), and leads to a completely different evaluation of the model. There must be a directed path that passes through all the decision nodes, defining the order in which decisions are made. The only exception to this condition is when the ordering of the decisions is not relevant Nielsen and Jensen (1999), namely when unordered decisions do not influence the expected utility for the rest of unordered decisions.

Nodes in an MPAD are connected by directed links. Links to chance variables denote causal influence. For example, *State [1]* in Figure 4.2 depends causally on *State [0]*, *Therapy* and *Time in treatment [1]*. Links to utility nodes reflect that the respective utility function depends on the node at the source of the link. For example, the effectiveness (*Life years*), the direct medical costs and the community care costs only depend on the current state. However, the drug cost also depends on the therapy type and on the time in treatment, because the patients being treated with combination therapy receive lamivudine only during the first two years of treatment. It is worth noting that, as it is common in temporal models, links cannot be drawn from a node in a time slice to a node in a previous time slice (or to an atemporal node).

We say the potential of a variable  $X$  is stationary after  $k$  if it remains unchanged after the  $k$ -th slice. A temporal model is stationary after  $k$  if all its potentials are

stationary after  $k$  and there is no other  $k' < k$  that makes the model stationary. Such a model admits a *compact* representation that contains only the first  $k$  slices. It is then possible to unroll the model up to its time horizon  $h$  by cloning the last slice  $h - k$  times. Usually, as it is the case in our example, temporal models are stationary after the first two time slices; therefore, the graph in Figure 4.2 can be interpreted as the compact representation of the model. A model that is stationary after  $k$  may contain variables whose potentials become stationary earlier, i.e., after  $k'$ , with  $k' < k$ . In this case, these nodes can be removed from the compact representation without losing information to produce the *concise* representation. Thus, in 4.2 we could remove the utility nodes from the second slice without loss of information. MPADs are usually defined specifying their concise representation.

It is important to stress the difference between a transition diagram Equation (4.1) and the graph of a PGM Equation (4.2). In a transition diagram each node represents a state and each link a transition, while in a PGM each node represents a variable. When an MPAD contains a variable representing states, the transition diagram can be understood as the inner structure of that variable. Another difference is that a transition diagram may contain cycles—in particular, an arrow connecting a node to itself means that the patient might remain in the same state—while the graph of a PGM cannot have cycles.

### 4.2.2. Numerical parameters

An MPAD contains two types of numerical parameters: probabilities and utilities. Each chance node defines a probability distribution conditioned on its parents in the graph. In our example, the probability distribution for the parentless node *State* [0] is not conditioned on any variable, and is therefore the a priori probability  $P(\textit{State} [0])$ . In a cohort interpretation of the model, this probability determines the proportion of the cohort in each state at the beginning of the treatment. As we assume that all the patients are initially in state A, the  $P(\textit{State} [0] = \textit{state A}) = 1$  and  $P(\textit{State} [0] \neq \textit{state A}) = 0$ .

*State* [1], whose parents are *State* [0], *Therapy*, and *Time in treatment* [0], has an associated conditional probability distribution,  $P(\textit{State} [1] \mid \textit{State} [0], \textit{Therapy}, \textit{Time in treatment} [1])$ , which can be defined as a four-dimensional matrix or as a probability tree. Analogously, each utility node defines a utility function that indicates how its value depends on its parents in the graph. In our example, the utility function for node *Life years* [0], whose only parent is *State* [0], can be represented by a one-dimensional matrix (a vector),  $U_{\textit{Life years}} [0](\textit{State} [0])$ , taking the value 1 when the patient is alive (states A, B and C) and 0 when he is dead (state D). Similarly, the function for *Drug cost* [0], which has three parents in the graph, is  $U_{\textit{Drug cost}} [0](\textit{State} [0], \textit{Therapy}, \textit{Time in treatment} [0])$ .

### 4.2.3. Criteria

MPADs support multiple decision criteria. In multi-criteria decision analysis with PGMs, each utility variable needs to be assigned a criterion. In fact, cost-effectiveness analysis is a special case of multi-criteria evaluation. In our example, the criterion for *Life years [0]* and *Life years [1]* is effectiveness, while for the other utility nodes it is cost. MPADs might also have a single criterion. Multi-criteria MPADs can be evaluated as unicriterion models by selecting one criterion as the target criterion and assigning a conversion factor from each of the other criteria to the target criterion.

## 4.3. Evaluation of MPADs

The goal of evaluation in influence diagrams is usually to find the strategy that will maximize global expected utility. However, in a multiple criteria environment, a strategy that maximizes a criterion does not necessarily maximize others and therefore an optimal strategy need not exist. In fact, our goal might be to maximize some criterion (effectiveness) and to minimize others (e.g. cost). Therefore, the goal of the evaluation of an MPAD is to calculate the global expected utility, which is the sum of the expected value of all the utility nodes, for all decisions and criteria. An optimal strategy for the decisions can be calculated if conversion factors are provided for the criteria, as explained in the previous section. In the case of CEA, once we know the expected cost and effectiveness of each intervention, we can calculate their respective net monetary benefit for a certain  $\lambda$  and pick the intervention that maximizes the NMB, as explained in Section 2.3.1.

### 4.3.1. Expansion of the MPAD

Given that MPADs are usually defined using concise or compact representations, the first step in the evaluation of an MPAD consists of building the expanded version. If the concise representation is used, the first step is to add the nodes and links necessary to get the compact representation. Once we have the compact representation, the model needs to be unrolled up to the desired time horizon by cloning the nodes in the last time slice and the links pointing at them as many times as necessary. An expanded representation of an MPAD is an influence diagram and can therefore be evaluated using influence diagram evaluation algorithms.

### 4.3.2. Exact inference

In exact inference, the probability of a node in the expanded graph is computed by multiplying its probability distribution by the joint probability of its parents. In a similar way, the expected value of a utility node is calculated by summing the

utility for each configuration of parents multiplied by their probability. It is, in fact, similar to cohort analysis (or cohort simulation) in Markov models. Efficient calculation of expected utility can be done using inference algorithms such as variable elimination Luque and Díez (2004). The efficiency of variable elimination may be highly dependent on the heuristic used to determine the elimination order. Especially tailored heuristics where variables are eliminated starting from the last time slice and eliminate all variables in a time slice before progressing to the previous one improve the efficiency achieved with common heuristics. Exact inference is usually only applicable with categorical or finite-state variables, but there are very useful cases where exact inference can be used with numeric variables.

#### 4.3.2.1. Exact inference with numerical variables

OpenMarkov's exact inference algorithms are based on operations with table potentials. Therefore, exact inference is possible only if, given the available evidence, every potential can be transformed into a table. All potentials defined on finite-states variables can be transformed easily into tables. Inference with numeric variables is possible only when each of them has a finite set of possible values, because in this case every potential can be transformed into a table. A numeric variable  $Y$  has a finite set of values when:

1. It has an associated finding.
2. It has no parents and its potential is a delta<sup>3</sup>, because this potential induces a finding for the variable.
3. The value of  $Y$  depends deterministically on its parents,  $y = f(\mathbf{x})$ —where  $\mathbf{X} = X_1, \dots, X_n$  is the set of parents of  $Y$ —and each parent has a finite number of possible values.

In the third case,  $Y$  can be replaced with a finite-states variable  $Y'$  such that its states represent the possible values of  $Y$ . The conditional probability for  $Y'$  is:

$$P(y | \mathbf{z}) = \begin{cases} 1 & \text{if } y = f(\mathbf{z}) \\ 0 & \text{otherwise.} \end{cases} \quad (4.1)$$

The algorithm for converting the numerical variables of a network is as follows:

1. Assign all the possible findings, including those induced by the potentials.
2. Get a list of the variables sorted in topological order, i.e., if  $X$  is a parent of  $Y$ , then  $X$  appears before  $Y$  in the list.
3. For each variable  $Y$  in the list,

---

<sup>3</sup>Also known as Dirac's delta function, where the whole probability density concentrates in a single point.

- a) if  $Y$  is numeric, convert it into finite-states and convert its potential into a table, as explained below;
- b) if  $Y$  is not numeric but some of its parents is numeric, convert the potential into a table.

The conversion of a numeric variable  $Y$  is as follows:

1. Create a new finite-state variable  $Y'$ .
2. Create an empty list (the list of values of  $Y'$ ).
3. For each configuration  $\mathbf{x}$ , add  $f(\mathbf{x})$  to the list (if it is was not already therein).
4. For each  $y'$  in the list, create a new state of  $Y'$  having  $y'$  (a string) as its name and  $y'$  as its numeric value.
5. Create a table potential  $P(y' | x)$ .
6. For each configuration  $x$  and each value  $y'$ , if  $y' = f(\mathbf{x})$ , then  $P(y' | x) = 1$ , else  $P(y' | x) = 0$ .

In the above procedure, if any of the  $X_i$  's is a finite-state variable obtained by converting a numeric variable, the computation of  $f(\mathbf{x})$  must be based on the numeric value of  $X_i$ .

If a variable  $Z$  is finite-states but some of its parents in the original network were numeric, then it is necessary to convert the potential  $P(z | \mathbf{y})$  into a table  $P(z | \mathbf{y}')$ , where  $\mathbf{Y}$  and  $\mathbf{Y}'$  are the sets of parents of  $Z$  in the original network and in the converted network. Let  $\mathbf{Y}_{num}$  be the subset of numeric variables in  $\mathbf{Y}$  and  $\mathbf{Y}'_{num}$  the corresponding variables in  $\mathbf{Y}'$ . If  $\mathbf{Y}_{fs}$  is the set of finite states variables in  $\mathbf{Y}$ , then  $\mathbf{Y}_{fs} = \mathbf{Y} \setminus \mathbf{Y}_{num} = \mathbf{Y}' \setminus \mathbf{Y}'_{num}$ .

The conversion of the potential is done as follows:

1. For each configuration  $y'_{num}$  (a configuration of finite-states variables), we take  $y_{num}$  as a set of numerical findings and project  $P(z | y')$  onto a table which will depend on  $Z$  and  $\mathbf{Y}_{fs}$ .
2. We combine all these tables into a bigger table that will depend on  $\mathbf{Y}'$ .

### 4.3.3. Approximate inference

In approximate inference algorithms, joint probabilities and expected utilities are calculated by running multiple iterations where the value of each variable is sampled following a topological order (all parents are sampled before their child) using its conditional probability distribution: the value of a chance node with parents is sampled from the probability distribution defined by the configuration of sampled states of its parents. The value of a utility node for each iteration is given by the sampled states of its parents and its utility function. The induced probability for a chance variable is estimated by calculating the proportion of iterations each value



has been sampled. For example, if we have run 1,000 simulations and for variable  $X$  the value  $x_1$  has been drawn 238 times, the estimated value for  $P(X = x_1)$  is 0.238. The expected value of a utility node can be estimated as the average of the utilities sampled for that node.

It is analogous to individual level microsimulation in state transition models, where virtual patients are modeled by assigning a state or a numerical value to each variable in the model. In contrast with exact inference, approximate inference can be used with numeric variables that do not fulfill the conditions described above. The drawback of approximate inference is that it only gives approximate results; the higher the precision we desire, the more iterations we need.

#### 4.3.4. Analysis of the evaluation

Using these algorithms, OpenMarkov can compute and display graphically the temporal evolution of a variable, namely, the probabilities if it is a chance variable and the utility (either instantaneous or cumulative) if it is a utility node, as shown in Figure 4.3 and Figure 4.4.




Figure 4.3.: Temporal evolution of the variable *State* for combination therapy

The total economic cost of a strategy is the sum of the cumulative utility of the nodes that represent economic cost; the effectiveness is obtained in the same way. After computing the cost and effectiveness of each strategy, it is possible to perform a cost-effectiveness analysis to rule out the dominated strategies and to find the incremental cost-effectiveness ratios (ICERs).




Figure 4.4.: Aggregate value of the variable *Life years* over time.


## 4.4. Overcoming the Markov property

A common criticism of Markov models is their inability to model patient history in a natural way due to the Markov property (also called Markov assumption). In real life, transition probabilities depend strongly on the time elapsed since a certain event. For example, the probability of death increases with time for patients suffering from cancer or AIDS. Sometimes it is the number of certain events that affects transition probabilities: for example, the probability of death is higher after recurrence of cancer or for subsequent myocardial infarctions. This section explains how MPADs can represent patient history without the modeling tricks typically used in Markov models.

### 4.4.1. Representing time-dependency

In some medical problems the transition probabilities depend on how long the patient has been in a certain state. For example, the longer an HIV patient has had a low CD4 count (state B), the higher the probability of developing AIDS (state C) or dying (state D). In order to represent the increasing transition probability, Chancellor's model can be refined by dividing state B into three sub-states: B1, B2 and B3, such that B1 means that the individual just entered state B, B2 that he entered state B in the previous cycle and B3 that he has been in state B for at least two cycles. States B1 and B2 are called tunnel states because they can be visited only in a fixed sequence, analogous to passing through a tunnel Sonnenberg and Beck (1993) : even if an individual is still in state B (low CD4 count without

AIDS), he will necessarily transition from sub-state B1 to B2 or from B2 to B3, as shown in Figure 4.5.



**Figure 4.5.:** A refined transition diagram for Chancellor’s model. Sub-states B1, B2 and B3 indicate how long the patient has been in state B. B1 and B2 are called tunnel states because the patient cannot be in them for more than one cycle.

An alternative to using tunnel states is to represent the state with two variables: one for the clinical state—having four values or states: A, B, C and D, as in the original model—and the other for indicating how long the patient has spent in the current state. Implementing such a model in a spreadsheet would be complicated (even after discretizing the numeric variable representing time into a finite set of intervals), because spreadsheets cannot represent multidimensional matrices in a natural way. For this reason, when Hawkins et al. (2005) built a Markov model for the treatment of epilepsy and wanted to reflect the decrease in the probability of treatment failure over time, they used R instead. In contrast, it is easy to build an MPAD having one variable for the clinical state and another one for the time spent in the current state. The auxiliary variable *State change [1]* tracks whether there has been a transition in cycle 1. The link from *Time in current state [0]* to *State [1]* implies that transition probabilities do not only depend on the therapy type, but also on the time spent in current state. In this case, only the transitions from state B to state C or D will depend on the time spent in the current state.

A variable, such as *Time in current state*, that contains information about the history of the individual is called a *memory variable*. Memory variables allow the model to satisfy the Markov property, as the transition probabilities only depend on the variables that describe the current state, while rebutting the widespread notion that Markov models do not have memory.

Another way to encode time dependency is to use parametric survival models, which can be fitted to data using statistical techniques. These models learn a probability

density function  $f(t)$  from data, whose associated cumulative density function is  $F(t) = P(T \leq t)$ . We define the survivor function  $S(t)$  as the probability of surviving for a period of time greater than  $t$ , and therefore as the complement of the cumulative density function:

$$S(t) = P(T > t) = 1 - P(T \leq t) .$$

Transition probability between time-points  $t - u$  and  $t$  is calculated using the following formula:

$$tp(t_u) = 1 - S(t)/S(t - u) .$$

The most used probability distributions are Weibull Weibull (1951) and Gompertz Gompertz (1825). According to Gray et al. (2011), the Gompertz distribution is often used to model fatal events, whereas the Weibull distribution is used to model the distribution of non-fatal events.


#### 4.4.2. Representing past events

In some cases transition probabilities depend on the number of times an event has occurred in the history of an individual. For example, a more up-to-date model for HIV should allow backward transitions from state C (AIDS) to A or B, due to the efficacy of modern treatments; transition probabilities could depend on how many relapses the individual has had (defined here as transitions to state C). As in the previous example, this problem might be modeled by multiplying the number of states, such that they do not only represent the current health condition, but also the number of relapses, as shown in 4.6.

An alternative approach would be to have one variable representing the clinical state and another one representing the number of relapses. Again, it would be difficult to implement that model in a spreadsheet, even after discretizing that variable into a finite number of intervals, but the implementation as an MPAD is straightforward.

### 4.5. Probabilistic sensitivity analysis with MPADs

As explained in Section 2.3.1, the objective of sensitivity analysis is to analyze how the uncertainty in the input parameters and translates into uncertainty over the outputs of interest and therefore into decision analysis. Probabilistic sensitivity



**Figure 4.6.:** An alternative refinement of Chancellor’s model that includes backward transitions. The subindex of a state indicates the number of relapses, i.e., the number of times the patient has transitioned into state C; subindex 2 means that the patient has had AIDS at least twice.

analysis (PSA) is a special type of sensitivity analysis where some point estimates are replaced by probability distributions.

Probabilistic sensitivity analysis might be encoded in MPADs by introducing probability distributions in the potentials, instead of numerical values. Transition probabilities as well as utility functions may be assigned different probability distributions including Beta, Gamma, Dirichlet, Normal, and Triangular distributions, as shown in 4.7. In order to account for parameter correlation, when covariance relationships of parameters are known, such as in regression frameworks, the Cholesky decomposition is used. If known, the Cholesky decomposition can be specified or else calculated from the covariance matrix. The Cholesky decomposition provides correlated samples from a multivariate normal distribution.

Each simulation in PSA results in an estimated cost and effectiveness for each of the interventions, which differ from the results of other simulations due to the random sampling from the probability distributions. Therefore, instead of representing each intervention with a single point in the cost-effectiveness plane as in Figure 2.5, each intervention is represented by a cloud of points, as shown in Figure 4.8, one point for each simulation and intervention.

Expected value of information (EVI) analysis, which we explained in Section 2.3.1, can also be performed using MPADs in `OpenMarkov`. In order to perform this analysis, we need to specify the size of the population the intervention affects per

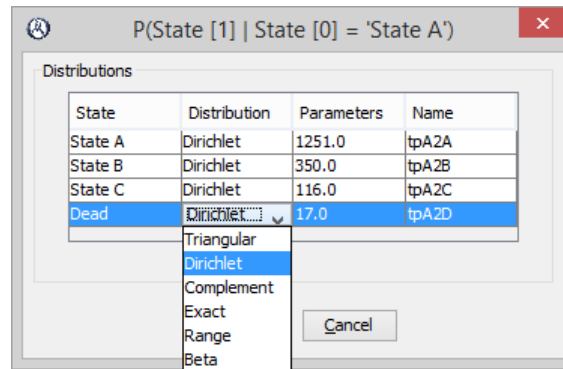


Figure 4.7.: Assigning probability distributions to transition probabilities


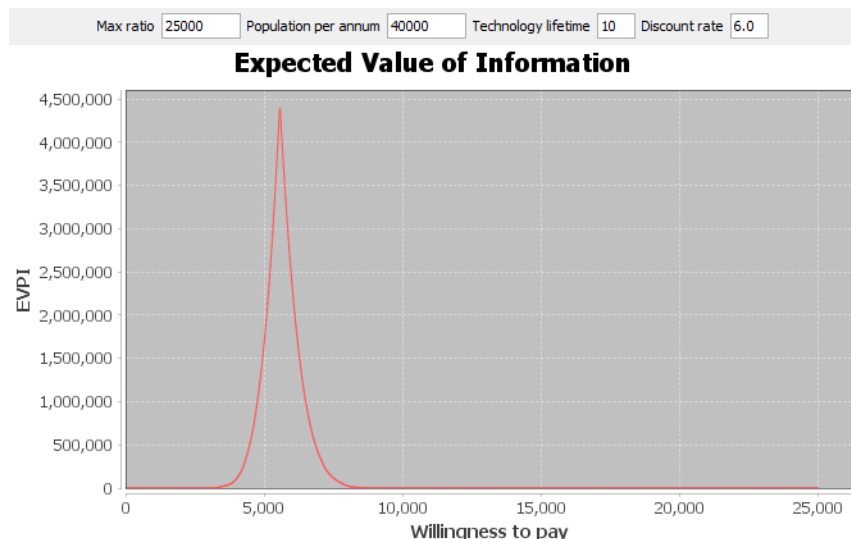


Figure 4.8.: Cost-effectiveness plane for the HPV vaccine model, evaluated using PSA.

year, the number of years the technology we are evaluating is going to be relevant, and the discount rate to be used in the analysis. As can be seen in Figure 4.9, expected value of perfect information peaks around the ICER.

## 4.6. Comparison of MPADs with other tools

In this section we compare MPADs with the main techniques for building state-transition models and performing cost-effectiveness analyses. The current guidelines for NICE technology appraisal includes only Microsoft Excel, TreeAge Pro, R or WinBUGS as standard software packages NICE (2012).



**Figure 4.9.:** Expected value of perfect information for the hip replacement model Briggs et al. (2004).

#### 4.6.1. MPADs vs. spreadsheets

Most state-transition models are implemented in spreadsheet applications (commonly Microsoft Excel). The main advantages of such applications are their general availability, sparing the need to buy dedicated software, and the low level of expertise needed to write (simple) formulas. But when the logic behind a transition probability or a utility function becomes complex, as it is usually the case in real-world applications, the length and obscurity of formulas grow exponentially. The use of indirect references to cells in the spreadsheet rather than the use of meaningful variable names makes complex formulas much harder to understand than the equivalent source code in a high-level programming language, such as Java or C++. Besides, implementing sensitivity analysis usually requires writing macros in a high-level programming language (such as Visual Basic for Microsoft Excel), therefore requiring programming skills.

For a modeler, the main drawback is the need to encode the structure, the parameters and the code for the evaluation of the model in the form of tables and formulas (or macros), which cannot be reused from previous models. Encoding a model and its evaluation in a spreadsheet, in addition to demanding a considerable effort, is a process prone to errors, which are especially hard to detect. In fact, we detected bugs in two of the four spreadsheet models whose results we tried to replicate with MPADs, even though their authors were experienced decision analysts. Fortunately the impact of those mistakes was irrelevant, but there are notorious examples of bugs in spreadsheet models that led to wrong conclusions Herndon et al. (2013); Baggerly and Coombes (2009). In contrast, *OpenMarkov*'s source code is periodically tested using modern software-engineering techniques and tools, and errors while using a

graphical user interface are more unlikely and easier to spot than when writing formulas.

For an evaluation agency, the main advantage of a model implemented as a spreadsheet is that the same file contains the model, the code and all the numerical results of its evaluation, thus making tracing its execution step by step possible. On the other hand, spreadsheet models usually include complex formulas and several macros. Without a close collaboration with the modeler, it is usually very complicated to decipher the logic implicit in obscure formulas which might even encode structural assumptions. This makes model validation time-consuming and prone to error. In contrast, MPADs are more transparent because their causal graphs clearly show the structural assumptions and because the model can be inspected without reading a single line of code.

### 4.6.2. MPADs vs. Markov cycle trees

Markov cycle trees are a particular type of decision trees in the same way as MPADs are a particular type of PGMs. They both are based on graphs containing three types of nodes: chance, decision and utility. The main difference is that in decision trees the links outgoing from a node represent the values (states) of the corresponding variable, while links in MPADs represent causal relations. In general, every variable appears several times in the tree, and the number of branches grows exponentially with the number of variables. For this reason many decision analysts find Markov cycle trees inadequate for building large models. In contrast, in an MPAD each variable can appear at most twice, which leads to much more compact models and hence to the possibility of solving much more complex problems.

Additionally, when the characteristics of the model require to take patient's history into account, Markov cycle trees have two possibilities: to multiply the number of states, as described in 4.4, or to use tracker variables. The first option leads to cumbersome models, just as when using spreadsheets. The second requires individual-level simulation to evaluate the model, which entails an extra computing cost, as described in 2.3.2.2, which make Markov cycle trees with tracker variables unaffordable for large problems. In contrast, many MPADs having memory variables can be efficiently evaluated with cohort analysis, as explained in 4.4.

With respect to the software tools, the main package for building Markov cycle trees and performing cost-effectiveness analysis is **TreeAge Pro**, a mature commercial product. Older programs, such as **WinDM** and **SMLTREE**, are no longer available. **OpenMarkov**, in contrast, is still a research prototype, but in addition to supporting MPADs and other types of PGMs, has the main advantage of being open-source, which implies that it can be used for free, and also that its code can be inspected and extended to address the user's needs. **TreeAge Pro** offers many features that are not available on **OpenMarkov**, but in turn the latter has many more features not available in the former.



### 4.6.3. MPADs vs. programming languages

Some state-transition models for cost-effectiveness analysis have been implemented in C++, Matlab, R, Stata or a language in the BUGS family, such as WinBUGS, OpenBUGS or JAGS. The main advantage of this approach is the possibility of building virtually any model and any evaluation algorithm. The drawback is that, as in the case of spreadsheets, everything must be programmed by the modeler, from the combination of transition probabilities to the sensitivity analysis and the graphical display of results. Statistical languages, such as R or BUGS, offer many validated statistical functions which might not be available as libraries for languages like C++ and they are higher level programming languages with a milder learning curve. Nevertheless, writing and debugging the code in such statistical languages requires an expert modeler and a lot of time.

Building and evaluating an MPAD with OpenMarkov is much faster and easier, but some models may require some features that are not yet available in this tool. In some cases, extending OpenMarkov with such features will still be a smaller task than implementing the model in R or BUGS from scratch. However, as we are continuously adding features, we believe this scenario will grow ever more unlikely in the future.

From the point of view of an evaluation agency, assessing the validity of a model encoded in a programming language may be difficult and time consuming, especially if the code lacks a good design, clear in-line comments and sufficient additional documentation.

Finally, we should mention that BUGS is especially suited for health technology assessment based on clinical trials because of its ability to integrate statistical inference and cost-effectiveness analysis. This is something that currently cannot be done with MPADs. An interesting line for future development is the integration of JAGS (the most portable version of BUGS) and OpenMarkov, thus combining the statistical power of the former for evidence synthesis with the facilities for model building and analysis of the latter.

### 4.6.4. MPADs vs. discrete event simulation

The advantages and disadvantages of discrete event simulation (DES) over Markov models have been extensively and passionately discussed in the literature Karnon (2003); Caro (2005); Caro et al. (2007); Simpson et al. (2009). Most of the arguments on this discussion are also valid when comparing DES with MPADs.

However, MPADs include some of the advantages of DES over Markov models, such as the ability to use several variables to represent the evolution of the patient—including memory variables, as explained in Section 4.4— without the need to resource to state multiplication; additionally, both DES and MPADs can handle numeric variables with similar ease. MPADs also include some of the advantages of Markov models over DES, such as the easiness to build, debug, and communicate

the models. Another advantage of MPADs over both DES and Markov models is the use of causal graphs, which facilitate the construction of the model, as the decision analyst can specify and modify its structure before addressing the quantitative details. The causal graph also offer the evaluator an overall vision of the model, with the possibility of focusing on a particular node to further analyze the relations between the variables and the numerical parameters.

From the point of view of the computational cost, many MPADs—including all those mentioned in this chapter—can be evaluated with cohort analysis, which is in general more efficient than DES. The extra computational cost of DES might hinder in some casethe conduction of probabilistic sensitivity analysis Griffin et al. (2006).

There is a wide range of software packages available for DES, as it is used in a variety of applications including engineering and finance. **Arena®**, a proprietary software, seems to be the most widely used tool to build DES models for CEA Caro (2005); Simpson et al. (2009). **OpenMarkov** has the advantage of being open source, as explained when comparing it to **TreeAge Pro**, and that offers functionality especially tailored for CEA, which, to the best of our knowledge, no DES software does.

DES is especially useful to model many entities that interact between them and with a set of constrained resources; for example, several beds in an emergency service to be occupied by patients who arrive over time. However, if the model does not include such special requirements, the disadvantages of DES—such as the complexity, the high number of required input parameters, and the higher computational cost—outweigh its advantagesKarnon (2003); Griffin et al. (2006).

## 4.7. Conclusions

Markov processes with atemporal decisions are a new type of probabilistic graphical model especially designed for cost-effectiveness analysis within the state-transition modeling framework.

We have argued that most state-transition models implemented as spreadsheets or decision trees can be built much more easily using MPADs. In particular, MPADs can use memory variables to summarize the patient's history without an explosion in the number of states. They can also combine numeric and categorical variables in the same model. These features make it possible to implement as MPADs complex models that exceed by far the expressive power of spreadsheets.

For a modeler, building an MPAD and evaluating it with **OpenMarkov** is much easier and faster than implementing it as a spreadsheet or a with a programming language. In particular, its graphical user interface allows the modeler to save considerable time and reduces drastically the probability of mistakes. In the case of complex models, it is also easier and faster to build an MPAD than a Markov cycle tree.

For an evaluation agency, understanding and assessing the correctness of MPADs is much easier than in the case of other representation techniques for two reasons:

because MPADs are based on causal graphs that summarize the structure of the model and because each model can be inspected separately from the source code used to evaluate it.

With respect to computational efficiency, many MPADs—including all the examples mentioned in this chapter that combine categorical and numeric variables—can be evaluated with cohort analysis, which is more efficient than stochastic individual-level simulation; this is an advantage over Markov cycle trees using tracker variables and DES models. However, there are other types of problems for which DES or dynamic transmission models are more adequate.

The main drawback of MPADs is that their applicability in practice is limited by the features offered by **OpenMarkov**, the only tool currently available for building and evaluating MPADs. However, this tool already offers advanced representational features—such as the Weibull distribution and generalized linear models with fractional polynomials, and the corresponding correlation matrices—as well as several types of deterministic and probabilistic analysis. This tool is under active development at the CISIAD (UNED) and new features are added when needed to address new medical problems.



# 5. Tuning networks

We have defined a type of network that models the problem of finding the set of adjustments needed for the optimal performance of a system with tunable parameters.

## 5.1. Basic properties

### 5.1.1. Elements of the network

Tuning networks consist of a DAG and a probability distribution and can contain decision and utility nodes. Decision nodes represent tunable parameters in tuning networks. The values of these nodes can represent either setting the value of the parameter to a certain value in absolute terms or to do it in relative terms, that is, they represent the change in the value of the parameter (increase, decrease or no change). In both cases, we define the term of the *neutral state*, for relative variables as the "no change" value and for the absolute variables as the ideal value of the parameter. Tuning networks also allow the definition of restrictions as described in Section 3.1.1.

### 5.1.2. Interpretation of the network

A PGM representing the system described above has at least two possible interpretations:

The *diagnostic interpretation* assumes that the optimal tuning is unique, i.e., there is only one configuration of the parameters (the inputs) that makes the system perform optimally. The three values of each variable are interpreted as  $\{decreased, optimal, increased\}$ . The current output of the system is introduced as evidence into the Bayesian network and the goal of inference is to *diagnose* which parameters are not properly tuned. Therefore, inference proceeds down-up, i.e., from the observed outputs to the inputs. This method has two advantages: first, it does not require a global gain function, and second, inference is more efficient, because it evaluates the network only once, while the variational interpretation need to evaluate the network once for each change in one of the parameters.

The *variational interpretation* tries to *predict* the impact that a change in the value of the parents (the causes or the inputs) will have on the children (the effects or

the outputs). The three values of each variable,  $-$ ,  $0$ ,  $+$  are interpreted as  $\{\text{decrease, no-change, increase}\}$ . Initially all the variables take on the value *no-change* by definition. The process of inference consists of computing the posterior probability of each output variable for each change that the user may impose on the input variables. Therefore, in this interpretation inference proceeds top-down, i.e., from (the possible adjustments of) the inputs to the outputs, using predictive reasoning. The changes that lead to an improved performance will be applied. When an improvement in some of the output variables comes together with a worsening in others, it is necessary to have a utility function that measure the global gain in performance.

## 5.2. The tuning model

The *tuning model* is an ICI model (see Section 2.1.4.1) that represents how change in some variables (the parents) affects another variable (the child). The motivation for the proposal of this new model is that, when building a PGM for programming cochlear implants, we needed a model that could represent how changes in some of the parameters of the physical device affect other properties of the system, which in turn may affect the subject's hearing performance. Given that none of the existing models fitted our needs, we devised the tuning model.

### 5.2.1. Mathematical definition of the tuning model

A noisy ICI model is defined by three elements: the domains of the variables, the function  $f$ , and some constraints on the values of  $P(z_i | x_i)$ .

In the tuning model, all the variables have the same **domain**,  $-$ ,  $0$ ,  $+$ , where “ $-$ ” represents a decrease in the value of the variable, “ $+$ ” represents an increase, and “ $0$ ” means “no change”. If the variable is denoted by  $V$ , we will sometimes write  $v^-/v^0/v^+$  instead of  $-/0/+$  to make it clear what variable we are speaking of.

The **function** of the tuning model is defined as follows:

$$f_{\text{tuning}}(\mathbf{z}) = \begin{cases} y^+ & \text{if } n^+(\mathbf{z}) > 0 \\ y^0 & \text{if } n^+(\mathbf{z}) = 0 \\ y^- & \text{if } n^+(\mathbf{z}) < 0, \end{cases} \quad (5.1)$$

where  $n^+(\mathbf{z})$  is a function that returns the number of variables that take the value  $+$  in configuration  $\mathbf{z}$  minus the number of those that take the value  $-$ . For example,  $n^+(z_1^+, z_2^+, z_3^+) = 3$ ,  $n^+(z_1^+, z_2^-, z_3^0) = 0$ , and  $n^+(z_1^+, z_2^-, z_3^-) = -1$ . Therefore,  $f_{\text{tuning}}(z_1^+, z_2^+, z_3^+) = y^+$ ,  $f_{\text{tuning}}(z_1^+, z_2^-, z_3^0) = y^0$ , and  $f_{\text{tuning}}(z_1^+, z_2^-, z_3^-) = y^-$ .

A **constraint** that we impose on  $P(z_i | x_i)$  is that

$$P(z_i^0 | x_i^0) = 1, \quad (5.2)$$

which implies that  $P(z_i^+ | x_i^0) = P(z_i^- | x_i^0) = 0$ . Therefore, if we introduce four parameters for each link,  $c_i^{++}$ ,  $c_i^{+-}$ ,  $c_i^{-+}$ , and  $c_i^{--}$ , the CPT for that link has the form shown in Table 5.1.

$P(z_i   x_i)$	$x_i^-$	$x_i^0$	$x_i^+$
$z_i^+$	$c_i^{-+}$	0	$c_i^{++}$
$z_i^0$	$1 - c_i^{-+} - c_i^{--}$	1	$1 - c_i^{++} - c_i^{+-}$
$z_i^-$	$c_i^{--}$	0	$c_i^{+-}$

**Table 5.1.:** Conditional probability table for link  $X_i \rightarrow Y$  in the tuning model.

It is possible to prove from Equation 2.16 that when all the  $X$ s take the value 0, then  $Y$  takes the value 0 with absolute certainty. When  $X_i$  takes the value + and the other  $X$ s take the value 0, then  $Y$  takes the value + with probability  $c^{++}$  and the value – with probability  $c^{+-}$ . Similarly, when  $X_i$  takes the value – and the other  $X$ s take the value 0, then  $Y$  takes the value + with probability  $c^{-+}$  and the value – with probability  $c^{--}$ . Therefore, the four  $c$ -parameters quantify the individual impact of  $X_i$  on  $Y$ .

### 5.2.2. Classes of interactions

We have seen that the general form of the conditional probability table associated with link  $X_i \rightarrow Y$  is as shown in Table 5.1. However, it is possible to impose a second **constraint** for each variable  $X_i$  and each value of this variable,  $P(z_i^+ | x_i) = 0$  or  $P(z_i^- | x_i) = 0$ ; put another way:

$$(c_i^{-+}=0 \vee c_i^{--}=0) \wedge (c_i^{-+}=0 \vee c_i^{--}=0). \quad (5.3)$$

Therefore, when this constraint holds for a link  $X_i \rightarrow Y$ , only two parameters are different from 0—in contrast with the general case, which requires four independent parameters—and that link must belong to one of four classes: direct, inverse, always increasing, and always decreasing.

The *direct* class is shown in Table 5.2. The values in the  $x_i^0$  column are imposed by the first constraint of the tuning model (Eq. 5.2). The  $x_i^-$  column implies that a decrease in  $X_i$  causes a decrease in  $Y$  with a probability  $c_i^{--}$ , such that  $c_i^{--} > 0$ . It may occur, with probability  $1 - c_i^{--}$ , that a decrease in  $X_i$  fails to cause a change in  $Y$ ,

but that decrease can never cause an increase in  $Y$ . Similarly, the  $x_i^+$  column means that an increase in  $X_i$  causes an increase in  $Y$  with a probability  $c_i^{++}$ . Therefore, this class represents a positive influence of  $X_i$  on  $Y$  Lacave et al. (2006) and a positive correlation between both variables.

$P(z_i   x_i)$	$x_i^-$	$x_i^0$	$x_i^+$
$z_i^+$	0	0	$c_i^{++}$
$z_i^0$	$1 - c_i^{--}$	1	$1 - c_i^{++}$
$z_i^-$	$c_i^{--}$	0	0

**Table 5.2.:** Conditional probability table for a link  $X_i \rightarrow Y$  of the *direct* class.

Similarly, the *inverse* class is characterized by  $c^{++} = c^{-+} = 0$ ,  $c^{-+} > 0$ , and  $c^{+-} > 0$ , which implies that a decrease in  $X_i$  causes an increase in  $Y$ , and vice versa, thus leading to a negative correlation between both variables.

The relations that define the *always decreasing* class are:  $c^{-+} = c^{++} = 0$ ,  $c^{--} > 0$ , and  $c^{+-} > 0$ . Therefore, any change in  $X_i$  will cause a decrease in  $Y$ . The properties of the *always increasing* class are analogous.

We may impose a third **constraint**: the *symmetry* of the influence. In the case of an *direct* interaction, it implies that  $c_i^{++} = c_i^{--}$ , i.e., the probability that an increase in  $X_i$  causes an increase in  $Y$  is the same as the probability that a decrease in  $X_i$  causes a decrease in  $Y$ . In the case of an *always decreasing* interaction, the probability of a decrease in  $Y$  is the same for an increase in  $X_i$  as for a decrease:  $c_i^{--} = c_i^{++}$ . A link satisfying the condition of symmetry requires only one parameter.

Several kinds of interaction may coexist within the same family.

**Example.** In a family with four parents, the link  $X_1 \rightarrow Y$  might be general (i.e., free from the second and third constraints, as shown in Table 5.1),  $X_2 \rightarrow Y$  might be a direct interaction,  $X_3 \rightarrow Y$  might be direct and symmetric, and  $X_4$  might be always decreasing. The total number of parameters for this model would be  $4 + 2 + 1 + 2 = 9$ .

If the interaction of this family did not use any canonical model, its conditional probability table would require  $3^5 = 243$  parameters, but given that there are  $3^4 = 81$  constraints among them, this family would require  $243 - 81 = 162$  independent parameters. Obtaining those parameters from a database is unreliable, unless in the case of a huge database, because many of the configurations of the  $X$ s will not be represented. Obtaining those parameters from an expert would be impossible in practice not only for the amount of time required, but mainly because estimating the probability of  $Y$  for each configuration of the  $X$ s exceeds by far the cognitive capabilities of the human mind.



### 5.2.3. Causal interpretation of the tuning model

In the tuning model  $Y$  represents a parameter of a system whose value depends on the values taken on by other parameters,  $\{X_1, \dots, X_n\}$ . Each auxiliary variable  $Z_i$ , associated with link  $X_i \rightarrow Y$ , as shown in Figure 2.4, indicates whether a change in  $X_i$  (from  $x_i^0$  to  $x_i^+$  or  $x_i^-$ ) has induced a change in  $Y$ :  $z_i^+$  indicates an increase (from  $y^0$  to  $y^+$ ) while  $z_i^-$  indicates a decrease (from  $y^0$  to  $y^-$ ).

The first constraint, given by Equation 5.2, means that when there is no change in  $X_i$ , then there is no change in  $Y$ .

The second constraint, given by Equation 5.3, means that a change in  $X_i$  may cause either an increase or a decrease in  $Y$ , but not both; this assumption seems reasonable for some domains, but there might be others in which an increase (or a decrease) in  $X_i$  sometimes produces an increase in  $Y$  and sometimes a decrease.

The function  $f_{\text{tuning}}(\mathbf{z})$ , given by Equation 5.1, means that when some of the  $X_i$ s induce an increase in  $Y$  and others cause a decrease, the global effect depends on whether there are more increases than decreases, or vice versa, or there is a tie.

The tuning model is used to build probabilistic networks in which each variable represents a property of the system. In these networks, a node without parents represents a physical parameter that can be adjusted by the user, while a node  $Y$  with parents  $\{X_1, \dots, X_n\}$  represents a parameter or a property of the system whose value depends on other parameters (its parents).

### 5.2.4. Inference with the tuning model

Inference on the tuning model can be done either with exact or stochastic inference in quite different ways:

**Exact Inference** Using exact inference, conditional probabilities of the tuning model are obtained by marginalizing out the  $Z_i$ s as in Equation 2.12. The computational cost of exact inference is  $O(\exp(n))$  where  $n$  is the number of parents.

**Stochastic inference** In stochastic inference a sample is picked randomly for each variable. In order to pick a sample from the tuning model, one could first calculate the conditional probability as in exact inference and pick a sample from that probability distribution, but it is much more efficient to pick a sample for all the  $Z_i$ s and then calculating  $y = f(\mathbf{z})$ ,  $f$  function in Equation 5.1. The computational cost of stochastic inference is  $O(n)$  where  $n$  is the number of samples.

### 5.3. Construction of a tuning network

The construction of a tuning network for the tuning of a physical system begins by selecting the variables. Some of them will represent variations in the parameters of the system. The domain of each of these variables will be  $\{-, 0, +\}$ , which implies a discretization of a continuous variable. The value 0 might indicate that the value of the parameter has not changed at all, and  $-/+$  might represent any increase/decrease in its value, no matter how small it might be. However, in practice it is better that 0 indicates “no significant change”,  $-$  represents a significant decrease and  $+$  represents a significant increase. It is the knowledge engineer, in collaboration with human experts, who must determine what constitutes a significant change. For instance, the threshold might be  $\pm 5\%$  of the absolute value of the parameter represented by the variable; for a different variable tuned with higher or lower precision the threshold might be  $\pm 2\%$  or  $\pm 10\%$ , respectively. This threshold might be different for each variable in the network, but it must be very clearly defined, because it will affect the elicitation of the conditional probabilities.

The second step in the construction of a tuning network is to draw causal links between the variables, which is usually the easiest task in the construction of the network.

The third step is to analyze for each family in the tuning network the possibility of applying a canonical model. The conditions for applying an OR, a MAX, a MIN, or an XOR model are discussed in Díez and Druzdzel (2006), while the conditions for applying a tuning model have been described in the previous section.

The fourth step is to obtain the numerical parameters, i.e., the conditional probabilities for each family in the Bayesian network.

#### 5.3.1. Elicitation of a tuning model

The first thing to do when considering to apply a tuning model to a relation in a family is to check whether the requirements are met. The first condition for applying a tuning model is that all the variables involved in the family must have the same domain:  $\{-, 0, +\}$ . The second condition is that the effects of the parents can be combined by applying the function  $f_{\text{tuning}}$  defined in Equation 5.1, which basically states that each change in one of the  $X$ s may produce an increase or a decrease in  $Y$ , and the resulting value of  $Y$  depends on whether there are more increases than decreases, or vice versa, or there is a tie. The third condition is that the increase or decrease produced by each  $X_i$  only depends on the value taken by this variable, not on the values of the other  $X$ s; this condition seems difficult to assess for a human expert because in fact the individual effects are combined by the function  $f_{\text{tuning}}$ , and consequently it is difficult to think of the individual effects “before” being combined. Therefore, it is reasonable to give the third condition for granted and assume that the tuning model can be applied whenever the first two conditions hold.

Once we know we can apply the tuning model, the next thing to do is to obtain the numerical parameters, i.e., the conditional probabilities for the family. In the tuning model, each link  $X_i \rightarrow Y$  must be analyzed independently of the others. The first question is: “Does the second constraint, given by Equation 5.3, hold for this link, or is it possible that the same change in  $X_i$  sometimes causes an increase in  $Y$  and other times a decrease?” In the latter case it will be necessary to obtain four parameters:  $c^{++}$ ,  $c^{+-}$ ,  $c^{-+}$ , and  $c^{--}$ . However, some of the parameters might coincide; for example, using causal knowledge we might state that  $c^{++} = c^{--}$  and  $c^{-+} = c^{+-}$  (assumption of symmetry). This would reduce the number of independent parameters to be estimated.


On the contrary, if the second constraint holds, the next question to be asked is: “What class of interaction is this: direct, inverse, always increasing, or always decreasing?” (see Sec. 5.2.2). The last question is about symmetry. For example, in the case of a *direct* interaction, the question is: “Does a decrease in  $X_i$  cause a decrease in  $Y$  with the same probability that an increase in  $X_i$  causes an increase in  $Y$ ?” If there is symmetry, we only need to elicit one parameter; otherwise, we need two.

Then, we have to estimate the numerical value(s) of the parameter(s) of each link. The question to be asked for each parameter can be derived from the comment in the last paragraph of Section 5.2.1. For example, the question for parameter  $c^{++}$  is: “What is the probability that an increase in  $X_i$  causes an increase in  $Y$  when there is no change in the other parents of  $Y$ ?” The questions for the other parameters are analogous.

Finally, we must consider for that family whether a noisy tuning model suffices or it is necessary to apply a leaky tuning model. The question is: “Is it possible that a change in some of the physical parameters not explicitly represented in the Bayesian network causes a change in  $Y$ ?” If the answer is affirmative, the question: “What is the probability that they cause an increase in  $Y$  (when none of the explicit parents changes)?” will give us an estimate for the leak parameter  $c_L^+$ . The question: “What is the probability that they cause a decrease in increase in  $Y$ ?” will yield  $c_L^-$ .

### 5.3.2. Conditioned interactions

It may occur in practice that the effect of a certain variable—say  $X_1$ —on  $Y$  depends on the value of a third variable,  $C$ . For example,  $X_1$  may represent a change in the value of a physical parameter (an increase or a decrease) while  $C$  represents the absolute value of that parameter. In this situation we can apply a modeling trick consisting of adding an auxiliary variable  $A_1$ , as shown in Figure 5.1. The interaction between  $A_1$  and  $Y$  is given by the identity matrix:  $P(a_1 | y) = \delta_{a_1, y}$ , where  $\delta$  is Kronecker’s delta function; put another way, the link  $A_1 \rightarrow Y$  is a deterministic symmetric *direct* interaction (see Table 5.2) with  $c_i^{--}$ .



**Figure 5.1.:** Conditioned interaction: the effect of  $X_1$  on  $Y$  depends on the value of variable  $C$ . The combined effect of  $X_1$  and  $C$  is modeled by the auxiliary variable  $A_1$ .

Table 5.3 shows with a hypothetical example how the effect of  $X_1$  on  $Y$  may depend on a conditioning variable,  $C$ . When the value of  $C$  is *low*, the interaction is bottom-up and symmetric: in 90% of cases, a decrease in  $X_1$  causes a decrease in  $Y$ , and vice versa. When  $C = \textit{medium}$ , the interaction is also bottom-up and symmetric, but the effect is qualitatively smaller, i.e., it occurs in a lower proportion of cases. When  $C = \textit{high}$ , the effect of a change in  $X_1$  is unpredictable: it may cause an increase in  $Y$  but it may also cause a decrease, and the variability is asymmetric: it is higher for an increase in  $X_1$  than for a decrease.

$P(a_i   x_i)$	$C = \textit{low}$			$C = \textit{medium}$			$C = \textit{high}$		
	$x_i^-$	$x_i^0$	$x_i^+$	$x_i^-$	$x_i^0$	$x_i^+$	$x_i^-$	$x_i^0$	$x_i^+$
$a_i^+$	0	0	0.9	0	0	0.4	0.04	0	0.26
$a_i^0$	0.1	1	0.1	0.4	1	0.6	0.68	1	0.67
$a_i^-$	0.9	0	0	0.6	0	0	0.28	0	0.07

**Table 5.3.:** Conditional probability table for the auxiliary variable  $A_1$ .

## 5.4. Inference: finding a near-optimal strategy

As in influence diagrams, the goal of inference in tuning networks is to find the optimal strategy. In the case of tuning networks, the optimal strategy is the set of adjustments or changes in the tunable parameters that will maximize the global expected utility of the model. As common inference algorithms for influence diagrams would be computationally unaffordable, we propose a greedy algorithm that explores the space of possible strategies searching for the optimal strategy. The algorithm can be divided in two parts: score and search.

**Search** The algorithm performs a greedy search over the space of possible strategies. The search space is limited by link restrictions, described in Section 5.1.1. The search is initialized setting all policies for all decision nodes to the *neutral state*. The greedy search then looks for the single change in a decision node's policy (also referred to as an adjustment) that maximizes the utility. After changing the policy for that node, it repeats the search for the next change that will maximize the utility.

**Score** The computation of the score is the computation of the global expected utility given a certain strategy. In this case, the computation of probabilities and utilities uses the same algorithms as influence diagrams.

---

**Algorithm 2:** Search for near-optimal strategy

---

**Data:** Set of decision nodes, PGM

**Result:** Set of adjustments

set the imposed policies of all decision nodes to the neutral state;

find best adjustment;

**while** *there is a better adjustment* **do**

    add current best adjustment to result set;

    add current best adjustment to imposed policies;

    find best adjustment;

**end**

---

## 5.5. Discussion

The tuning model arose from a need encountered when building a Bayesian network for real-world problem: the programming of cochlear implants. In this domain of application, it soon became clear that the diagnostic approach was inappropriate, hence we decided to use the predictive approach positive results.

The fact that the tuning model arose from a real-problem is a difference with some of the canonical problems proposed in the literature, which came out from mathematical speculation and have never been implemented on a software tool nor used in practice.



**Part III.**

**APPLICATIONS TO MEDICINE**





## 6. Cost-effectiveness analysis with MPADs

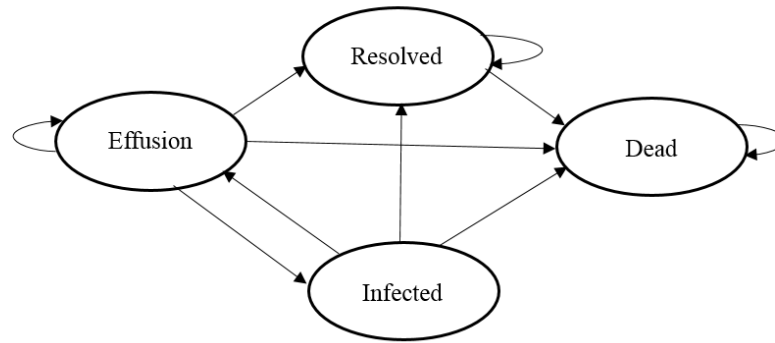
### 6.1. Cost-effectiveness of PleurX® vs. talc

#### 6.1.1. Background

Malignant pleural effusions (MPE) are common among oncologic patients. The median prognosis for patients with MPE is poor (3 to 6 months) and it causes pleuritic pain and dyspnea, leading to a reduced quality of life. Conventional treatment options of MPE include thoracentesis, thoroscopic chemical pleurodesis with talc, and placement of a long-term indwelling catheter. Thoracentesis, the removal of fluid from the pleural space, is effective in relieving acute symptoms but has a 98-100 recurrence rate at 30 days (Neragi-Miandoab, 2006). Pleurodesis, the procedure to artificially obliterate the pleural space, leads to effusion resolution and prevents its recurrence. Previous studies have concluded that chest tube insertion using talc as sclerosing agent is effective in achieving pleurodesis, but it requires hospitalization, it is usually painful, and can be fatal (Kennedy and Sahn, 1994; Dresler et al., 2005). Long-term indwelling pleural catheters, which may be placed in an outpatient setting and are less invasive, have been used to treat MPE for more than a decade with positive outcomes (Putnam et al., 2000), but are slower in achieving pleurodesis and have a higher complication rate (Dresler et al., 2005; Tremblay et al., 2007).

The objective of the current study is to analyze the relative cost-effectiveness of two common treatment options for MPE: chest tube placement with talc and placement of an indwelling pleural catheter, such as PleurX®, for home-based drainage of effusions. Cost-effectiveness analysis is performed using a Markov process with atemporal decisions (MPAD) —see Chapter 4. The evaluation of the MPAD results in the estimation of the quality-adjusted life years (QALYs) and costs derived from both treatments in order to determine the incremental cost-effectiveness of the PleurX® treatment.

This analysis is the result of the collaboration with Dr. María José Roca, a medical specialist in thoracic surgery at the Hospital Universitario Virgen de la Arrixaca in Murcia and associate professor in the Department of Thoracic Surgery at the University of Murcia Medical School. She proposed this research and acted as the expert in the elicitation process.



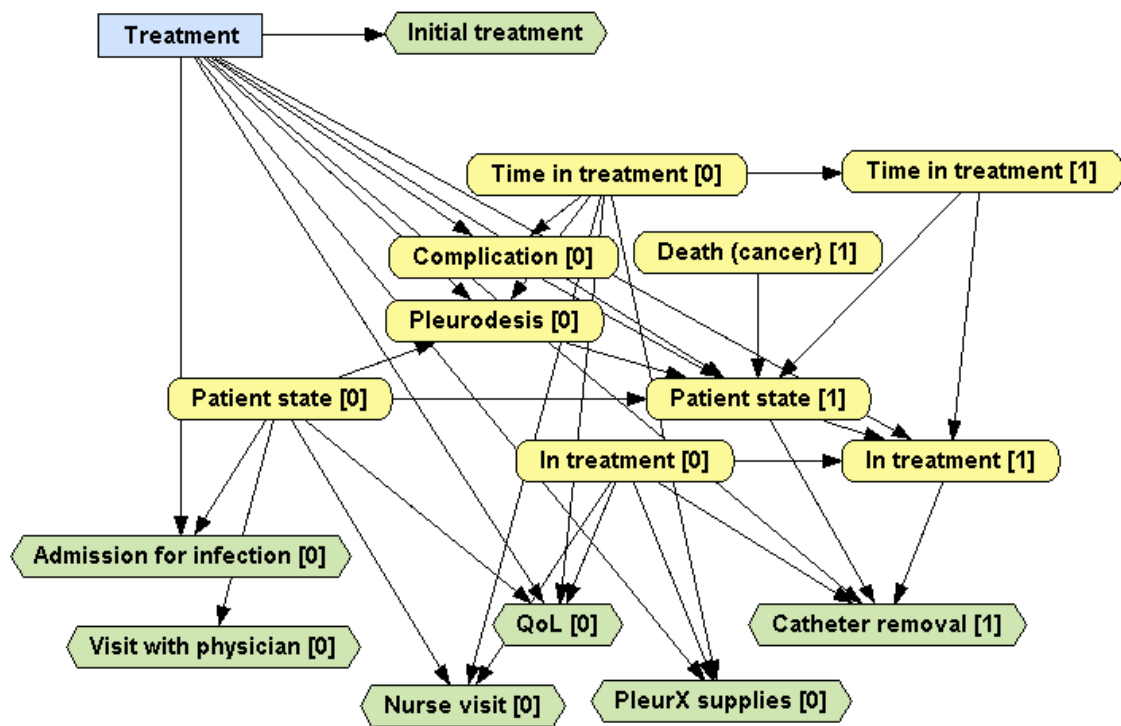
**Figure 6.1.:** State transition diagram for patients with MPE

## 6.1.2. Description of the model

The target population of the analysis are patients with any type of cancer and recurrent MPE. We assume that the diagnosis of MPE has been established at the time of the intervention, that both treatments have no statistically significant effects on mortality, and that all effects are expressed in terms of patient’s utilities achieved through relief of symptoms. The model was built from a health provider perspective. A previous atemporal model built by Olden et al. (Olden and Holloway, 2009) for the cost-effectiveness analysis of talc and PleurX® was used as inspiration for our model. Instead, we built a state transition model encoded as an MPAD.

### 6.1.2.1. States and transitions

We assume that patients are in one of the states shown in the transition diagram illustrated in Figure 6.1. *Effusion* is the initial state, as patients enter the model just after diagnosis. At the end of each cycle, a patient can either stay in the *effusion* state or progress to one of the other states: if the treatment is successful achieves pleurodesis, the patient progresses to *resolved*; if there is a complication related to the treatment, the patient might progress either to *infected* or *dead*. Infection and death were the only complications included in the model—the former because it was the most frequent complication and the latter for its impact. A patient in the *infected* state might progress to *resolved* if pleurodesis is achieved or to *effusion* otherwise. Patients in the *resolved* state will remain therein until death. Finally, patients in any other state might transition to the *dead* state. Due to the short life expectancy after diagnosis, we set the duration of each cycle to a week.



**Figure 6.2.:** Concise representation of the structure of the MPAD, containing the relevant variables of the model.

### 6.1.2.2. Structure of the MPAD

In order to build the MPAD, we first built the causal graph shown in Figure 6.2. The *Treatment* variable represents the decision over which treatment to apply and has two values: *pleurX* and *talc*. The *Patient state* temporal variable represents the state of the patient at each cycle, and has the a value for each state of the transition diagram described in the previous section. *Time in treatment* keeps the count of the time passed since the beginning of the treatment, since it is relevant for example, for the probability of pleurodesis at each cycle. The amount of liquid evacuated through the catheter will also change over time, decreasing as the time passes, and therefore the frequency of drainage will also decrease. The variable *Pleurodesis* determines whether the patient achieved pleurodesis in that cycle, or put it with other words, if the effusion was resolved in that cycle. The temporal variable *Complication* determines whether the patient has a complication in a particular cycle. Similarly, *Death (cancer)* represents whether the patient dies of cancer in a particular cycle. *In treatment* is an auxiliary variable that keeps track of whether the patient is still under treatment. The rest are utility nodes: the temporary *QoL* utility node represents the quality of life of the patient and the rest represent the various costs incurred throughout treatment.

### 6.1.2.3. Quantitative information


Table 6.1 contains the list of numeric parameters (probabilities and utilities) used in the model, along with their source and the range used in one-way sensitivity analysis, for those parameters for which sensitivity analysis was carried out.

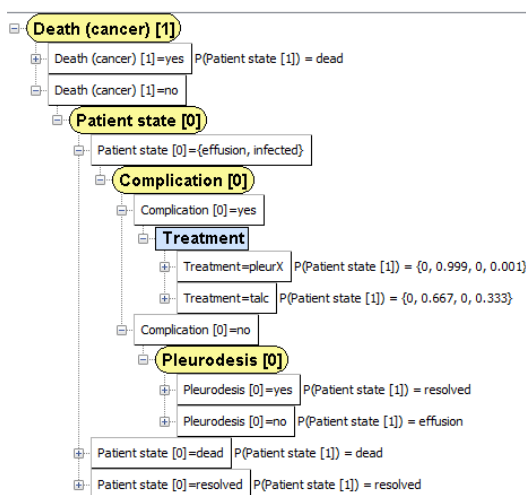
**Probabilities** Given that all the consulted sources reported total rates, we had to calculate weekly transition probabilities from the reported rates and the length of treatments. Whereas in (Olden and Holloway, 2009), all patients were assumed to live exactly as long as their life expectancy, we modeled the mortality rate of patients (the probability potential of *Death (Cancer)*) with a Weibull distribution (Weibull, 1951). With the help of an expert clinician, we estimated the shape parameter of the distribution to be equal to 3 and calculated the scale parameter so that the median of the distribution would match the assumed life expectancy, which for the base case is 6 months (26 weeks). Figure 6.3 shows the weekly mortality and the cumulative survival during the first year after diagnosis for a cohort with a that life expectancy.

The probability for *Pleurodesis* depends on the type of treatment. Patients treated with talc spend a week in the hospital and have a 0.8 probability of pleurodesis or effusion resolution within a week. On the contrary, patients treated with an indwelling pleural catheter do not usually achieve pleurodesis so early. Instead,

**Table 6.1.:** Model parameters

	Variable	Treatment	Estimate	Range	Source
Probabilities	Pleurodesis	Talc	0.8	0.62-0.91	(Olden and Holloway, 2009)
		PleurX®	0.43	0.21-0.70	(Tremblay and Michaud, 2006)
	Pleurodesis after infection	Talc	0.9		(Olden and Holloway, 2009)
		PleurX®	0.65		(Olden and Holloway, 2009)
	Complication	Talc	0.015		(Dresler et al., 2005)
		PleurX®	0.075	0.02-0.13	(Tremblay and Michaud, 2006)
	Infection, given complication	Talc	0.333		(Tremblay and Michaud, 2006)
		PleurX®	0.9999		(Olden and Holloway, 2009)
Death, given complication	Talc	0.667		(Tremblay and Michaud, 2006)	
	PleurX®	0.0001		(Olden and Holloway, 2009)	
Utilities	Effusion resolved		0.599		(Nafees et al., 2008)
	Effusion not resolved		0.473		(Nafees et al., 2008)
	Under treatment	Talc	0.4		(Olden and Holloway, 2009)
		PleurX®	0.58	0.56-0.59	(Liem et al., 2008)
	Infection		0.4		(Olden and Holloway, 2009)
Costs	Admission	Talc	\$4331		Medicare DRG
	Admission for infection		\$9121		Medicare DRG
	Home visit from nurse		\$94		(Olden and Holloway, 2009)
	Visit with physician		\$110		(Olden and Holloway, 2009)
	Placement of catheter	PleurX®	\$2417		Medicare DRG
	Catheter removal	PleurX®	\$485		Medicare DRG
	PleurX® supplies (10 boxes)	PleurX®	\$750		Denver Biomedical, Inc.

**Figure 6.3.:** Weekly mortality rate and cumulative survival for a cohort with a 26-week life expectancy following the Weibull probability distribution.



**Figure 6.4.:** Conditional probability distribution of *Patient state* [1].

pleural liquid is drained every few days until the effusion is resolved. According to (Tremblay and Michaud, 2006), pleurodesis is achieved in 42.9% of the patients, the median time before catheter removal being 59 days. Even if we lack survival data, we believe it is reasonable to assume pleurodesis with PleurX® also follows a Weibull distributions where the probability density is highest around its median. Therefore, we modeled the probability of pleurodesis with PleurX® using a Weibull distribution with a shape parameter of 2 and a median of 8.42 weeks.

The probability potential behind *Complication* uses the probabilities in Table 6.1 to determine whether a patient develops a complication (infection or death) in that cycle depending on the treatment. We assume complications only can happen in the first cycle after the interventions. This is a simplification, given that patients treated with PleurX® might develop a complication later in the treatment. We deem this simplification assumable given that most complications, which are rare, happen at the beginning of the treatment.

Figure 6.4 shows the probability potential of *Patient state*, which is considerably simplified thanks to the *Pleurodesis* and *Complication* variables. If the patient is alive (that is, she was not dead in the previous cycle and does not die in this cycle) and in the *effusion* state, she might develop a complication, in which case she will progress either to the *infected* or the *dead* state. If on the other hand, the patient achieves pleurodesis in that cycle, she will progress to the *resolved* state. In any other case, she will remain in the *effusion* state.

**Utilities** Effectiveness is measured in quality adjusted life years (QALYs). Each state of the Markov model has a different quality of life, obtained from different sources, as specified in Table 6.1. The exception to this rule is that patients in

the *effusion* state have a different utility depending on whether they are under treatment or not. We could not find in the literature any estimate of quality of life with PleurX® measured on a preference-based scale. Following the decision taken by (Olden and Holloway, 2009), we used the EQ5D score measured for patients with peritoneal dialysis (Liem et al., 2008) as an approximation.

Costs are calculated in 2014 USDs and using US prices. The cost of the initial intervention, i.e., the placement of the catheter or the chest tube, was estimated using Medicare outpatient facility rates CPT 32550 (“Insertion of indwelling tunneled pleural catheter”) and Medicare DRG 188 (“Pleural effusion w/o CC/MCC”) respectively. The cost of infection was estimated using Medicare DRG 186 (“Pleural effusion w/MCC”). For PleurX® patients, we assume that the catheter is removed after a successful pleurodesis or if there has been a complication; the cost was obtained from Medicare outpatient facility rates CPT 32552 (“Removal of indwelling tunneled pleural catheter with cuff”).

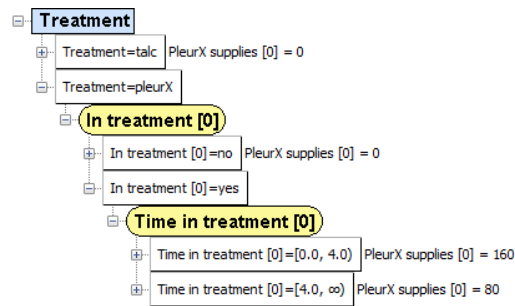
An important assumption of the model, based on the experience of the expert, is that patients under the PleurX® treatment complete drainage twice a week during the first four weeks and once a week for the rest of the treatment. This assumption is in line with that of (Puri et al., 2012), who assumes a weekly drainage for the length of the treatment. The cost of PleurX® supplies, such as vacuum bottles, was calculated from the price of 10 bottle boxes and obtained from (Olden and Holloway, 2009). We also assumed that 50% of the patients with PleurX® would receive home visits from a nurse to assist with each drainage, whereas the rest would do it by themselves or with the help of their family. We believe this assumption, which is also adopted by (Puri et al., 2012), to be a reasonable compromise between common practice in many countries, such as Spain, where patients are instructed to perform drainage at home, and the assumption that all patients will receive home visits from a nurse. In the case of patients whose treatment failed, we assume nurse visits to occur once per week. We analyze the impact of these two assumptions in Section 6.1.4.2. Finally, we assumed that patients in both treatments would visit the physician once a month. The costs of home visits from a nurse and monthly visit to the physician were obtained from (Olden and Holloway, 2009).

All costs were corrected for inflation using the inflation calculator of the website of the US Department of Labor<sup>1</sup>. We used the typical annual discount rate of 0.03, which was converted to its weekly equivalent of  $5.686 \cdot 10^{-4}$  using the following formula:

$$discount\ rate_{weekly} = \sqrt[52]{(1 + discount\ rate_{annual})} - 1.$$

---

<sup>1</sup>[http://www.bls.gov/data/inflation\\_calculator.htm](http://www.bls.gov/data/inflation_calculator.htm)



**Figure 6.5.:** Utility function for the cost of PleurX supplies, assuming we need two vacuum bottles during the first four weeks and only one thereafter.

### 6.1.3. Limitations

Our model has a number of limitations. Survival was modeled with a Weibull distribution whose parameters were estimated by an expert; we could not perform a statistical regression because of a lack of data. However, we think that the impact of a possible imprecision is not significant because the two treatments analyzed in our study do not affect survival, but only the quality of life. Another limitation is that our model does not include the cost of medications or oxygen therapy for patients whose treatment was unsuccessful, due to the complexity to calculate it. Finally, utilities were not obtained from patients with malignant pleural effusions, but from others in similar situations: patients with peritoneal dialysis (Liem et al., 2008) for the utility of patients undergoing PleurX® treatment and patients with non-small cell lung cancer (Nafees et al., 2008) for the rest. Further research on quality of life of patients with MPE undergoing both treatments could improve the precision of predicted effectiveness.


### 6.1.4. Analysis

The result of the analysis presents the estimated total costs and QALYs associated with both treatments. Cost-effectiveness was assessed in terms of the incremental cost per QALY gained using the PleurX® catheter treatment compared with talc pleurodesis. All estimates were obtained as a result of the evaluation of the MPAD.

#### 6.1.4.1. Base case analysis

The base-case analysis was performed assuming a six-month (26-week) life expectancy. The time horizon was set to two years to make sure that only a negligible proportion of the cohort remained outside the *dead* state by the time horizon was reached; given that the cycle length was one week, we evaluated the model for 104 cycles.





**Figure 6.6.:** Temporal evolution of the cohort for both treatments. Only states *effusion* and *resolved* are shown for clarity.

**Table 6.2.:** Results of the cost-effectiveness analysis for the base case.

Treatment	Cost	Effect (w)	Effect (y)	Inc. Cost	Inc. Effect	ICER
Talc	5 562	15.078	0.290	-	-	-
PleurX	6 548	15.582	0.300	986	0.00967	101 914


Figure 6.6 shows the temporal evolution of the cohort for both treatments. Only states *effusion* and *resolved* are shown for the sake of clarity, but given that only a minimal portion of the cohort is *infected*, we can assume the missing proportion of the cohort is in the *dead* state. It is visible how most patients treated with talc achieve pleurodesis and therefore progress to the *resolved* state after the first week, whereas patients treated with PleurX® slowly achieve pleurodesis throughout the first 15 weeks.

Table 6.2 provides the results of the cost-effectiveness analysis for the base case. The treatment with PleurX® results in 0.00967 QALYs more than talc but at an extra cost of \$986, resulting in an ICER (see Section 2.3.1) of \$101 920/QALY, considerably above the cost-effectiveness threshold, which is assumed to be around \$50 000/QALY in the USA (Grosse, 2008).

#### 6.1.4.2. Sensitivity analysis

We performed one-way sensitivity analyses for those parameters showing the highest variability in the literature. We also analyzed the model varying our assumptions in order to assess the impacts of such assumptions in the results of our analysis.

Patients diagnosed with malignant pleural effusion usually have a short expected



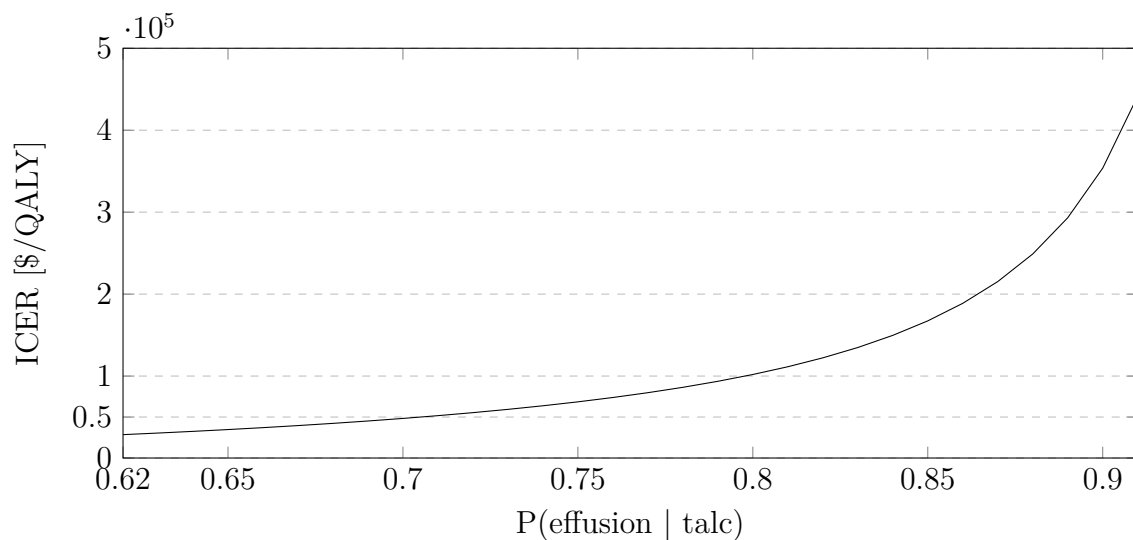
**Figure 6.7.:** Evolution of the ICER for patients with an expected survival ranging from 1 to 52 weeks.

survival but there is great heterogeneity depending on cancer type and other factors (Clive et al., 2014; Tremblay and Michaud, 2006). We performed a one way sensitivity analysis evaluating the model with patients with expected survival ranging from 1 week to a year. In this analyses, shown in Figure 6.7, PleurX® dominates talc (i.e., it is more expensive and less efficient) for patients with an expected survival shorter than nine weeks. PleurX® can be considered cost-effective (the ICER is below \$50 000/QALY) when expected survival is thirteen weeks (around three months) or less.

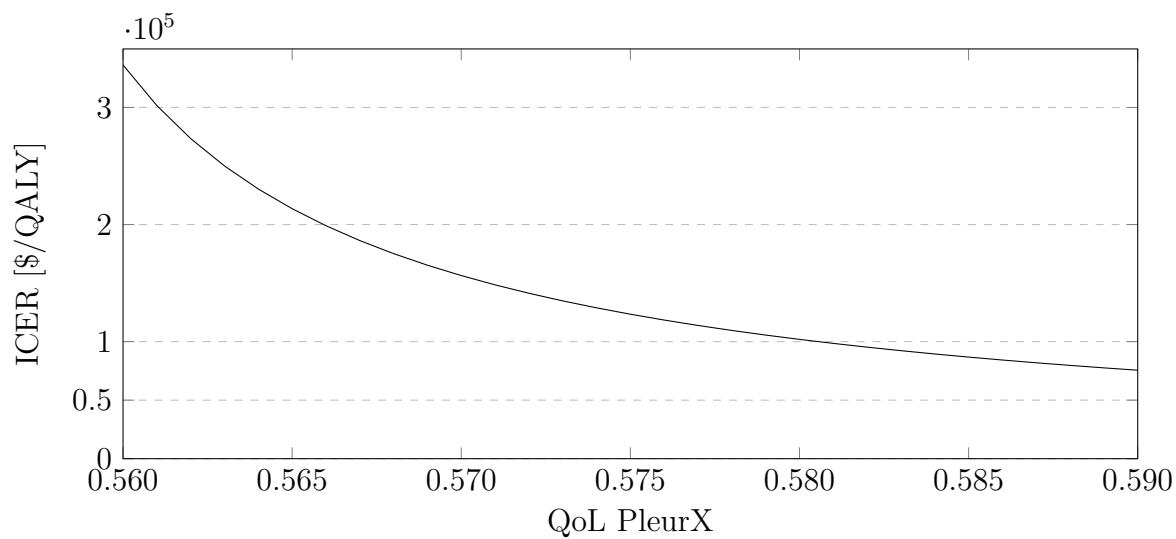
Reported rates of pleurodesis with talc range from 0.62 (Shaw and Agarwal, 2004) to 0.91 (Kennedy and Sahn, 1994). The ICER of PleurX® vs. talc increases significantly as the rate of effusion resolution increases (see Figure 6.8). Eventually talc dominates PleurX, but only at an unlikely rate of 0.95.

In Section 6.1.2.3, we mentioned that we could not find in the literature an estimate of the quality of life with PleurX®, we used that of peritoneal dialysis as a surrogate. In order to assess the validity of this approximation, we analyzed the impact of this parameter on the results of the CEA. This sensitivity analysis (Figure 6.9) showed that talc pleurodesis dominates PleurX® if a utility lower than 0.553 is assumed and that the ICER is highly sensitive to small changes in the value of this variable.


Finally, we analyzed the impact of our assumptions about the number of drainage sessions per week for patients treated with PleurX® and the home visits from a nurse to assist with the drainage. Following the advice of the expert, in the base



**Figure 6.8.:** Evolution of the ICER of PleurX® vs. talc for different effusion resolution rates with talc pleurodesis.



**Figure 6.9.:** Evolution of the ICER of PleurX® vs. talc for different quality of life scores for patients undergoing PleurX® treatment.




**Figure 6.10.:** Evolution of the ICER for patients with an expected survival ranging from 1 to 52 weeks with nurse visits, without them, and only 50% of patients receiving visits (base case).

case we assumed that drainage sessions took place twice a week in the first four weeks and once a week thereafter. However, (Olden and Holloway, 2009) assumes three visits from a nurse per week to assist with the drainage, which implies three weekly drainage sessions. On the contrary, (Puri et al., 2012) assumed a single drainage each week and assistance from a nurse for half of the patients. Figure 6.10 shows how big the impact of the home nursing visits to assist with the drainage is. For example, for a patient with 6 months of expected survival, the ICER is \$184 748/QALY with visits and \$19 080/QALY without them. We should note however that due to the lack of data about the impact of home nursing visits on the quality of life, these ICERs are only approximations. Similarly, Figure 6.11 shows how different frequencies of drainage for patients treated with PleurX® impacts the ICER.

### 6.1.5. Conclusions

Treatment of malignant pleural effusions in patients with an expected survival of 6 months with the indwelling catheter PleurX® is slightly more effective than treatment with talc and it entails a high increase in economic cost, with an ICER of \$101 914/QALY, above cost-effectiveness threshold. Treatment with PleurX® might be cost-effective if the patient's expected survival is thirteen weeks (ICER=\$49 272/QALY) or less, or if we exclude the cost of home visits from a nurse (ICER = \$19 080/QALY). The ICER of PleurX® increases over time because the talc treatment includes a high initial cost (and low quality of life) derived from the chest tube insertion procedure



**Figure 6.11.:** Evolution of the ICER for patients with an expected survival ranging from 1 to 52 weeks assuming drainage sessions once a week (base case), twice a week and thrice a week.

as well as hospitalization, but features low maintenance costs whereas PleurX® incurs in significant costs throughout the length of the treatment. Among the costs included throughout PleurX® treatment, the price of vacuum bottles and home visits from a nurse are the most important.

The novelty of our study with respect to previous cost-effectiveness analyses for MPE (Olden and Holloway, 2009; Puri et al., 2012) is the use of a Markov model instead of an atemporal model. Atemporal models are known to suffer from important limitations: given their inability to represent the elapse of time, time to death was assumed to be of a fixed length for all patients (instead of modeling it with survival analysis techniques such as a Weibull distribution) and they were not able to model the evolution of patients throughout time. Another improvement from previous analyses is that we have incorporated recently published evidence into the model.

The conclusions of our study are in line with those of (Olden and Holloway, 2009) and (Puri et al., 2012): PleurX® is cost-effective for patients with short expected survivals (ranging from six weeks (Olden and Holloway, 2009) to three months (Puri et al., 2012)) but not for patients with longer expected survivals (six (Olden and Holloway, 2009) or twelve months (Puri et al., 2012)), but our results are more precise because we explicitly calculated the ICER of PleurX® vs. talc as a function of life expectancy, as shown in Figure 6.7. Given that expected survival is a major factor in determining the cost-effectiveness of PleurX®, factors predicting survival of patients with MPE should be considered when deciding upon which treatment to apply. For example, based on median survivals reported in (Tremblay and Michaud, 2006), we could deduce that PleurX® is cost-effective for patients with non-small cell lung cancer (median 108 days; 95% CI, 60 to 155 days) and for patients with

ovarian cancer (median 95 days; 95% CI, 41 to 149 days) whereas it is apparently too expensive for breast cancer (median 218; 95% CI, 132 to 303 days) and mesothelioma patients (median 203 days; 95% CI, 162 to 244 days). Our analysis also concluded that treatment of malignant pleural effusions with PleurX® might be cost-effective for patients with longer expected survivals if costs were reduced by excluding home visits from a nurse or due to a decrease in the price of vacuum bottles.

Our study has used pioneering artificial intelligence techniques for CEA, namely a new type of PGM, called Markov process with atemporal decisions (MPAD) developed at the UNED and implemented in **OpenMarkov**, as explained in Chapter 4.

## 6.2. Cost-effectiveness of mammography by risk factors for breast cancer

### 6.2.1. Background

Breast cancer is still an important health issue for its frequency and mortality. It is the deadliest neoplastic malignancy for the female population. According to GLOBOCAN 2012 (Ferlay et al., 2012), currently 232 714 malignant breast tumors are diagnosed in the USA and around 367 090 in the European Union. Incidence of breast cancer in Spain in 2012 was of 25 215 cases. The number of deaths caused by breast cancer in 2012 was 43 909 in the USA, 91 495 in the European Union and 6 075 in Spain.

In 2006, the European Union recommended a biennial mammography for women between 50 and 69 years (European Union, 2003). In 2009, the United States Preventive Services Task Force (USPSTF) updated its recommendations and proposed to extend mammography to women aged 50 to 74 (U.S. Preventive Services Task Force, 2009). However, mammography has its limitations and its effectiveness in women under 50 is still under debate. There is no general consensus amongst international medical organizations about the desirable age of start of mammography screening.

According to some studies, the early detection of breast cancer using mammography reduces mortality between 24% and 29% in women aged 50 or more (Castells et al., 2007). However, a Cochrane review estimated the mortality reduction to be just 15% and the USPSTF estimated it to be around 16% (Gøtzsche and Nielsen, 2009).

Many studies are analyzing the customization of screening, since risk factors, such as family history or genetic factors, along with undesirable side effects of screening (such as false negatives, false positives, etc.) question the effectiveness of current population-wide screening programs. For this reason there is a tendency to build predictive models that might identify the most efficient scenarios and thus tailor screening programs.

Cost-effectiveness analysis is a method to examine the trade-off between costs and effects of alternative treatments. The effectiveness of treatments is evaluated in terms of survival time and quality of life attributable to the treatment. The outcome of any health intervention can then be calculated as the product of the increase in utility that it may cause and the time over which it may be enjoyed.

Cost-effectiveness analyses of early detection of breast cancer in different countries with different strategies have already been published. Results of these analyses show that biennial screening through mammography improves quality of life in patients and is cost-effective for women older than 50 years compared to other alternatives (such as clinical exams or selective screening) and, according to some studies carried out in Brazil, Slovenia or Switzerland, it may be cost-effective even for women

between 40 and 50. In Spain there are very few studies about this issue/problem. . López Bastida et al. (2009) analyzed the cost-effectiveness of a program for the early detection of breast cancer depending on age. Carles et al. (2011) carried out a study in Catalonia, where optimal strategies for breast cancer screenings were estimated using a stochastic model. On the other hand, the technical reports on population-wide screenings in Spain (breast cancer included), elaborated by the Spanish agency for health technology assessment (Agencia de Evaluación de Tecnologías Sanitarias) provide information about the efficiency, effectiveness, and quality of the early detection programs developed by the different regional health agencies.

## 6.2.2. Methods

A literature review was carried out to identify the most relevant publications on the cost-effectiveness of mammography. The search was done in the PubMed and MEDLINE databases, and completed by tracking the studies referenced in them. The studies reviewed include those that analyzed the influence of risk factors on breast cancer mortality and hence on the effectiveness of breast cancer screening for different screening frequencies.

We built an MPAD, loosely based on the model developed by (Schousboe et al., 2011) in the USA, which takes into account different risk factors such as age, breast density, previous biopsies, and familiar history and analyzes the cost-effectiveness of different frequencies (yearly, biennial, every 3-4 years or not at all). The model was built and evaluated using OpenMarkov.

### 6.2.2.1. Structure of the model

Figure 6.12 shows the concise graph of our model. The decision on the frequency of mammographic screening is represented by the atemporal variable *Mamm frequency*.

The variable *Age* represents the age of the patient when she enters the model. It is divided into four ranges: 40 to 49, 50 to 59, 60 to 69, and 70 to 79. Risk factors are represented by chance variables: *Breast density*, *Family history*, and *History of breast biopsy*. Breast density was classified into two levels according to the Breast Imaging Reporting and Data System (BI-RADS). The higher the breast density, the more difficult it is to detect signs of cancer using mammography and therefore the less likely it is to detect cancer at an early stage. BI-RADS divides breast density into four categories, where category 1 represents the lowest density and 4 the highest. In our model, we divided these categories into two groups: low density (categories 1 and 2) and high density (categories 3 and 4). The variable *Family history* has only two values, representing whether the patient has had a case of breast cancer in the family or not. *History of breast biopsy* also takes two values depending on whether the patient has had at least one positive biopsy in the past. As shown in



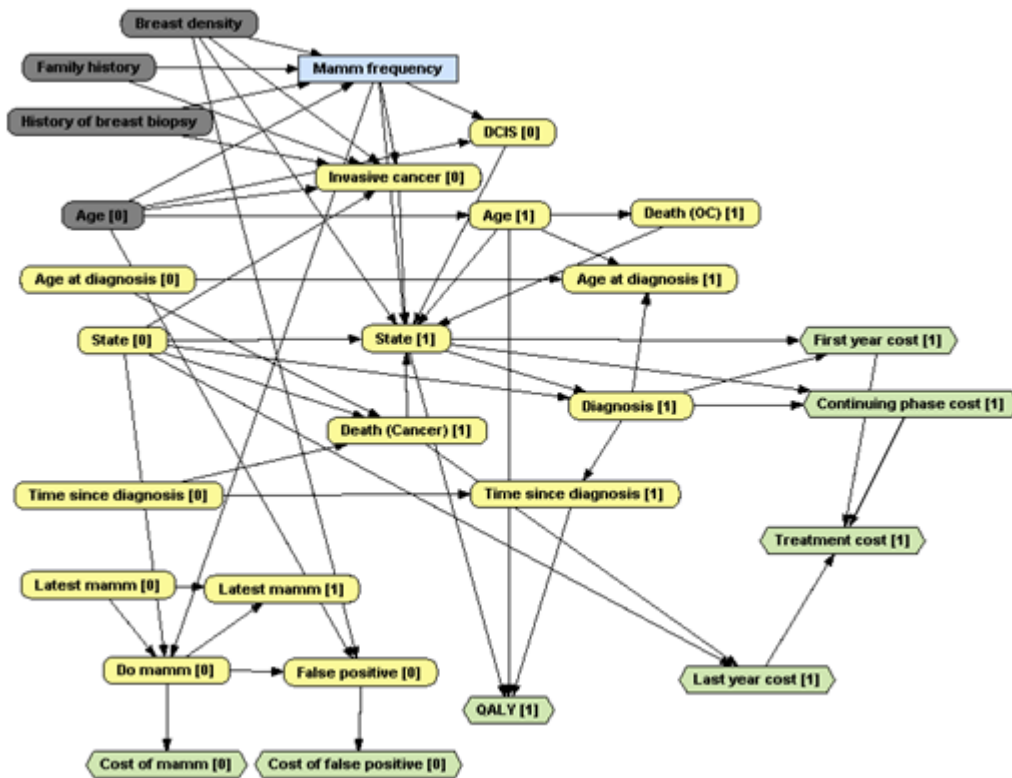



Figure 6.12.: MPAD for breast cancer screening through mammography.



**Figure 6.13.:** State transition diagram from the cancer diagram.

Figure 6.12, these three risk factors, together with *Breast density*, determine the incidence of *Invasive cancer*.

The state of the patient is represented by the temporal variable *State*, which can take six values depending on the patient's current state: all patients start in the *healthy* state; *dcis*, if ductal carcinoma in situ (DCIS) has been detected in the patient; a patient diagnosed with invasive cancer, may be in *local*, *invasive* or *distant* states, depending on the stage of the cancer at diagnosis; and finally *dead*, if the patient has passed away, be it due to breast cancer or other causes. Figure 4.1 shows the possible transitions between these states. It is worth noting that, once the patient has entered a state of cancer, she cannot transition to other states of cancer. This simplification is motivated by the availability of data, or rather the lack thereof. We have included four auxiliary variables in the MPAD to model transitions to different stages: *DCIS*, *Invasive cancer*, *Death (Cancer)* and *Death (OC)*. *DCIS* models the transition from *healthy* to *dcis*. *Invasive cancer* models the transitions from *healthy* or *dcis* to one of the states of invasive cancer. *Death (Cancer)* models the transition from any state of invasive cancer to *dead* due to breast cancer, whereas *Death (OC)* represents the transition from any state to *dead* due to other causes.

*Age at diagnosis* and *Time since diagnosis* are memory variables that modulate the transition probabilities based on factors of the patient's history.

The variable *False positive* represents the probability of a false positive occurring in

a mammography.

Utility or value nodes model the cost of mammography (including the costs derived from false positives), the costs of cancer treatment (which differ depending on the stage of cancer found at the moment of diagnosis and on the stage of treatment), and the health outcomes, measured in QALYs.

### 6.2.2.2. Numerical parameters

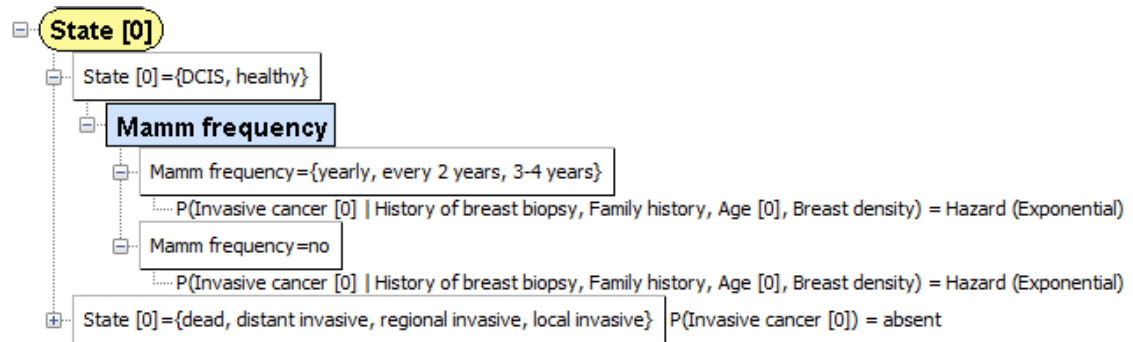
**Diagnosis of DCIS** The probability of a DCIS diagnosis was modeled based on age-dependent rates reported by the Surveillance, Epidemiology and End Results (SEER) of the National Cancer Institute (USA) (Howlander et al., 2011). We assumed in our model that DCIS could be diagnosed in two ways: either through mammography or, in the absence of mammography, by clinical examination. According to Schousboe et al. (2011), the relative risk of a DCIS diagnosis is of 1.393 for patients undergoing a periodical mammography and 0.34 for the rest.

**Incidence of invasive cancer** Schousboe et al. (2011) constructed an invasive cancer baseline risk equation from the SEER data as a function of age:

$$Rate_{IC}(Age) = -3.717 \cdot 10^{-4} \cdot Age + 9.9 \cdot 10^{-6} \cdot Age^2 - 6.53 \cdot 10^{-8} \cdot Age^3 + 4.1328 \cdot 10^{-3}.$$

Following Tice et al. (2008), we assumed that the relative risk of invasive cancer was 1.454 in the presence of a family history of breast cancer in a first degree relative and 0.938 in its absence. We also assumed that the relative risks of invasive breast cancer for a positive and negative history of prior breast biopsy are 1.495 and 0.906, respectively. DCIS is also a risk factor for invasive breast cancer. According to Ernster et al. (2000) and Kerlikowske et al. (2003) the risk of subsequent invasive breast cancer for patients whose DCIS was discovered with mammography is less than when discovered by clinical examination. The relative risk for high breast density was based on Schousboe et al. (2011). Finally, Schousboe et al. (2011) calculated a relative risk of 1.9 for patients with DCIS diagnosed through mammography and 3.4 for the rest.

Applying relative risks to probability potentials is not straightforward, especially when numeric variables such as *Age* are involved in the calculation of the baseline probability. The probability potential for *Invasive cancer* was encoded with two exponential hazard potentials with log relative hazards, one for patients not being screened and the other one for the rest, in order to accommodate different relative risks for DCIS. In the exponential hazard potential with log relative hazards, the log hazard is defined as a linear combination of explanatory variables. In this case, the probability is calculated as:



**Figure 6.14.:** Probability potential for incidence of invasive cancer.

$$Rate_{IC} = \exp(\log(Rate_{IC}(Age)) + \log(RR_{biopsy}) + \log(RR_{family}) + \log(RR_{density}) + \log(RR_{DCIS}))$$

$$P(+invasive\ cancer) = 1 - \exp(-Rate_{IC})$$

**Stage distribution of invasive cancer at diagnosis** In our cost-effectiveness analysis, the decision about mammography frequency mainly affects the stage distribution of invasive cancer at diagnosis. Schousboe et al. (2011) reports stage distributions according to age, mammography frequency and breast density (low or high). These distributions represent the proportion of patients in each state at the moment of diagnosis and have been calculated from the Breast Cancer Surveillance Consortium data. The proportion of patients diagnosed with local cancer is higher for patients going through mammography every two years than that of patients having a yearly mammography, as shown in Table 6.3t. It is also the case that the proportion of patients diagnosed with distant cancer is lower in the former group than in the latter. This is probably caused by a selection bias, since patients were not randomly assigned a mammography frequency; their physicians took that decision, probably assigning a higher frequency to patients at higher risk.

**Mortality from cancer** Mortality from invasive cancer was calculated using SEER data (Howlander et al., 2011), depending on the stage diagnosis and the time elapsed since then, as the complement to the survival ratio. We assumed no cancer-related mortality from DCIS.

**Table 6.3.:** Cancer stage distributions according to age, mammographic frequency, and breast density.

Breast density		-	Low			High		
Mammography frequency		None	1 year	2 years	3-4 years	1 year	2 years	3-4 years
Age 40-49	Local	0.515	0.709	0.713	0.676	0.652	0.657	0.617
	Regional	0.431	0.267	0.269	0.296	0.321	0.323	0.353
	Distant	0.054	0.025	0.018	0.028	0.028	0.020	0.031
Age 50-59	Local	0.484	0.733	0.737	0.703	0.681	0.685	0.647
	Regional	0.443	0.248	0.249	0.276	0.299	0.299	0.330
	Distant	0.073	0.019	0.014	0.021	0.021	0.015	0.023
Age 60-69	Local	0.496	0.768	0.771	0.740	0.720	0.724	0.689
	Regional	0.407	0.209	0.213	0.234	0.254	0.258	0.283
	Distant	0.097	0.023	0.016	0.026	0.025	0.018	0.028
Age 70-79	Local	0.533	0.800	0.803	0.775	0.759	0.762	0.731
	Regional	0.378	0.184	0.187	0.209	0.226	0.227	0.252
	Distant	0.090	0.014	0.01	0.015	0.015	0.011	0.017

**Mortality from other causes** An all-cause mortality risk function created by Schousboe et al. (2011) from 2003 death rates for all females in the USA (Arias, 2007) was used to calculate transition probabilities from any state to *dead*:

$$Mortality_{all\ cause} = 1.2229 \cdot 10^{-3} + 3.09 \cdot 10^{-17} \cdot Age^8$$

**Cost of mammography** As an estimation for the cost of mammography we used the reimbursement fee for 2D mammography screening of 135 USD established by the Centers for Medicare and Medicaid Services (2014). We considered false positives due to mammography as an indirect cost of mammography. The mean cost of follow up after receiving a false positive mammogram was reported to be 330 USD in 2005 (Tosteson, 2008); after correcting for inflation, it amounted to 400 USD in 2014.

**Cost of cancer treatment** We used the treatment costs reported by Yabroff et al. (2008), based on cancer stage at diagnosis and stratified into clinically relevant periods: initial period (first year) following diagnosis, continuing period (subsequent years until the last year of life), and the terminal period (last year of life). These costs have also been adjusted for inflation.

**Health utilities** We used the QALY values calculated in Schousboe et al. (2011), which were based on Lidgren et al. (2007). Different health related values were estimated for the different stages of invasive cancer (localized, regional, and distant) and the time since diagnosis (first year after diagnosis and later years). These values were then weighed using average QALY values for each age group.

### 6.2.3. Analysis

Following Schousboe et al. (2011), we divided the range of age between 40 and 80 into four 10-year periods. This was motivated by the variability of breast cancer incidence and risk factors, such as breast density, changing over time.<sup>2</sup> We analyze the cost-effectiveness of different frequencies of mammography for each 10-year period to allow for re-assessment of mammogram frequency. We ran our model four times, for patients aged 40, 50, 60, and 70 respectively. The most naïve way to do the analysis would have been to use a time horizon of 10 years for each run, but this would considerably underestimate the benefits of higher mammography frequencies. In order to model the health outcomes of each intervention more accurately, we used a 100-year lifetime horizon. First, we ran the model for the oldest cohort, with an average age of 70 years at the start of the model and a time horizon of 30 years. Assuming a cost-effectiveness threshold of \$50 000/QALY (Grosse, 2008), we registered the most effective mammography frequency that was cost-effective for that age range. Then, we ran the model with a starting age of 60 and a time horizon of 40 years, setting the mammography frequency between 70 and 80 to be the most cost-effective as calculated by the previous evaluation. We would then proceed in a similar fashion for cohorts aged 50 and 40. We performed this analysis for cohorts with high and low breast density, with and without a family history of breast cancer, and with and without a previous positive biopsy.

The results of the analysis can be seen in Table 6.4. The most cost-effective option for each case is written in bold type. Yearly mammography is dominated by biennial mammography in all age ranges. This is very likely a result of the selection bias in the data from which we calculated the transition probabilities, which were gathered through years of medical practice, not from a randomized controlled trial. Therefore, instead of randomly assigning patients with a mammography frequency, their physicians made that decision, probably assigning a higher frequency to patients at higher risk.

### 6.2.4. Limitations

One of the limitations of our model is that it does not take all risk factors into account, such as genetic mutations. In the presence of certain genetic mutations, such as those in the genes BRCA1 and BRCA2, the optimal screening frequency might be higher, and the optimal screening type might not be mammography. Another limitation is that, due to the nature of available data, our model is not causal but observational, which might bias the results of our analysis. Another important limitation is the lack of Spanish data to feed the model. This has prevented us to adapt this model to the Spanish population and compare it with the results of other studies done in our country.

---

<sup>2</sup>Over a 10 year period breast density will change significantly in only 10% of women

**Table 6.4.:** ICER for different mammography frequencies according to age range, breast density, family history, and history of biopsy.

Age	Breast density	Family history	Previous biopsy				
				No	Yearly	Biennial	Every 3-4 years
40-49	Low	No	No	base	dom	118776	<b>38441</b>
			Yes	base	dom	66136	<b>15475</b>
		Yes	No	base	dom	72152	<b>18132</b>
			Yes	base	dom	<b>base</b>	dom
	High	No	No	base	dom	70802	<b>42317</b>
			Yes	base	dom	<b>33927</b>	13756
		Yes	No	base	dom	<b>7803</b>	1692
			Yes	base	dom	<b>12422</b>	base
50-59	Low	No	No	base	dom	95789	<b>173</b>
			Yes	dom	dom	<b>36665</b>	base
		Yes	No	dom	dom	<b>40754</b>	base
			Yes	dom	dom	<b>15761</b>	base
	High	No	No	base	dom	<b>36107</b>	base
			Yes	dom	dom	<b>12819</b>	base
		Yes	No	dom	dom	<b>15367</b>	base
			Yes	dom	dom	<b>base</b>	dom
60-69	Low	No	No	base	dom	66158	<b>base</b>
			Yes	dom	dom	<b>28669</b>	base
		Yes	No	dom	dom	<b>32389</b>	base
			Yes	dom	dom	<b>10654</b>	base
	High	No	No	dom	dom	<b>30423</b>	base
			Yes	dom	dom	<b>9477</b>	base
		Yes	No	dom	dom	<b>11917</b>	base
			Yes	dom	dom	<b>base</b>	dom
70-79	Low	No	No	base	dom	158160	<b>8443</b>
			Yes	dom	dom	81353	<b>base</b>
		Yes	No	dom	dom	88974	<b>base</b>
			Yes	dom	dom	<b>39452</b>	base
	High	No	No	base	dom	74682	<b>base</b>
			Yes	dom	dom	<b>30831</b>	base
		Yes	No	dom	dom	<b>35180</b>	base
			Yes	dom	dom	<b>6936</b>	base

### 6.2.5. Conclusions

Our analysis suggests that, contrary to common practice, the decision about the frequency of mammography should be personalized on the basis of a woman's age and well-known risk factors, such as breast density, history of breast biopsy, and family history of breast cancer. Assuming a cost-effectiveness threshold of \$50 000/QALY, a mammography every two years is the most cost-effective intervention in most of the cases. However, a mammography every three to four years is the optimal intervention for women aged 40-49 or 70-79 with a low breast density and with less than two risk factors, and for those with a high breast density and no risk factors. It is also the most cost-effective strategy for women of all ages with low breast density and no risk factors.

The current analysis proves that MPADs are a powerful tool that can be used to build very complex models. Even though we did not have access to the model by Schousboe et al. (2011), which was implemented in *TreeAge Pro*, we can imagine the tremendous effort necessary to build it, while our MPAD can be built in a few hours. Building such a model with spreadsheets would have been almost impossible. Building it with R, STATA or BUGS would have required advanced programming knowledge and many hours of coding and debugging. We also proved it is possible and fairly easy to apply relative risks, which are very popular amongst health economists, when defining probability potentials in in MPADs.




# 7. Cochlear implant programming

## 7.1. Fundamentals

### 7.1.1. Anatomy and physiology of the cochlea

The cochlea, the organ responsible for human hearing, is a spiral-shaped cavity in the bony labyrinth inside the inner ear and is filled with nearly incompressible fluids. Two membranes along the length of the cochlea divide it into three compartments: Reissner's membrane and the basilar membrane. On top of the basilar membrane, the actual receptor organ for hearing (spiral organ or organ of Corti) is found. It contains the auditory sensory cells responsible for converting mechanical waves into electrical signals to be transmitted to the brain. These cells are called hair cells because of their stereocilia, hair-like structures that respond to the mechanical motion of the basilar membrane. There are two types of hair cells: inner hair cells (IHC) and outer hair cells (OHC) which, despite their similarities, serve quite different functions. It is the IHCs, that generate action potentials in the neurons of the cochlear nerve that eventually deliver the perception of sound. The OHCs on the other hand, act as local, frequency sensitive amplifiers, and they highly contribute to the wide range of sound levels that the human ear is able to process and the high frequency resolution it exhibits.

The place on the basilar membrane where there is a higher displacement depends on the frequencies of the sound captured by the ear. This is caused by the mechanical properties of the basilar membrane: it is relatively narrow (100  $\mu\text{m}$ ) and rigid at the base but it becomes wider (up to 500 $\mu\text{m}$ ) and less stiff towards the apex. There is also a gradual increase from base to apex in the length of hair cells as well as a gradual decrease in their diameter. These differences cause each frequency to be mapped to a specific point on the basilar membrane that displaces maximally (Figure 7.1). Maximal displacement of the basilar membrane happens nearer to the base of the cochlea for higher frequencies and nearer to the apex for lower frequencies. Receptor cells at the place of maximal displacement are stimulated the strongest. The spatial distribution of the cochlea depending on its maximal frequency sensitivity is called tonotopy.




**Figure 7.1.:** A: schematic of basilar membrane displacement for pure tones of 2 frequencies  $f_1$  and  $f_2$  where  $f_1$  is lower than  $f_2$ . B: the tonotopy of the cochlear is logarithmic, shifting an octave each 3mm.

### 7.1.2. Hearing loss

Hearing impairment is the most frequent sensory deficit among humans, affecting more than 250 million people in the world. It is estimated that one child in a thousand is born with severe to profound hearing loss. Consequences of hearing impairment include inability to interpret speech, often provoking a reduced ability to communicate, delay in language acquisition, economic and educational disadvantage, social isolation, and stigmatization. Most congenital and childhood-onset hearing losses follow from various disease and injury causes. The leading causes of adult-onset hearing loss are age related hearing loss followed by noise-induced hearing loss. The most common type of permanent hearing loss is sensorineural hearing loss and its most common cause is damage to the structures within the cochlea. Cochlear hearing loss not only reduces the ability to hear soft sounds, but also speech loud enough may sound unclear, distorted, or muffled. Another common symptom of cochlear hearing loss is that, while having a reduced audibility for soft sounds, sounds at high levels are often perceived as having the same loudness as they would for an unimpaired listener.

In general, cochlear hearing loss is related to damage to OHCs and/or IHCs. Their stereocilia may be distorted or destroyed, or the cells themselves dead. OHCs are the most vulnerable, and their damage results in a decreased amplification of soft sounds and deterioration of the frequency selectivity. In a cochlea without functioning OHCs, response thresholds are elevated due to decreased sensitivity. When the IHCs are damaged, sensitivity is reduced, resulting in elevated detection thresholds. When both IHCs and OHCs are severely damaged, sensitivity and tonotopical tuning are greatly reduced.



**Figure 7.2.:** Main parts of a cochlear implant. The external part, consisting of (A) a speech processor (1) that can be worn behind the ear and a transmitter coil (2). The internal part (B) comprises the receiver-stimulator (3) and the electrode array (4), which is positioned in the inner ear (C).

### 7.1.3. Cochlear implants

Cochlear implants (CI) have become a viable treatment for individuals who suffer from severe to profound hearing loss. A cochlear implant is an electrical device that is surgically inserted into the ear and provides the user with a sense of sound electrically stimulating auditory nerve cells. It consists of an internal part (Figure 7.2 B) which is surgically implanted, and a removable external part (Figure 7.2 A) that is usually worn behind the ear.

The external part contains a microphone, which captures sounds from the environment, and a speech processor, which converts these analogue sound waves into a digitally coded signal and sends them via a radio frequency transmitting coil to the internal part, a receiver-stimulator implanted under the skin. The receiver-stimulator decodes the signals transmitted by the sound processor into an electrical stimulation pattern that is delivered to the cochlea through its electrode array.

This array is positioned in the cochlea surgically and consists of a number of electrode contacts (currently ranging from 12 to 22, depending on the manufacturer) running along the length of the cochlea (Figure 7.2 C), which provide direct stimulation of auditory nerve cells. In order to obtain a tonotopical organization for electrical stimulation, electrodes are positioned along the length of the cochlea; each electrode is responsible for conveying the sound at a certain frequency band. This is accomplished by implementing a series of band-pass filters, each passing a portion of the spectral input to their corresponding channel.

**Table 7.1.:** Features of the ICF and parameters across manufacturers

	Cochlear	Advanced Bionics	MED-EL	Neurelec
EDR minimum	T	T	THR	Min
EDR maximum	C	M	MCL	Max
IMR	T-SPL, C-SPL	IDR	-	-
Input gain	Sensitivity	Sensitivity	AGC Sensitivity	Analog Gain
Output gain	Gain	Gain	-	Gain
Input compression	Autosensitivity	AGC	AGC Compression Ratio	-
Output compression	Loudness Growth	-	Maplaw Compression	Volume

### 7.1.3.1. Manufacturers

Four major CI manufacturers have commercially available systems in the European market as per 2015. The market leader is the Australian company Cochlear Ltd., the manufacturer of the Nucleus cochlear implant, the first to be approved by the U.S. Food and Drug Administration (FDA). MED-EL is the Austrian CI manufacturer founded by the scientists who developed the first micro-electronic multichannel CI. Advanced Bionics' (AB) headquarters have long been in California, but the company has been acquired by the Swiss holding Sonova in 2009. Neurelec, the smallest player, is located in France and has recently been purchased by the Danish William Demant holding (Oticon Medical).

### 7.1.3.2. Parameters and intensity coding

Cochlear implant parameters determine how sound waves are converted into electric signals, also known as the intensity-coding function (ICF) (Vaerenberg et al., 2014b). As electrodes are placed along the cochlea, they may face very different conditions when targeting auditory nerve fibers for stimulation: the surviving neural population in the vicinity of specific electrode contacts, or the physical distance to auditory nerve cells may vary considerably throughout the array. This is why some parameters are adjustable per channel or electrode, whereas others affect the behavior of the whole system. Even though the set of parameters varies across manufacturers, they mainly define a set of features of the intensity coding function, common to all manufacturers. These features and the set of parameters available in different devices are summarized in Table 7.1 and described next.

The Electrical Dynamic Range (EDR) is defined as the electrical range between the smallest detectable level (EDR minimum) and the maximum level that is still comfortable (EDR maximum). Both EDR minimum and EDR maximum are defined as separate parameters whose name differs across manufacturers, as can be seen in Table 7.1. All manufacturers offer a parameter to adjust the microphone sensitivity or input gain, which is called Sensitivity in Cochlear and AB devices, AGC Sensitivity in MED-EL, and Analog Gain in Neurelec. In addition, all manufacturers except

for MED-EL give the chance to adjust channel gains, which are applied to the input in Cochlear devices and to the output in AB.

Cochlear implants need to compress the acoustical range available ( $>100$  dB) to normal hearing into an electrical range that is an order of magnitude smaller (a typical EDR is from 6 to 12 dB). This is usually accomplished by two stages of compression: (1) a long term compression of the broadband input signal and (2) an instantaneous compression of the band limited signal within each channel.

Long term compression of the broadband input is typically performed at the CI's front end processing stage (Input compression in Table 7.1). First, a range of sound levels to be processed at any given time is defined: although different terms exist for this range, in this dissertation we use the term Instantaneous Mapping Range (IMR). Input compression adapts the overall sensitivity of the CI such that the IMR is shifted based on environmental sound levels through Automatic Gain Control (AGC) systems. The size of IMR is determined by IDR in AB devices and is fixed in MED-EL and Neurelec (to 55 dB and 85 dB respectively), and its instantaneous position determined in AB and MED-EL by AGC. In Cochlear, on the other hand, T-SPL and C-SPL determine the lower and upper bounds of the IMR, thus defining its size and position, but the position is also affected by microphone sensitivity and Autosensitivity (similar to AB's AGC).

Instantaneous compression (Output compression in Table 7.1) is performed within every channel using highly compressive logarithmic mapping functions. The mapping function in all CI systems is approximately linear, but Cochlear, MED-EL, and Neurelec provide an additional parameter to adjust the percentage of EDR that is allocated to higher/lower intensity ranges, thus curving the function.

Finally, pulse width is configurable in AB and Cochlear devices, but with different effects. In Cochlear, doubling the pulse width will cause the channel to deliver twice the charge, whereas in AB it will remain equal as the pulse amplitude is automatically decreased.

## 7.2. Cochlear implant programming

After surgical implantation, the parameters described above must be appropriately adjusted for the individual. This process is called cochlear implant programming or fitting, and its aim is to ensure that the electric pattern generated by the internal device in response to sound stimulation yields an optimal auditory percept, both in terms of performance and comfort. Even if cochlear implants have been available for more than 25 years and despite the general agreement that fitting is crucial in obtaining good results (Shapiro and Bradham, 2012), to date there is neither well described and commonly adopted good clinical practice for this task nor evidence-based materials to distinguish efficient procedures from less efficient ones (Vaerenberg et al., 2014c).

A set of values assigned to the parameters of a CI at a given point in time is called a *MAP*<sup>1</sup> or a program. Therefore, we could define CI programming as the search for the optimal MAP. In fact, some also denote CI programming with the term MAPping (Pedley et al., 2007).

Fitting is usually restricted to setting two basic psychophysical measures for each electrode: the auditory threshold (T level) defined as the softest level at which the patient is stimulated more than 50% of the time, and the most comfortable level loudness level (M/C level), defined as the loudest sound that is still comfortable. The two most popular approaches to threshold measurement are behavioral and objective measures. Behavioral measurements are taken through interaction between the audiologist and the patient: the patient is instructed to react when a sound is audible for T levels and to express discomfort to set M/C levels. This might turn to be a time-consuming task, especially in children.

Objective measures arose as a means to avoid the error introduced by the patient's subjective response, by measuring electrophysiological responses (neural or muscular) elicited by a stimulus. Most popular objective measures include evoked stapedius reflex threshold (ESRT) and (ECAP) and they provide the audiologist with a confirmation of an audible stimulation level for a particular electrode. Many studies have investigated whether objective measures could be used for predicting T and M/C levels (Brown et al., 2000; Hughes et al., 2000). The general conclusion of these studies is that such measures might be indicative of T and M/C levels but there is no consensus on the usability of this information. While some claim correlations are not strong enough to be used for fitting (Smoorenburg et al., 2002), others defend the benefits of fitting implants with thresholds derived from ECAPs (Botros and Psarros, 2010). According to Vaerenberg et al. (2014c), 59% of the CI fitting centers that took part in the survey use ECAP measurements, mainly to set the profile of T and M/C levels across the electrode array.

As part of the CI programming process, electrodes can be disabled or dropped by the fitting expert if they are showing poor performance. Poor performance in specific electrodes can be caused by the area of the cochlea where the electrode lays being damaged or containing an unusually low population of functioning hair cells. As a consequence of disabling an electrode, the frequency bands assigned to the rest of electrodes stretch and shift, so that the whole spectrum is still covered.

This traditional approach to CI programming has remained essentially unchanged since the beginning. Generally, it is oriented towards patient comfort and not measurable performance (Vaerenberg et al., 2014c). Recently, some clinicians have advocated for setting and reaching measurable performance targets (Vaerenberg et al.,

---

<sup>1</sup>We decided to write MAP in capital letters in this dissertation for two main reasons: because Cochlear Ltd., the original proponent of the term, and most of the scientific community use this form and because we think that the documents produced by cartography would be a misleading metaphor of the concept. For all we know, MAP might be an acronym whose meaning has been lost.

2014a; van der Beek et al., 2015), thus stressing the importance of cochlear implant programming.

### 7.2.1. Measuring performance

Hearing performance can be divided into three categories: sound detection, discrimination and identification. In test conditions, detection is measured by asking the patient to react to any auditory stimulus. Discrimination tests measure the ability to tell different sounds apart. Identification does not depend exclusively on the hearing performance, but also on cognitive processes. The following tests are designed to measure these three aspects of hearing.

#### 7.2.1.1. Tone audiometry


The aim of tone audiometry (also called simply audiometry) is to find the auditory threshold, defined as the lowest sound level detectable. The patient is presented with a series of pure or warble tones of different loudness and asked to react to sound. The test is repeated with tones centered at different frequencies distributed across the spectrum covered by cochlear implants (usually 250, 500, 1000, 2000, 4000, 6000 and 8000 Hz) in order to measure the audiometric thresholds for different frequencies. According to Skinner et al. (2002), sound threshold levels from 250 to 6000 Hz must be below 30 dB for CI users to be able to perceive soft speech.

#### 7.2.1.2. Phoneme discrimination

Phoneme discrimination is a special case of spectral discrimination (Govaerts et al., 2006) that tests the patient's ability to tell two phonemes apart. Phonemes are presented at 70 dB in an oddity paradigm: one phoneme is presented repeatedly (background phoneme) in fixed intervals and every now and then the background phoneme is replaced by the other phoneme (stimulus phoneme). The patient is instructed to react to the stimulus phoneme. This test assesses the patient's frequency resolving power, which is one of the cochlea's most important functions.

#### 7.2.1.3. Loudness scaling

The loudness scaling test is performed using one-third octave narrow band noises centered at 250, 1000 and 4000 Hz. A 1876ms stimulus is presented twice at each level and scored on a visual analogue scale ranging from 0 (inaudible) to 6 (too loud). Levels are randomly presented at 5dB increments between 30 and 80dB. Figure 7.3 shows the result of a loudness scaling test at 4000 Hz. Each dot represents the loudness perception of the patient at different loudness levels. The loudness scaling of a patient is deemed normal when all dots are inside the gray area.



**Figure 7.3.:** The result of a loudness scaling test in Audiqueen©.

#### 7.2.1.4. Speech audiometry

Speech audiometry tests the patient’s ability to correctly identify phonemes and words. The test involves the acoustic presentation of a list of short CVC (consonant, vocal, consonant) words; e.g., “dog”. The list of words is chosen as to cover the set of phonemes in the patient’s language. Words are presented at different loudness levels: 40, 55, 70, and 85 dB. Scoring is usually done based on percentages of correctly guessed words and phonemes.

### 7.2.2. Previous applications of AI to CI programming

Since the commercial introduction of cochlear implants, there has been a few attempts to apply artificial intelligence or machine learning techniques to fitting. For example, (Chang et al., 2001) built and trained neural networks to mimic patients’ vowel recognition ability in order to simulate how their ability would change for different values of a parameter. Unfortunately they did not verify whether CI users showed similar improvements in vowel recognition as the trained neural networks. Below we provide an overview of the most interesting attempts.

#### 7.2.2.1. AutoNRT™

AutoNRT™ was developed by engineers in Cochlear Limited and presented as an automated system to measure ECAP thresholds “via machine intelligence” (Botros et al., 2006). Its authors claimed it was “the first machine intelligent system (...) within the whole cochlear implant industry” (Botros, 2010). It consists of two classification trees<sup>2</sup> automatically built from a dataset with the C5.0 algorithm (Quinlan,

<sup>2</sup>Even though classification trees are often called “decision trees”, they should not be confused with the decision trees used in probabilistic decision analysis (Raiffa and Schlaifer, 1961; Raiffa, 1968).



1993). According to its authors, “performed as accurately as the ‘average’ clinician” when determining ECAP thresholds and saved time.

### 7.2.2.2. The HÉVÉA project

The objective of the HÉVÉA project (Collet et al., 2009), funded by the French Ministry of Health, was the optimization of CI fitting using evolutionary algorithms. Evolutionary algorithms try to solve optimization problems by defining a population of solutions whose quality or fitness improves over time through crossover (mating) and mutation by means of the “survival of the fittest” principle. The individuals of the population were defined as sets of values (two per electrode) for the T and M/C levels, randomly chosen within the range of behaviorally measured thresholds. The evaluation function defined in order to assess each individual’s fitness was a speech understanding test, therefore needing interaction with the patient. Therefore each individual in the population (each combination of T and M/C levels) needed to be evaluated by performing a speech test on the patient and thus, the number of individuals in the populations as well as the number of iterations allowed was very limited. For the first set of experiments (Bourgeois-Republique et al., 2005), the evaluation function was a test consisting of ten sentences extracted from the corpus elaborated by Pr. Lafon (Lafon, 1964), in order to reduce testing time. Those experiments were conducted with a single patient whose score seemed to improve after a few iterations, but only in about half of the algorithm runs. Further tests (Legrand et al., 2007) done with more elaborate evaluation tests in two different patients failed to prove any improvement, allegedly because longer testing times made impossible (due to patient’s fatigue) to perform sufficient iterations for the algorithm to converge.

### 7.2.2.3. FOX

FOX (Govaerts et al., 2010), which stands for “Fitting to Outcomes eXpert”, is a software tool developed by Otoconsult, a software company associated with the Eargroup, a hearing clinic located in Antwerp (Belgium). It is an intelligent agent that assists the audiologist by recommending a set of adjustments to the CI parameters. A fundamental principle of FOX is to base parameter adjustments on performance (auditory test results) rather than on comfort measurements, which is the traditional approach. FOX has been under development since 2001 and the first working prototypes were used in 2002. Its development was motivated by the desire to improve the quality of programming, systematize it, and reduce its variability, while also reducing the required resources by automating the process.

FOX’s knowledge base consists of a set of rules encoded by its main author, an expert clinician, to cover the different situations faced in CI programming. Each rule is divided into a condition and an action. The condition determines the set

of requirements the input has to fulfill in order for the rule to be fired and as a consequence, the action to be taken. The input consists of the current set of parameter values (the current MAP) and the set of auditory test results (outcomes) achieved with the current MAP. The output is a recommendation of adjustments of the parameters to try to bring the patient's auditory performance closer to target. Whenever more than one rule is fired, conflict resolution takes place by averaging the parameter adjustments proposed by each of the rules fired.


Another special feature of FOX is that instead of using behaviorally measured thresholds or ECAP measurements for the initial fitting, it provides the patient with a set of incrementally louder MAPs for the initial period of adaptation. These MAPs are calculated through statistical analysis of the set of MAPs whose recipients reached outcome targets.

The results achieved through the use of FOX have been well documented. A study (Battmer et al., 2014) showed that computer-assisted fitting since switch-on, that is, from the very beginning, can be successfully achieved with FOX with significant time savings during the first two weeks and including more measures of performance than with the standard fitting approach. Another study, this time with long-term CI users (Buechner et al., 2014), claimed that lower audiometric thresholds were achieved after fitting with FOX. Improvements in median scores of speech perception scores turned out not to be statistically significant, perhaps due to the small number of patients in the study. However, being able to match the results of current world experts is already a considerable achievement for an automated approach. Finally, Vaerenberg et al. (2014a) showed that patients were able to reach 57% of the targets at switch on with the map provided by FOX, and after fitting with FOX the percentage rose by 24%.

### 7.3. The Opti-FOX project

The objective of the Opti-FOX project (Szlávik et al., 2011) was to improve the performance of FOX by overcoming its limitations. The first approach explored consisted in using machine learning techniques to infer the complex relation between CI parameters and hearing performance from historic data. However the domain proved to be too complex for the size of the available data set, and models learned from data using different algorithms were not able to meet the expectations.

Because of the difficulty of building a model using a fully data-driven approach, the leaders of the Opti-FOX project considered the possibility of applying PGMs, which are suited to combine expert knowledge with data to perform reasoning in uncertain domains. This was the reason for inviting the researchers of the CISIAD (UNED, Spain) to join this project. This section describes the probabilistic model we built for tuning CIs, which is a particular case of the tuning networks described in Chapter 5.



**Figure 7.4.:** Prototype of the Opti-FOX network.

### 7.3.1. Description of the model


This section has been omitted from the public version of this dissertation because it concerns proprietary or confidential information subject to IP.

### 7.3.2. Evaluation and results

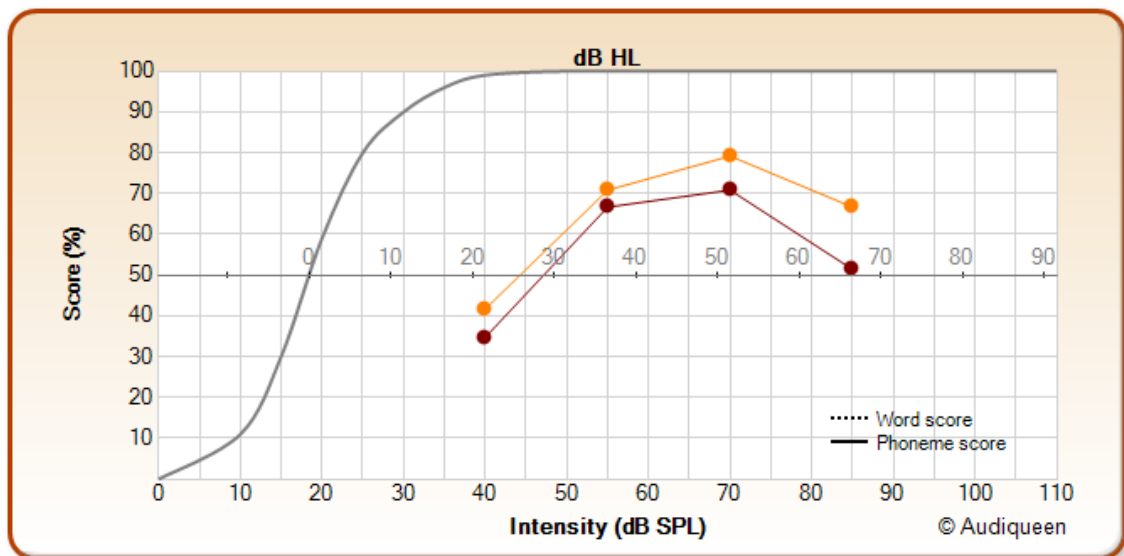
The first version of our network only covered the lower frequencies (up to 1000Hz); we later extended it to the mid and high frequencies by creating new instances of the above mentioned classes. Various (around 10) tests were done with data from real patients with poor performance in the lower frequencies. In these tests the adjustments recommended by our model were evaluated by comparing them against those of FOX, and those of the expert.

After the expert fine tuned some of the probabilities, the model reached the point where it usually agreed with the outcome of FOX. Also, in most cases the outcome has been what the expert expected but even in the case of unexpected recommendations, the expert described them as being “smart”, “intelligent”, or “worth trying”.

The prototype was tried only once on a CI user: a middle-aged woman whose CI device had been fitted by expert audiologists at the Eargroup clinic, supported by FOX, for more than two years; even though her audiogram was inside the range of normality, she showed a poor performance in speech audiometry. On July 31st 2012, when those audiologists programmed her CI following the recommendations of our probabilistic model, both her loudness perception and her ability to recognize speech increased, as shown in Figures 7.5 and 7.6, respectively. Obviously, it is not possible to draw conclusions from one particular case, but it was a promising result.



**Figure 7.5.:** Loudness scaling of the patient mentioned in the text, before (red) and after (orange) applying the MAP proposed by the Opti-FOX prototype. The improvement is considerable given that all the scores from 30 to 65 dB in the orange curve are inside the gray area.



**Figure 7.6.:** Speech audiometry for the same patient before (red) and after (orange) applying the MAP proposed by the Opti-FOX prototype. The vertical axis represents the percentage of phonemes recognized. The patient performed consistently better across all loudness levels (differences of 5-15%).

### 7.3.3. Conclusions

In spite of the satisfactory results described above, the expert expressed the feeling that this model had reached a plateau of performance. The complexity of the model had grown to the point where every little refinement took him hours of work, despite OpenMarkov's reasoning explanation facilities. Besides, the expert expressed his concern about the model's frequent recommendation of increasing the T levels for different frequency bands, even when these were already high.

A big limitation of the model was its lack of precision in the size of the changes recommended. The model would recommend the direction of the change (up or down), but it would not specify by how much. Even if we deemed this limitation affordable at first, it was increasingly clear as we learned about the domain that this was a key factor of the quality of the advice given.

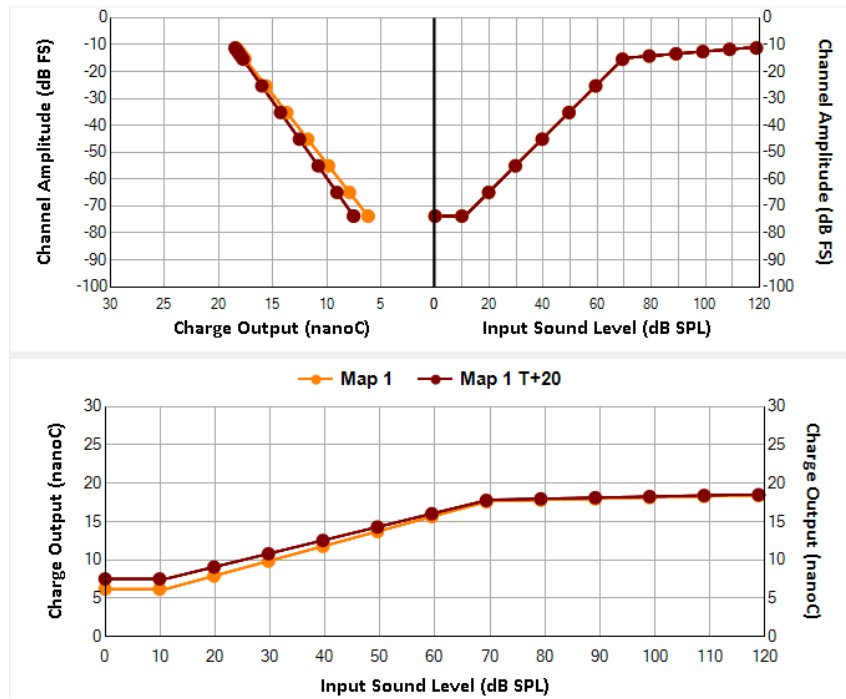
Another limitation of the model was that it only considered the set of outcomes that was available for the current map, ignoring all previous test results. This is a considerable shortcoming, given that due to time constraints, patients do not usually go through the whole battery of tests in a single visit.

Also, there was a growing concern that given the features of the network, learning from data was not going to be effective. The PGM contained a series of hidden variables with a high number of parents, which led to very complex conditional probability distribution. Learning the conditional probability function of a hidden variable is difficult enough, but learning a number of them with such complex conditional probabilities would demand a vast amount of cases.

## 7.4. Hearing Minds

Almost a year after the end of the Opti-FOX project, we continued our work on the decision-support system for fitting CIs within the Hearing Minds project, funded by the Marie Curie Action "Industry-Academia Partnerships and Pathways" Programme. In the meantime, Vaerenberg et al. (2014b) published a paper on intensity coding in CIs (ICCI) that offered a clear insight into how CI parameters determined the electrical charge provided by the electrodes for all four manufacturers mentioned in Section 7.1.3.1. They presented a software application for visualizing the intensity-coding function, also allowing to visualize the impact of any change in the many MAP parameters. Figure 7.7 shows a screenshot of the software application presented in (Vaerenberg et al., 2014b).

Even if the authors of the paper admit that it is a simplification, this model provided a more accurate prediction of how changes in the parameters could affect intensity than we could have represented with a PGM. Therefore, we decided to replace the layer of the PGM that calculated the changes in electrical intensity given changes



**Figure 7.7.:** Screenshot of the ICCI software showing how a 20 point increase in the T parameter affects the output charge (in an Advanced Bionics device). The bottom graph summarizes the upper two: its horizontal axis represents the input sound level in decibels and the vertical axis represents the electrical charge output in nanocoulombs. The two lines in the bottom graph represent the intensity-coding function of two different MAPs.

in the parameters, which happened to be its more complex part. Moreover, we realized that based on the expert's knowledge, we could build predictive models to predict test results based on past test results and changes in energy as calculated by Vaerenberg et al. (2014b).

The inspiration for the new approach was to combine this physical simulation model with probabilistic graphical models to predict the impact of electrical changes on the test results probabilistically.

### 7.4.1. Description of the model

This section has been omitted from the public version of this dissertation because it concerns proprietary or confidential information subject to IP.

### 7.4.2. Results

The decision-support system has been used in the last few months to assist in the fitting of CIs in the Eargroup hearing clinic in Antwerp, Belgium. It has been tested in 15 experienced patients, none of which were newly implanted. They all had been through the fitting process before and most of them had been coming to the clinic for periodical fitting sessions for years. Due to the time and effort needed from both patients and audiologists to perform a full batch of tests, most of the patients were only tested with the MAPs recommended by our model. For four of these patients our model considered that the risk of worsening the test outcomes outweighed the benefit that a change in the map might contribute; therefore no further action was recommended.

All patients but three had audiometries considered to be on target, that is, the audiometric thresholds measured were equal or below target across all frequencies except for three patients. The audiometries of two of those three patients improved to thresholds that were on target (one of them in two iterations), whereas two thresholds of the remaining patient remained out of target after an iteration.

Three of the fifteen patients failed to discriminate at least one phoneme pair. Two of them failed to discriminate a single contrast, and one of them failed three. From the two patients who failed a single contrast, one managed to successfully discriminate it after the system's intervention, and the other one kept failing. The patient who failed to discriminate three contrasts was able to discriminate them all with the MAP proposed by our model.

Eight patients had at least one loudness scaling score that was off the acceptable range in at least one frequency. From those eight patients, only three patients kept at least one score off target after the improvements to their MAPs recommended by our system. Put another way, from the 16 scores that were off target and the 10

that were almost on target, only 7 remained off target and 4 almost on target after using the MAP proposed by our system.

Improvements are more subtle in the speech audiometry results. Thirteen speech audiometries were conducted using the MAPs recommended by our system. Only two of them achieved an overall improvement for different loudness levels that improved the average score across all loudness levels by more than five points. Other five speech audiometries included an improvement of more than 10 points in at least one loudness level. The average scores were slightly higher after our intervention, but the improvement was not statistically significant. Another positive aspect is that the differences between scores in neighboring loudness levels are smaller (a median of -4.2 both for the difference between scores at 40 dB compared to 55 dB and 55 dB compared to 70 dB) after applying the changes proposed by our model.

## **7.5. Conclusions and future work**

This section has been omitted from the public version of this dissertation because it concerns proprietary or confidential information subject to IP.



**Part IV.**  
**CONCLUSION**



# 8. Conclusions

## 8.1. Main contributions

### 8.1.1. New types of PGMs

We have studied three new types of PGMs: DANs, MPADs, and tuning networks. Even if the development of the first two types of models were defined at the CISIAD of the UNED before the beginning of my doctoral research, I have contributed to both. Additionally, I have developed a new type of model, called tuning networks.

DANs are relevant for decision analysis because they can represent and solve not only symmetric decision problems—with the same ease as influence diagrams (IDs)—but also asymmetric problems for which IDs are inappropriate. As we argued in Section 3.1.5, in the case of asymmetric problems DANs are more suitable than any other formalism proposed so far. For this reason we think that DANs might replace IDs as the standard tool for decision analysis.

My contributions to the development of DANs are as following. First, when implementing the conversion of DANs into equivalent decision trees in `OpenMarkov`, I found a few loose ends in the definition of DANs and in the conversion algorithm. The current definition of DANs, as detailed in the paper that we have submitted to *IEEE Transactions on Pattern Analysis and Machine Intelligence* (impact factor: 5.694), contains some of the modifications resulting from my work. Second, I designed and implemented a new evaluation algorithm, which is described in Section 3.3 and in that paper. We compared the efficiency of this algorithm with the expansion of an equivalent decision tree and with another algorithm for DANs developed and implemented by Prof. Manuel Luque, another member of our research group. The results, shown in Section 3.4, indicate that the algorithm proposed in this thesis performs significantly better than the other two for large DANs; in particular, it has been able to resolve the  $n$ -tests problem with  $n = 12$ , which is remarkable, considering that Jensen and Vomlelová (2002) only managed to get to  $n = 9$  in a similar problem (an UID with no structural constraints) and its a problem whose complexity grows superexponentially. To our knowledge, no one had ever solved such a large asymmetric decision problem.

With respect to MPADs, I completed the implementation of algorithms for cost-effectiveness analysis (CEA) and reimplemented sensitivity analysis almost from

scratch. We were able to finally evaluate the HPV model<sup>1</sup>, built by CISIAD two years before. During my three-month stay at the Centre for Health Economics of the University of York in 2013, I implemented as MPADs several models published in the literature based on their corresponding Excel and R files: the cost effectiveness studies of lamivudine/zidovudine combination therapy in HIV infection (Chancellor et al., 1997), total hip replacement (Briggs et al., 2004), and cotrimoxazole prophylaxis in HIV-infected children in Zambia (Ryan et al., 2008). Representing these models required the analysis and implementation of new types of potentials, as well as an extension of some of the existing algorithms. We proved that MPADs are a powerful tool to build and evaluate models for CEA and have significant advantages compared with Excel and R. My PhD supervisor and I gave seminars at the Centre for Health Economics of the University of York and at the School of Health and Related Research of the University of Sheffield, two of the centers leading the development of new methods for health technology assessment. I collaborated on a paper, co-authored by members of the CISIAD and of the University of York, which presents MPADs to the health economics community; it will be submitted to *Medical Decision Making* (impact factor: 2.698).

During the participation of our research group in the Opti-FOX project, financed by the EU 7th Framework Programme, we proposed a new type of PGM, called tuning networks, as no existing PGM suited our needs. We also defined a new canonical model that represents the effect of multiple parameters on a single outcome. Finally, during my stay in the University of Aalborg I designed and implemented object-oriented probabilistic networks as an extension of object-oriented Bayesian networks. We believe tuning networks, along with the tuning model and combined with the object-oriented probabilistic networks, will be very useful to anyone interested in modeling a system with a high number of tunable parameters. These contributions were presented at the AIME2013, a conference rated as CORE A.

### 8.1.2. Application of MPADs to cost-effectiveness analysis

We carried out two CEAs with MPADs. First, we analyzed the cost-effectiveness of the PleurX® catheter vs. talc pleurodesis for the treatment of malignant pleural effusions, in collaboration with an expert surgeon. Our results showed that PleurX® dominates the talc treatment for patients with a life expectancy shorter than nine weeks and is cost-effective for patients with a life expectancy lower than 14 weeks. For patients with longer life expectancy, talc is the most cost-effective option. This paper will also be submitted to *Medical Decision Making* after the paper on MPADs.

We also built an MPAD to analyze the cost-effectiveness of mammographic screening for different age ranges and for women with different risk factors. The results of our analysis mostly match those of Schousboe et al. (2011), which was heavily based on the correlations observed in the SEER database of the National Institute

---

<sup>1</sup>Available at <http://www.probmodelxml.org/networks/mpad/MPAD-HPV.pgm.xml>.

of Health (USA). Unfortunately, we could not adapt the model to Spain, as we originally intended, because their model is not causal and we lacked the corresponding observational data for the Spanish population given that there is no database similar to that of the SEER study in our country. If the original model were causal, we might have combined the data available for Spain mixed with the those of the original model. Nevertheless, we believe that the network we have built is a good example of the ability of MPADs to represent complex models.

These two MPADs are available online<sup>2</sup>. The possibility of reproducing the results of the analyses is even more important than in the case of algorithms, because there are well-known cases of very influential papers based on models that later proved to be wrong Herndon et al. (2013); Baggerly and Coombes (2009). The number of papers whose conclusions are derived from models is ever increasing, especially in the field of CEA, and the scientific community is increasingly aware that research should be reproducible, but only in very exceptional cases is the corresponding model publicly available.

### 8.1.3. Assisting cochlear implant programming

We have built a decision-support system for fitting cochlear implants (CIs) in collaboration with Otoconsult, a software company located in Antwerp (Belgium) specialized in audiological products. First, inside Opti-FOX, a project funded by the European Union (FP7-SME-2010-1 262266), we built a tuning network that relates the performance of the CI with the audiological outputs and recommends the set of adjustments to the parameters of the CI estimated to lead to optimal performance. Thanks to the success of this model, the collaboration of Otoconsult and the UNED continued with the support of a new European grant (Marie Curie Action “Industry-Academia Partnerships and Pathways”, FP7-PEOPLE-2012-IAPP). In this project we have built a new version of our decision-support system that combines a physical model of the CI with a probabilistic graphical model and has been integrated into a new version of FOX—described in Section 7.2.2.3.

The great variability in current fitting practices and the fact that it is usually comfort driven (Vaerenberg et al., 2014c), lead experts to believe that many CI users suffer from suboptimal performance. A study on the use of the old version of FOX (Vaerenberg et al., 2014a), which was based on deterministic rules, proved to reduce the variability and improve the results achieved through manual fitting. According to the authors of FOX, the decision-support system described in Section 7.4, offers better recommendations than the rule-based expert system, and for this reason it is currently used at the Eargroup, a hearing clinic associated with Otoconsult, as the main tool to assist in the fitting of CIs. Considering that over 320 000 people in the world have received cochlear implants (NIH, 2013), we are confident that

---

<sup>2</sup>At [www.probmodelxml.org/networks](http://www.probmodelxml.org/networks).

the decision-support system developed in this doctoral thesis will have a significant impact in the hearing performance of many people.

### 8.1.4. OpenMarkov, an open-source tool for PGMs

The new types of models and the algorithms developed in this doctoral thesis have been implemented in the open-source program **OpenMarkov**, that can be downloaded from [www.openmarkov.org](http://www.openmarkov.org). This is important for two main reasons. First, because the experiments that compare the efficiency of the algorithm for DANs (Section 3.4) can be reproduced by other researchers using our source code. As mentioned above the scientific community is increasingly convinced of the need for reproducibility in research. Second, because these contributions are now available to the communities of uncertainty in artificial intelligence, decision analysis, and health technology in a mature program that can be used for free, even by those that do not have a strong background on PGMs. In this aspect the models and algorithms described here differ from the ones that have never been implemented.

For these reasons, we think that **OpenMarkov** is a major contribution of the CISIAD to the field of PGMs. It is used at universities, research institutes, and private companies of at least 20 countries. My contributions to **OpenMarkov** range from refactoring and debugging existing code to implementing several new features, as explained in A.

One important contribution is the functionality for learning Bayesian networks interactively, which was the subject of my master thesis. During my doctoral research, in addition to improving the module for learning Bayesian networks with new features and refactoring and debugging the algorithms, I have co-authored a paper that describes this functionality and that is currently under evaluation at the *Journal of Statistical Software* (impact factor: 3.801).

## 8.2. Future work

### 8.2.1. New types of PGMs

In order to make DANs more appealing, we intend to offer two important features for their evaluation. First, we would like to develop an algorithm to generate a compacted representation of the optimal strategy of a DAN. The optimal strategy is best represented with a policy tree in which the order of variables in the tree minimizes its size. This task would consist in the adaptation of optimal policy trees described in Prada (2014) to DANs. Second, it would be interesting to develop reasoning explanation techniques for DANs, similar to the ones existing for BNs and IDs in Elvira Lacave et al. (2007). Explanation of the reasoning is important

for debugging purposes and to overcome users' reluctance to accept the advice of a model.

Another field of application of DANs is CEA. IDs have already been used for CEA (Baio et al., 2006; Arias and Díez, 2015), usually in symmetrized decision problems. The main advantage of DANs is that they can also represent order and structure asymmetry in a natural way, which can prove crucial in problems involving more than one decision; for example, when several tests are available or when more than one treatment can be applied sequentially.

Regarding MPADs our objective is to make them a well-known formalism for CEA so that they gain a prominent space in the modeler's toolbox. The first and most important step will be to publish the paper presenting MPADs in the *Medical Decision Making* journal, in order to introduce them to the health technology assessment community. Once the MPADs have been presented, we will perform cost-effectiveness analyses for other medical problems. The more models we build and publish, the more arguments we will have to convince the community of the advantages of using MPADs. Another important line of work is the development of new features in `OpenMarkov`, such as new potentials, to make sure that modelers do not miss the most important tools they normally use to build models. In this line, we intend to work towards making `OpenMarkov` more easily extensible, so that any modeler with notions of programming is able to develop any type of potential needed when building a model.

Finally, there are software tools that integrate Bayesian inference and cost-effectiveness analysis (Abrams et al., 2004; Baio, 2012; Welton et al., 2012). Therefore a very interesting line for future research is the integration of `OpenMarkov` with a some of the packages in the BUGS family, such `OpenBUGS` or `JAGS`, thus combining the statistical power of the latter with the facilities for model building and analysis of the former.

### 8.2.2. Cost-effectiveness analysis with MPADs

We believe the value of our CEA of the PleurX® catheter compared to chest tube placement with talc pleurodesis could increase if we combine our current results with predictive factors for the life expectancy of patients. Thus, we could recommend treatments based on these predictive factors, such as type of tumor, the lactate dehydrogenase level in the pleural fluid, and the neutrophil-to-lymphocyte ratio Clive et al. (2014).





**Part V.**  
**APPENDICES**



# A. Functionality implemented in OpenMarkov

This appendix describes the work done by the author of this dissertation as part of the development of OpenMarkov.

## A.1. Reorganization in Maven projects

The source code of OpenMarkov was included in a single project. This had a number of disadvantages. For example, there was no clear separation between different aspects of the program, such as the business logic and the GUI. Also, due to the collaborative nature of the development of OpenMarkov, where students extend its functionality, having all the code in a single project made it vulnerable to code check-ins that would compromise its stability.

For this reason, we decided to reorganize the source code of to a set of OpenMarkov into a set of projects linked by dependencies, each project containing a different aspect of the software. These projects are configured with Apache Maven, which is a project configuration and build automation tool. This distribution forced us to separate different aspects of the code and therefore create modules minimizing dependencies amongst them. Thanks to the new distribution we were also able to set different permissions for different parts of the code, thus ensuring a higher stability.

Besides, we combined this new distribution with the use of Java annotations to create extensibility points in OpenMarkov. Through these extensibility points, the functionality of OpenMarkov can be extended without the modification of its source code, only by placing jar files inside its classpath. Using this mechanism, users can add new inference algorithms, elimination heuristics, potentials, learning algorithms, metrics for search-and-score learning algorithms and even entire modules accessible through its GUI.

## A.2. Build and test automatization

The division of the code into a set of different projects brought the disadvantage of having to deal with a high number of them. Maintenance of the code base became

increasingly complicated. Tasks such as deploying a new version of the software became complex and time consuming. Making sure changes in some projects did not affect others became a cumbersome task. We automated build generation and periodical testing of the code using a **Jenkins** server, so that every time a change is submitted to the code repository, the tests of the modified projects are run. This allows us to detect errors in the code immediately.

### **A.3. Interactive learning of Bayesian networks**

The development of a framework for the interactive learning of Bayesian networks has been an important part of the work of the author of this thesis and has led to a publication in a conference (Bermejo et al., 2012) and the submission of another paper to the *Journal of Statistical Software*. It was left outside of this dissertation because it did not fit together with the rest of the doctoral research.

### **A.4. The tuning model**

The tuning model described in 5.2 was implemented as a new type of ICI potential in OpenMarkov.

### **A.5. Object-oriented probabilistic networks**

Object-oriented probabilistic networks (OOPNs) were conceived as a generalization of the object-oriented Bayesian networks proposed by Koller and Pfeffer (1997); Bangsø and Willemin (2000) to PGMs with decision and utility nodes, such as IDs. I added functionality in the GUI to define classes, to create instances, and to link instances to other instances and to nodes. The evaluation of OOPNs was done converting it into a plain PGM as a preprocessing step.

### **A.6. Likelihood weighting**

OpenMarkov had no approximate inference algorithms implemented, and the tuning network built for the Opti-FOX project had too big potentials to be evaluated with exact inference. We implemented *likelihood weighting* (Fung and Chang, 1990; Shachter and Peot, 1990) because it was a simple yet powerful algorithm that met our requirements at the time.

## A.7. Cost-effectiveness analysis with MPADs

The implementation of cost-effectiveness analysis with MPADs was already underway when the author of this thesis took over the task: MPADs could be created and edited using OpenMarkov's GUI; the preprocessing of MPADs (turning it into a multi-criteria influence diagram by expanding the network, adding an artificial node of decision criteria as a conditioning variable and applying the discounts) to be evaluated with an inference algorithm for IDs was already implemented, but was buggy.

I corrected a few bugs from the original implementation, implemented the half-cycle correction and the probabilistic sensitivity analysis almost from scratch, including the visualization of its results—the cost-effectiveness analysis curve and the graph of the expected value of information. In order to be able to evaluate certain models, I also implemented the inference with numerical variables in certain scenarios. Finally, I implemented a set of new potentials: Weibull hazard potential, the Exponential hazard potential, the linear combination, and the delta potential.

## A.8. Inference on DTs

Implementation of the rollback algorithm to calculate expected utility and optimal strategy in decision trees, as described in Raiffa (1968).

## A.9. Conversion of an ID into an equivalent DT

This functionality was already present in Elvira and it is useful as an introduction to IDs for people who are used to work with DT. It is also useful to compare how DTs generated from IDs compare to those generated with DANs when there is structural asymmetry involved. This conversion is only tractable for small IDs due to the exponential growth of the size of the equivalent DT.

## A.10. Conversion of a DAN into an equivalent DT

As a means to define the syntax of DANs, an algorithm to convert DANs to equivalent decision trees was proposed in Díez et al. (2012). We implemented it in OpenMarkov also as a means to have an inefficient yet simple way to be able to evaluate DANs.

## **A.11. Evaluation of DANs: recursive decomposition**

The evaluation of DANs by converting them into an equivalent decision tree and then using the rollback algorithm to evaluate it was very inefficient. Therefore, we decided to devise and implement more efficient algorithms to evaluate DANs. This algorithm evaluates DANs by creating a tree of symmetric DANs, evaluating those symmetric DANs with variable elimination.

# Bibliography

- Abrams, K. R., Myles, J. P., and Spiegelhalter, D. J. (2004). *Bayesian Approaches to Clinical Trials and Health-Care Evaluation*. John Wiley & Sons.
- Ahlmann-Ohlsen, K. S., Jensen, F. V., Nielsen, T. D., Pedersen, O., and Vomlelová, M. (2009). A comparison of two approaches for solving unconstrained influence diagrams. *International Journal of Approximate Reasoning*, 50:153–173.
- Arias, E. (2007). United states life tables, 2003. *National Vital Statistics Report*, 54(14):1–40.
- Arias, M. and Díez, F. J. (2015). Cost-effectiveness analysis with influence diagrams. *Methods of Information in Medicine*.
- Åström, K. J. (1965). Optimal control of Markov decision processes with incomplete state estimation. *Journal of Mathematical Analysis and Applications*, 10:174–205.
- Baggerly, K. A. and Coombes, K. R. (2009). Deriving chemosensitivity from cell lines: Forensic bioinformatics and reproducible research in high-throughput biology. *The Annals of Applied Statistics*, 3(4):1309–1334.
- Baio, G. (2012). *Bayesian Methods in Health Economics*. Chapman and Hall/CRC.
- Baio, G., Pammolli, F., Baldo, V., and Trivello, R. (2006). Object oriented influence diagram for cost-effectiveness analysis of influenza vaccination in the Italian elderly population. *Expert Review of Pharmacoeconomics & Outcomes Research*, 6:293–301.
- Bangsø, O. and Wuillemin, P. H. (2000). Top-down construction and repetitive structures representation in Bayesian networks. In *Proceedings of the Thirteenth International Florida Artificial Intelligence Research Society Conference (FLAIRS-2000)*, pages 282–286, Orlando, FL.
- Battmer, R., Borel, S., Brendell, M., Britz, A., Büchner, A., Cooper, H., Fielden, C., Gazibegovic, D., Goetze, R., Govaerts, P., Lenarz, T., Mosnier, I., Muff, J., Nunn, T., Vaerenberg, B., and Vana, Z. (2014). Assessment of “Fitting to Outcomes eXpert” FOX with new cochlear implant users in a multicentric study. *Cochlear Implants International*.
- Beck, J. R. and Pauker, S. G. (1983). The Markov process in medical prognosis. *Medical Decision Making*, 3:419–458.
- Becker, A. and Geiger, D. (1994). Approximation algorithms for the loop cutset problem. In de Mantaras, R. L. and Poole, D., editors, *Proceedings of the Tenth*

- Conference on Uncertainty in Artificial Intelligence (UAI'94)*, pages 60–68, San Francisco, CA. Morgan Kaufmann.
- Bermejo, I., Oliva, J., Díez, F. J., and Arias, M. (2012). Interactive learning of Bayesian networks with OpenMarkov. In Cano et al. (2012), pages 27–34.
- Bielza, C., Gómez, M., and Shenoy, P. P. (2011). A review of representation issues and modelling challenges with influence diagrams. *Omega*, 39:227–241.
- Bielza, C. and Shenoy, P. P. (1999). A comparison of graphical techniques for asymmetric decision problems. *Management Science*, 45:1552–1569.
- Bonissone, P., Henrion, M., Kanal, L. N., and Lemmer, J. F., editors (1990). *Uncertainty in Artificial Intelligence 6 (UAI'90)*, Amsterdam, The Netherlands. Elsevier Science Publishers.
- Botros, A. (2010). *The application of machine intelligence to cochlear implant fitting and the analysis of the auditory nerve response*. PhD thesis, University of New South Wales.
- Botros, A. and Psarros, C. (2010). Neural response telemetry reconsidered: I. the relevance of ECAP threshold profiles and scaled profiles to cochlear implant fitting. *Ear and Hearing*, 31:367–379.
- Botros, A., van Dijk, B., and Killian, M. (2006). AutoNRT™: An automated system that measures ECAP thresholds with the Nucleus® Freedom™ cochlear implant via machine intelligence. *Artificial Intelligence in Medicine*, 40:15–28.
- Bourgeois-Republique, C., Frachet, B., and Collet, P. (2005). Using an interactive evolutionary algorithm to help fitting a cochlear implant. In *Proceedings of the 7th Annual Workshop on Genetic and Evolutionary Computation, GECCO '05*, pages 133–139, New York, NY, USA. ACM.
- Boutilier, C. (2005). The influence of influence diagrams on artificial intelligence. *Decision Analysis*, 2:229–231.
- Boutilier, C., Dearden, R., and Goldszmidt, M. (1995). Exploiting structure in policy construction. In Mellish, C., editor, *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI'95)*, pages 1104–1111, Chambéry, France. Morgan Kaufmann.
- Briggs, A., Claxton, K., and Sculpher, M. (2006). *Decision Modelling for Health Economic Evaluation*. Oxford University Press, New York.
- Briggs, A., Sculpher, M., Dawson, J., Fitzpatrick, R., Murray, D., and Malchau, H. (2004). The use of probabilistic decision models in technology assessment: The case of total hip replacement. *Applied Health Economics and Health Policy*, 3(2):78–79.
- Briggs, A. H. (1999). A Bayesian approach to stochastic cost-effectiveness analysis. *Health Economics*, 8:257–261.



- Brown, C. J., Hughes, M. L., Luk, B., Abbas, P., Wolaver, A., and Gervais, J. (2000). The relationship between EAP and EABR thresholds and levels used to program the Nucleus 24 speech processor: data from adults. *Ear and Hearing*, 21:151–163.
- Buechner, A., Vaerenberg, B., Gazibegovic, D., Brendel, M., De Ceulaer, G., Goverts, P., and Lenarz, T. (2014). Evaluation of the “Fitting to Outcomes eXpert” FOX with established cochlear implant users. *Cochlear Implants International*.
- Cano, A., Gómez, M., and Nielsen, T. D., editors (2012). *Proceedings of the Sixth European Workshop on Probabilistic Graphical Models (PGM’12)*, Granada, Spain.
- Carles, M., Vilapriyo, E., Cots, F., Gregori, A., Pla, R., Román, R., Sala, M., Macià, F., Castells, X., and Rue, M. (2011). Cost-effectiveness of early detection of breast cancer in Catalonia (Spain). *BMC Cancer*, 11(1).
- Caro, J. (2005). Pharmacoeconomic analyses using discrete event simulation. *Pharmacoeconomics*, 23(4):323–332.
- Caro, J. J., Getsios, D., and Möller, J. (2007). Regarding probabilistic analysis and computationally expensive models: Necessary and required? *Value in Health*, 10(4):317–318.
- Castells, X., Sala, M., Ascunce, N., Salas, D., Zubizarreta, R., and Casamitjana, M. (2007). Descripción del cribado del cáncer en España. Technical report, Ministerio de Sanidad y Consumo. Agència d’Avaluació de Tecnologia i Recerca Mèdiques de Catalunya.
- Chancellor, J. V., Hill, A. M., Sabin, C. A., Simpson, K. N., and Youle, M. (1997). Modelling the cost effectiveness of lamivudine/zidovudine combination therapy in HIV infection. *Pharmacoeconomics*, 12(1):54–66.
- Chang, C. H., Anderson, T., and Loizou, P. C. (2001). A neural network model for optimizing vowel recognition by cochlear implant listeners. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 9:42 – 48.
- Claxton, K., Martin, S., Soares, M., Rice, N., Spackman, E., Hinde, S., Devlin, N., Smith, P., and Sculpher, M. (2015). Methods for the estimation of the national institute for health and care excellence cost-effectiveness threshold. *Health Technology Assessment*, 19.
- Claxton, K., Sculpher, M., McCabe, C., et al. (2005). Probabilistic sensitivity analysis for NICE technology assessment: not an optional extra. *Health Economics*, 14:339–347.
- Clive, A. O., Kahan, B. C., Hooper, C. E., Bhatnagar, R., Morley, A. J., Zahan-Evans, N., Bintcliffe, O. J., Boshuizen, R. C., Fysh, E. T. H., Tobin, C. L., Medford, A. R. L., Harvey, J. E., van den Heuvel, M. M., Lee, Y. C. G., and Maskell, N. A. (2014). Predicting survival in malignant pleural effusion: development and validation of the lent prognostic score. *Thorax*, 69(12):1098–1104.

- CMS (2014). CY 2015 Physician Fee Schedule. Technical report, Centers for Medicare and Medicaid Services.
- Collet, P., Legrand, P., Bourgeois-République, C., Péan, V., and Frachet, B. (2009). Using interactive evolutionary algorithms to help fit cochlear implants. In Siarry, P., editor, *Optimization in Signal and Image Processing*, chapter 13. ISTE, London, UK.
- Covaliu, Z. and Oliver, R. M. (1995). Representation and solution of decision problems using sequential decision diagrams. *Management Science*, 41:1860–1881.
- Cowell, R. G., Dawid, A. P., Lauritzen, S. L., and Spiegelhalter, D. J. (1999). *Probabilistic Networks and Expert Systems*. Springer-Verlag, New York.
- Demirer, R. and Shenoy, P. P. (2006). Sequential valuation networks for asymmetric decision problems. *European Journal of Operational Research*, 169:286–309.
- Díez, F. J. and Druzdzal, M. J. (2006). Canonical probabilistic models for knowledge engineering. Technical Report CISIAD-06-01, UNED, Madrid, Spain.
- Díez, F. J., Luque, M., and König, C. (2012). Decision analysis networks. In Cano et al. (2012), pages 83–90.
- Díez, F. J. and van Gerven, M. A. J. (2011). Dynamic LIMIDs. In Sucar, L. E., Hoey, J., and Morales, E., editors, *Decision Theory Models for Applications in Artificial Intelligence: Concepts and Solutions*, pages 164–189. IGI Global, Hershey, PA.
- Dresler, C., Olak, J., Herndon, J., Richards, W., Scalzetti, E., Fleishman, S., Kernstine, K., Demmy, T., Jablons, D., Kohman, L., Daniel, T., Haasler, G., Sugarbaker, D., Cooperative Groups Cancer and Leukemia Group B, Eastern Cooperative Oncology Group, North Central Cooperative Oncology Group, and Radiation Therapy Oncology Group (2005). Phase III intergroup study of talc poudrage vs talc slurry sclerosis for malignant pleural effusion. *Chest*, 127(3):909–915.
- Drummond, M. F., Sculpher, M. J., Torrance, G. W., O’Brien, B. J., and Stoddart, G. L. (2005). *Methods for the Economic Evaluation of Health Care Programmes*. Oxford University Press, third edition.
- Elvira Consortium, T. (2002). Elvira: An environment for creating and using probabilistic graphical models. In Gámez, J. A. and Salmerón, A., editors, *Proceedings of the First European Workshop on Probabilistic Graphical Models (PGM’02)*, pages 1–11.
- Ernster, V., Barclay, J., Kerlikowske, K., Wilkie, H., and Ballard-Barbash, R. (2000). Mortality among women with ductal carcinoma in situ of the breast in the population-based surveillance, epidemiology and end results program. *Archives of Internal Medicine*, 160(7):953–958.
- European Union (2003). L327/34 council recommendation of 2 december 2003 on cancer screening (2003/878/ec). *Official Journal of the European Union*.

- Ferlay, J., Soerjomataram, I., Ervik, M., Dikshit, R., Eser, S., Mathers, C., Rebelo, M., Parkin, D., Forman, D., and Bray, F. (2012). Globocan 2012 v1.0, cancer incidence and mortality worldwide: IARC CancerBase no. 11. <http://globocan.iarc.fr>. [Online; accessed 2015-03-16].
- Fung, R. and Chang, K. C. (1990). Weighing and integrating evidence for stochastic simulation in Bayesian networks. In Bonissone et al. (1990), pages 209–219.
- Goeree, R., O'Brien, B., Hunt, R., et al. (1999). Economic evaluation of long term management strategies for erosive oesophagitis. *Pharmacoeconomics*, 16:679–697.
- Gompertz, B. (1825). On the nature of the function expressive of the law of human mortality, and on a new mode of determining the value of life contingencies. *Philosophical Transactions of the Royal Society of London*, 115:513–583.
- Govaerts, P. J., Daemers, K., Yperman, M., De Beukelaer, C., De Saegher, G., and De Ceulaer, G. (2006). Auditory speech sounds evaluation (AŞE®): a new test to assess detection, discrimination and identification in hearing impairment. *Cochlear Implants International*, 7:92–106.
- Govaerts, P. J., Vaerenberg, B., De Ceulaer, G., Daemers, K., De Beukelaer, C., and Schauwers, K. (2010). Development of a software tool using deterministic logic for the optimization of cochlear implant processor programming. *Audiology and Neurotology*, 31:908–918.
- Gray, A. M., Clarke, P. M., Wolstenholme, J., and Wordsworth, S. (2011). *Applied Methods of Cost-effectiveness Analysis in Healthcare*. Oxford University Press, New York.
- Griffin, S., Claxton, K., Hawkins, N., and Sculpher, M. (2006). Probabilistic analysis and computationally expensive models: Necessary and required? *Value in Health*, 9(4):244–252.
- Grosse, S. D. (2008). Assessing cost-effectiveness in healthcare: history of the \$50,000 per qaly threshold. *Expert Review of Pharmacoeconomics & Outcomes Research*, 8(2):165–178.
- Gøtzsche, P. C. and Nielsen, M. (2009). Screening for breast cancer with mammography. *Cochrane Database Syst Rev*, 4(1).
- Hawkins, N., Sculpher, M., and Epstein, D. (2005). Cost-effectiveness analysis of treatments for chronic disease: Using R to incorporate time dependency of treatment response. *Medical Decision Making*, 25:511–519.
- Hazen, G. B. (2004). Dynamic influence diagrams: Applications to medical decision modeling. In Brandeau, M., Sainfort, F., and Pierskalla, W., editors, *Operations Research and Health Care*, volume 70 of *International Series in Operations Research & Management Science*, pages 613–638. Springer.
- Heckerman, D., Breese, J. S., and Rommelse, K. (1995). Decision-theoretic troubleshooting. *Communications of the ACM*, 38:49–57.

- Herndon, T., Ash, M., and Pollin, R. (2013). Does high public debt consistently stifle economic growth? A critique of Reinhart and Rogoff. *Cambridge Journal of Economics*.
- Hollenberg, P. (1984). Markov cycle trees: A new representation for complex Markov processes. In *Sixth Annual Meeting of the Society for Medical Decision Making*, volume 4, page 529. Medical Decision Making.
- Howard, R. A. (1984). The used car buyer. In Howard, R. A. and Matheson, J. E., editors, *Readings on the Principles and Applications of Decision Analysis*, pages 689–718. Strategic Decisions Group, Menlo Park, CA.
- Howard, R. A. and Matheson, J. E. (1984). Influence diagrams. In Howard, R. A. and Matheson, J. E., editors, *Readings on the Principles and Applications of Decision Analysis*, pages 719–762. Strategic Decisions Group, Menlo Park, CA.
- Howlander, N., Noone, A., Krapcho, M., Garshell, J., Miller, D., Altekruse, S., Kosary, C., Yu, M., Ruhl, J., Tatalovich, Z., Mariotto, A., Lewis, D., Chen, H., Feuer, E., and Cronin, K. (2011). Seer cancer statistics review, 1975-2011. Technical report, National Cancer Institute.
- Hughes, M. L., Brown, C. J., Luk, B., Abbas, P., Wolaver, A., and Gervais, J. (2000). Comparison of EAP thresholds with MAP levels in the Nucleus 24 cochlear implant: Data from children. *Ear and Hearing*, 21:164–174.
- Jensen, F. V. and Nielsen, T. D. (2007). *Bayesian Networks and Decision Graphs*. Springer-Verlag, New York, second edition.
- Jensen, F. V., Nielsen, T. D., and Shenoy, P. P. (2006). Sequential influence diagrams: A unified asymmetry framework. *International Journal of Approximate Reasoning*, 42:101–118.
- Jensen, F. V. and Vomlelová, M. (2002). Unconstrained influence diagrams. In Darwiche, A. and Friedman, N., editors, *Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence (UAI'02)*, pages 234–241, San Francisco, CA. Morgan Kaufmann.
- Karnon, J. (2003). Alternative decision modelling techniques for the evaluation of health care technologies: Markov processes versus discrete event simulation. *Health Economics*, 12(10):837–848.
- Karnon, J., Stahl, J., Brennan, A., Caro, J., Mar, J., and Möller, J. (2012). Modeling using discrete event simulation: A report of the ISPOR-SMDM modeling good research practices task force-4. *Value in Health*, 15:821–827.
- Kennedy, L. and Sahn, S. (1994). Talc pleurodesis for the treatment of pneumothorax and pleural effusion. *Chest*, 106(4):1215–1222.
- Kerlikowske, K., Molinaro, A., Cha, I., Ljung, B.-M., Ernster, V. L., Stewart, K., Chew, K., H. Moore, D., and Waldman, F. (2003). Characteristics associated with recurrence among women with ductal carcinoma in situ treated by lumpectomy. *Journal of the National Cancer Institute*, 95(22):1692–1702.

- Koller, D. and Pfeffer, A. (1997). Object-oriented Bayesian networks. In Geiger, D. and Shenoy, P., editors, *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence (UAI'97)*, pages 302–313, San Francisco, CA. Morgan Kaufmann.
- Lacave, C. and Díez, F. J. (2002). A review of explanation methods for Bayesian networks. *Knowledge Engineering Review*, 17:107–127.
- Lacave, C., Luque, M., and Díez, F. J. (2007). Explanation of Bayesian networks and influence diagrams in Elvira. *IEEE Transactions on Systems, Man and Cybernetics—Part B: Cybernetics*, 37:952–965.
- Lacave, C., Oniśko, A., and Díez, F. J. (2006). Use of Elvira’s explanation facilities for debugging probabilistic expert systems. *Knowledge-Based Systems*, 19:730–738.
- Lafon, J. C. (1964). *Le test phonétique et la mesure de l’audition*. Ed. Centrex, Eindhoven.
- Lauritzen, S. L. and Nilsson, D. (2001). Representing and solving decision problems with limited information. *Management Science*, 47:1235–1251.
- Legrand, P., Bourgeois-Republique, C., Péan, V., Harboun-Cohen, E., Levy-Vehel, J., Frachet, B., Lutton, E., and Collet, P. (2007). Interactive evolution for cochlear implants fitting. *Genetic Programming and Evolvable Machines*, 8:319–354.
- León, D. (2011). A probabilistic graphical model for total knee arthroplasty. Master’s thesis, Dept. Artificial Intelligence, UNED, Madrid, Spain.
- Lidgren, M., Wilking, N., Jönsson, B., and Rehnberg, C. (2007). Health related quality of life in different states of breast cancer. *Quality of Life Research*, 16(6):1073–1081.
- Liem, Y. S., Bosch, J. L., and Hunink, M. G. M. (2008). Preference-based quality of life of patients on renal replacement therapy: A systematic review and meta-analysis. *Value in Health*, 11(4):733–741.
- López Bastida, J., Bellas Beceiro, B., and Garcia Perez, L. (2009). Analisis coste-efectividad del cribado del cancer de mama mediante mamografía en diferentes grupos de edad (40 a 49, 50 a 69 y 70 a 75). Technical report, Servicio Canario de Salud.
- Luque, M. and Díez, F. J. (2004). Variable elimination for influence diagrams with super-value nodes. In Lucas, P., editor, *Proceedings of the Second European Workshop on Probabilistic Graphical Models (PGM'04)*, pages 145–152, Leiden, The Netherlands.
- Luque, M. and Díez, F. J. (2010). Variable elimination for influence diagrams with super-value nodes. *International Journal of Approximate Reasoning*, 51:615 – 631.
- Luque, M., Díez, F. J., and Disdier, C. (2009). A decision support system for the mediastinal staging of non-small cell lung cancer. In *31st Annual Meeting of the Society for Medical Decision Making*, Los Angeles, CA.

- Luque, M., Nielsen, T. D., and Jensen, F. V. (2008). An anytime algorithm for evaluating unconstrained influence diagrams. In Jensen, F. V. and Kjærulff, U., editors, *Proceedings of the Fourth European Workshop on Probabilistic Graphical Models (PGM'08)*, pages 177–184, Hirtshals, Denmark.
- Murphy, K. (2002). *Dynamic Bayesian Networks: Representation, Inference and Learning*. PhD thesis, Computer Science Division, University of California, Berkeley.
- Nafees, B., Stafford, M., Gavriel, S., Bhalla, S., and Watkins, J. (2008). Health state utilities for non small cell lung cancer. *Health and quality of life outcomes*, 6.
- Neragi-Miandoab, S. (2006). Malignant pleural effusion, current and evolving approaches for its diagnosis and management. *Lung cancer*, 54(1):1–9.
- NICE (2012). *Single technology appraisal (STA) - Specification for manufacturer/sponsor submission of evidence*. National Institute For Health and Care Excellence.
- Nielsen, S. H., Nielsen, T. D., and Jensen, F. V. (2007). Multi-currency influence diagrams. In Salmerón, A. and Gámez, J. A., editors, *Advances in Probabilistic Graphical Models*, pages 275–294. Springer, Berlin, Germany.
- Nielsen, T. D. and Jensen, F. (2000). Representing and solving asymmetric bayesian decision problems. In Boutilier, C. and Goldszmidt, M., editors, *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence (UAI'00)*, pages 416–425, San Francisco, CA. Morgan Kauffmann.
- Nielsen, T. D. and Jensen, F. V. (1999). Welldefined decision scenarios. In Laskey, K. and Prade, H., editors, *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence (UAI'99)*, pages 502–511, San Francisco, CA. Morgan Kauffmann.
- NIH (2013). NIH Publication No. 11-4798.
- Olden, A. M. and Holloway, R. (2009). Treatment of malignant pleural effusion: PleuRx® catheter or talc pleurodesis? a cost-effectiveness analysis. *Journal of palliative medicine*, 13(1):59–65.
- Olmsted, S. M. (1983). *On Representing and Solving Decision Problems*. PhD thesis, Dept. Engineering-Economic Systems, Stanford University, CA.
- Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, CA.
- Pedley, K., Psarros, C., Gardner-Berry, K., Parker, A., Purdy, S. C., Dawson, P., and Plant, K. (2007). Evaluation of NRT and behavioral measures for MAPping elderly cochlear implant users. *International Journal of Audiology*, 46(5):254–262.
- Pidd, M. (2004). *Computer Simulation in Management Science (5th ed)*. John Wiley & Sons, New York.

- Pitman, R., Fisman, D., Zaric, G. S., Postma, M., Kretzschmar, M. Edmunds, J., and Brisson, M. (2012). Dynamic transmission modeling: A report of the ISPOR-SMDM modeling good research practices task force working group-5. *Medical Decision Making*, 32:712–723.
- Prada, L. (2014). Optimization of policy trees for influence diagrams. Master’s thesis, Dept. Artificial Intelligence, UNED, Madrid, Spain.
- Puri, V., Pyrdeck, T. L., Crabtree, T. D., Kreisel, D., Krupnick, A. S., Colditz, G. A., Patterson, G. A., and Meyers, B. F. (2012). Treatment of malignant pleural effusion: A cost-effectiveness analysis. *The Annals of Thoracic Surgery*, 94(2):374–380.
- Putnam, J. B., Walsh, G. L., Swisher, S. G., Roth, J. A., Suell, D. M., Vaporciyan, A. A., Smythe, W. R., Merriman, K. W., and DeFord, L. L. (2000). Outpatient management of malignant pleural effusion by a chronic indwelling pleural catheter. *The Annals of Thoracic Surgery*, 69(2):369–375.
- Quinlan, J. R. (1993). *C4.5: Programs for Machine Learning*. Morgan Kaufmann, San Diego.
- Raiffa, H. (1968). *Decision Analysis. Introductory Lectures on Choices under Uncertainty*. Addison-Wesley, Reading, MA.
- Raiffa, H. and Schlaifer, R. (1961). *Applied Statistical Decision Theory*. John Wiley & Sons, Cambridge, MA.
- Roberts, M., Russell, L. B., Paltiel, A. D., Chambers, M., McEwan, P., and Krahn, M. (2012). Conceptualizing a model: A report of the ISPOR-SMDM modeling good research practices task force-2. *Medical Decision Making*, 32(5):678–689.
- Ryan, M., Griffin, S., Chitah, B., Walker, A. S., Mulenga, V., Kalolo, D., Hawkins, N., Merry, C., Barry, M. G., Chintu, C., Sculpher, M. J., and Gibb, D. M. (2008). The cost-effectiveness of cotrimoxazole prophylaxis in HIV-infected children in Zambia. *AIDS*, 22(6):749–757.
- Schousboe, J. T., Kerlikowske, K., Loh, A., and Cummings, S. R. (2011). Personalizing mammography by breast density and other risk factors for breast cancer: Analysis of health benefits and cost-effectiveness. *Annals of Internal Medicine*, 155(1):10–20.
- Shachter, R. and Peot, M. (1990). Simulation approaches to general probabilistic inference on belief networks. In Bonissone et al. (1990), pages 221–231.
- Shachter, R. D. (1986). Evaluating influence diagrams. *Operations Research*, 34:871–882.
- Shapiro, W. H. and Bradham, T. S. (2012). Cochlear implant programming. *Otolaryngologic Clinics of North America*, 45(1):111 – 127.
- Shaw, P. and Agarwal, R. (2004). Pleurodesis for malignant pleural effusions. *Cochrane Database of Systematic Reviews*, (1).

- Shenoy, P. P. (1992). Valuation based systems for Bayesian decision analysis. *Operations Research*, 40:463–484.
- Shenoy, P. P. (2000). Valuation network representation and solution of asymmetric decision problems. *European Journal of Operational Research*, 121:579–608.
- Siebert, U., Alagoz, O., Bayoumi, A. M., Jahn, B., Owens, D. K., Cohen, D. J., and Kuntz, K. M. (2012). State-Transition Modeling: A report of the ISPOR-SMDM Modeling Good Research Practices Task Force-3. *Medical Decision Making*, 32(5):690–700.
- Simpson, K. N., Strassburger, A., Jones, W. J., Dietz, B., and Rajagopalan, R. (2009). Comparison of Markov Model and Discrete-Event Simulation Techniques for HIV. *Pharmacoeconomics*, 27(2):159–165.
- Skinner, M. W., Binzer, S. B., Potts, L., Holden, L. K., and Aaron, R. J. (2002). Hearing rehabilitation for individuals with severe and profound hearing impairment: hearing aids, cochlear implants, and counseling. In Valente, M., editor, *Strategies for Selecting and Verifying Hearing Aid Fittings*, pages 311–344. Thieme, New York.
- Smith, J. E., Holtzman, S., and Matheson, J. E. (1993). Structuring conditional relationships in influence diagrams. *Operations Research*, 41:280–297.
- Smoorenburg, G. F., Willeboer, C., and van Dijk, J. E. (2002). Speech perception in Nucleus CI24M cochlear implant users with processor settings based on electrically evoked compound action potential thresholds. *Audiology and Neurotology*, 7(6):335–347.
- Sonnenberg, F. A. and Beck, J. R. (1993). Markov models in medical decision making: A practical guide. *Medical Decision Making*, 13:322–339.
- Szlávik, Z., Vaerenberg, B., Kowalczyk, W., and P., G. (2011). Opti-Fox: Towards the automatic tuning of cochlear implants. In *20th Machine Learning conference of Belgium and The Netherlands.*, pages 79–80.
- Tatman, J. A. and Shachter, R. D. (1990). Dynamic programming and influence diagrams. *IEEE Transactions on Systems, Man and Cybernetics*, 20:365–379.
- Tice, J. A., Cummings, S. R., Smith-Bindman, R., Ichikawa, L., Barlow, W. E., and Kerlikowske, K. (2008). Using clinical factors and mammographic breast density to estimate breast cancer risk: Development and validation of a new predictive model. *Annals of Internal Medicine*, 148(5):337–347.
- Tosh, J. and Wailoo, A. (2008). *Review of software for decision modelling*. NICE Decision Support Unit.
- Tosteson, A. o. (2008). Cost-effectiveness of digital mammography breast cancer screening: Results from ACRIN DMIST. *Annals of Internal Medicine*, 148(1):1–10.



- Tremblay, A., Mason, C., and Michaud, G. (2007). Use of tunnelled catheters for malignant pleural effusions in patients fit for pleurodesis. *European Respiratory Journal*, 30(4):759–762.
- Tremblay, A. and Michaud, G. (2006). Single-center experience with 250 tunnelled pleural catheter insertions for malignant pleural effusion. *Chest*, 129(2):362–368.
- U.S. Preventive Services Task Force (2009). Screening for breast cancer: U.s. preventive services task force recommendation statement. *Annals of Internal Medicine*, 151(10):716–726.
- Vaerenberg, B., De Ceulaer, G., Szlávik, Z., Mancini, P., Buechner, A., and Govaerts, P. J. (2014a). Setting and reaching targets with computer-assisted cochlear implant fitting. *The Scientific World Journal*, 2014.
- Vaerenberg, B., Govaerts, P. J., Stainsby, T., Nopp, P., Gault, A., and Gnansia, D. (2014b). A uniform graphical representation of intensity coding in current-generation cochlear implant systems. *Ear and Hearing*, 35:533–543.
- Vaerenberg, B., Smits, C., De Ceulaer, G., Zir, E., Harman, S., and Jaspers, N. (2014c). Cochlear implant programming: A global survey on the state of the art. *The Scientific World Journal*, 2014.
- van der Beek, F. B., Briaire, J. J., and Frijns, J. H. M. (2015). Population-based prediction of fitting levels for individual cochlear implant recipients. *Audiology and Neurotology*, 20(1):1–16.
- Weibull, W. (1951). A statistical distribution function of wide applicability. *ASME Journal of Applied Mechanics, Transactions of the American Society of Mechanical Engineers*, 18:293–297.
- Weinstein, M. C., O’Brien, B., Hornberger, J., Jackson, J., Johannesson, M., McCabe, C., and Luce, B. R. (2003). Principles of good practice for decision analytic modeling in health-care evaluation: Report of the ISPOR task force on good research practices-modeling studies. *Value in Health*, 6(1):9–17.
- Welton, N. J., Sutton, A. J., Cooper, N., Abrams, K. R., and Ades, A. E. (2012). *Evidence Synthesis for Decision Making in Healthcare*. Wiley-Blackwell.
- Yabroff, K. R., Lamont, E. B., Mariotto, A., Warren, J. L., Topor, M., Meekins, A., and Brown, M. L. (2008). Cost of care for elderly cancer patients in the united states. *Journal of the National Cancer Institute*, 100(9):630–641.