A decision support-system for the mediastinal staging of non-small cell lung cancer

Manuel Luque and **Francisco J. Díez**
Dept. Artificial Intelligence, UNED
Juan del Rosal, 16
28040 Madrid, Spain
{mluque,fjdiez}@dia.uned.es

Carlos Disdier
Pneumology Section
University Hospital
47005 Valladolid, Spain
cdisdier@separ.es

Abstract

Lung cancer is a very frequent tumor in the developed world and the leading cause of cancer death, with non-small cell lung cancer being the most prevalent type and with most difficult prognosis. In this paper we present a decision support system built for finding the optimal selection of tests and therapy for each patient. The system basically consists of an influence diagram with super value nodes. The parameter λ, which in cost-effectiveness analyses represents the amount of money that the decision maker is willing to pay to obtain a unit of effectiveness, has been included in the influence diagram, and has allowed us to find a trade-off between cost and effectiveness. Finally, given the uncertainty on the values of the parameters, we have assigned, with the expert's help, a probability distribution to each parameter of the model and have performed a probabilistic sensitivity analysis.

1 INTRODUCTION

Lung cancer is a very frequent tumor in the developed world and the leading cause of cancer death. Lung cancer can be classified into two major types: small-cell lung cancer (SCLC) and non-small cell lung cancer (NSCLC). The first one appears in 20% of cases, is usually inoperable and only treatable with chemotherapy or chemo-radiotherapy. In contrast, when limited to the lung, certain adjacent structures, and lymph nodes proximal to the lung, surgery resection remains the optimal treatment for NSCLC. However, more than 80% of NSCLC patients can not be treated with surgery because the disease is out of control due to an advanced local extension of the tumor or spreading to other parts of the body (metastasis). A disappointing fact is that a high percentage of patients that may benefit from surgery die of lung cancer. A correct assessment at an early stage of the disease and an accurate selection of patients (staging phase) is very important to apply surgery in good prognosis patients, and, in turn, to avoid dangerous, painful, and unnecessary surgery in bad prognosis patients.

When there are no distant metastases, mediastinal staging, i.e., determining whether malignant mediastinal lymph nodes are present or absent, is the most important prognostic factor in patients with NSCLC and, consequently, determines the therapeutic strategy. Different techniques are available to study the mediastinum. There are non-invasive imaging techniques, such as CT scan and PET, with high sensitivity but low specificity; there are also minimally invasive endoscopic techniques (TBNA, EBUS, EUS)[1], with low risk, high specificity and varying degrees of sensitivity, as well as more invasive surgical techniques, such as mediastinoscopy, which is considered as the gold standard.

The main treatment options for lung cancer include surgery, chemotherapy, radiation therapy, radiochemotherapy, and palliative and supportive care. The applicability of each treatment depends on the stage of the tumor.

Because of this variety of available tests and treatments, each one having pros and cons, there is a vivid debate among specialists about which technologies should be used (Fritscher-Ravens et al., 2003; Schimmer et al., 2006). Nease and Owens (1997) proposed an influence diagram for the mediastinal staging of NSCLC, which provides a strategy for a simplified version of the problem. We propose here a new ID, with important improvements. We also describe how we have searched for a tradeoff between cost and effec-

[1] CT scan stands for computer tomography, PET for position emission tomography, TBNA for transbronchial needle aspiration, EBUS for endobronchial ultrasound, and EUS for endoscopic ultrasound.

tiveness by including in the influence diagram the parameter λ, which represents the amount of money that the decision maker is willing to pay to obtain a unit of effectiveness. Finally, we present the probabilistic sensitivity analysis (PSA) that we have performed given the uncertainty on the values of the parameters.

2 CONSTRUCTION OF MEDIASTINET

In this section, we describe the construction of MEDIASTINET, an influence diagram (ID) for the mediastinal staging of non-small cell lung cancer (NSCLC).

2.1 STRUCTURE OF THE GRAPH

Influence diagrams (Howard and Matheson, 1984) are a framework for representing and solving decision problems. An ID consists of an acyclic directed graph having three kinds of nodes: decision (graphically represented by squares or rectangles), chance (circles or ovals), and utilities (diamonds). Each decision node represents to actions under the direct control of the decision maker. Each chance node represents a random variable. In medical IDs, utility nodes represent medical outcomes and costs (morbidity, mortality, economic cost...). We will use the terms node and variable indifferently.

We next describe how the ID has been built by exploiting expert knowledge.

2.1.1 Identification of variables

Chance variables Given that our objective is the mediastinal staging of NSCLC, we have included a variable representing the value of N factor in the TNM classification[2] (Lloyd and Silvestri, 2001). Even though the N factor takes on four possible values, from N0 to N3, we have modeled it as a binary variable because the cancer is operable for groups N0 and N1, but it is inoperable for N2 and N3. The variable has been named *N2_N3* (see Figure 1).

The laboratory tests that can be performed are represented by the binary variables *CT_scan*, *TBNA*, *PET*, *EBUS*, *EUS,* and *MED* (the result of the mediastinoscopy). We have also created the variable *MED_Sv*, which represents whether the patient has survived mediastinoscopy.

Decision variables The set of possible treatments is represented by the variable *Treatment.* Its states are *thoracotomy*, *radio-chemotherapy*, and *palliative*.[3]

The decisions about whether to perform the different laboratory tests have been represented by the variables with the prefix *Decision_* on the name.[4] These decisions forced us to add a new state *no_result* to the variables *TBNA*, *PET*, *EBUS*, *EUS*, and *MED*, to reflect that when we do not perform a medical test its result is not available.

Ordinary utility nodes The quality-adjusted life expectancy (QALE) (Weinstein and Statson, 1977) of the survivors to the medical tests (except the mediastinoscopy) and the treatment is represented by the node *Survivors_QALE*.

The morbidities due to TBNA, EBUS, EUS, and mediastinoscopy, are depicted by *TBNA_Morbidity*, *EBUS_Morbidity*, *EUS_Morbidity*, and *Med_Morbidity* respectively, and measured in QALYs.

Med_Survival indicates whether the patient has survived to the mediastinoscopy.

The probability of survival to the treatment is represented by *Immediate_Survival*.

Super value nodes The ordinary utility nodes presented above have been combined by using super-value nodes (SVNs), as proposed by Tatman and Shachter (1990). SVNs are either of type sum or product. The type of each SVN has been represented by attaching the corresponding sign of sum or product, as shown in Figures 1 and 2.

Nodes *Survivors_QALE* (QALE of the survivors to the medical tests and the treatments) and *Med_Survival* (probability of survival to the mediastinoscopy), have been combined into the product node *Net_QALE*.

Nodes *TBNA_Morbidity*, *EBUS_Morbidity*, *EUS_Morbidity*, and *Med_ Morbidity* have been combined with *Net_QALE* into the sum node *Total_QALE*.

2.1.2 Arcs of the graph

The influence diagram contains four kinds of arcs:

1. **Arcs into chance nodes.** They represent probabilistic dependencies. In our influence diagram,

[2]The TNM classification is a cancer staging system using three factors (T, N and M) to describe the extent of cancer in a patient's body. N factor describes regional lymph nodes that are involved.

[3]Other possible treatments are irrelevant from a medical point of view in the scenarios considered in this diagram.

[4]We do not include a node *Decision_CT_scan* because CT scan is always performed to a patient.

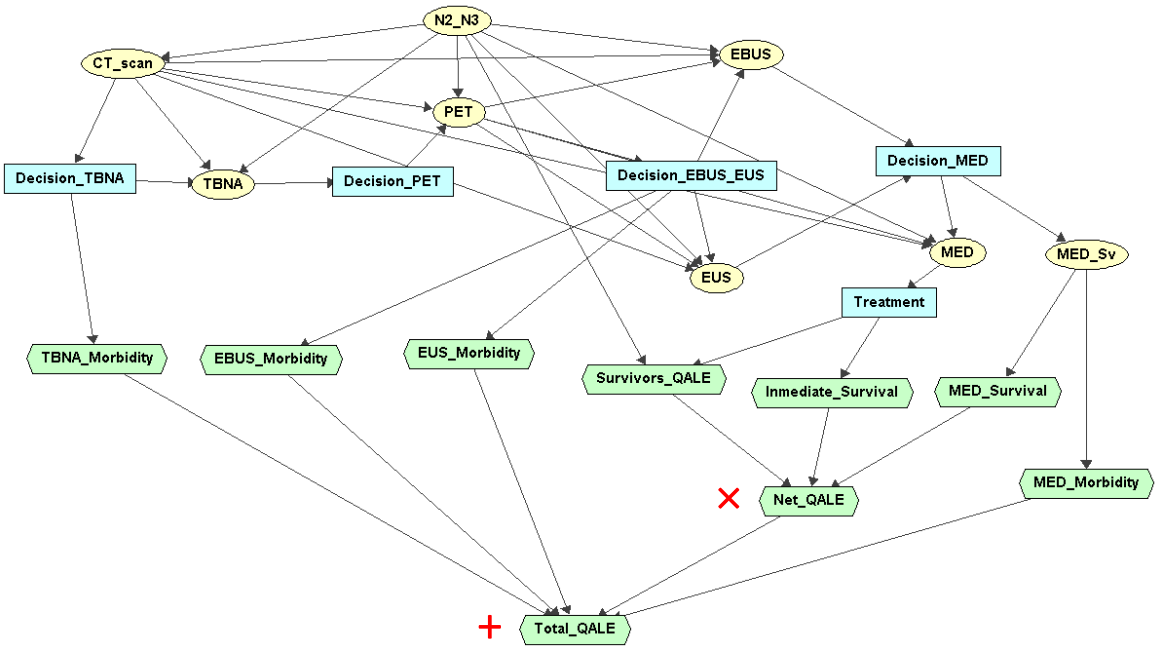

Figure 1: Influence diagram of MEDIASTINET.

an arc from a node representing the decision of a test, such as the arc *Decision_TBNA→TBNA*, indicates that the result (in this case *TBNA*) is only available when we perform the test (*Decision_TBNA=yes*)

2. **Arcs into decision nodes.** They imply informational precedence. Based on the "non-forgetting" assumption (Nielsen and Jensen, 1999), we have not drawn non-forgetting links, to make the influence diagram more clear. For example, the arc *CT_scan→Decision_PET* is not necessary due to the no-forgetting assumption.

3. **Arcs into ordinary utility nodes.** They represent functional dependencies. For example, the arcs into the node *Immediate_Survival* means that the domain of its utility function consists of nodes *N2_N3* and *Treatment*.

4. **Arcs into SVNs.** They indicate the set of utility nodes that are combined into the SVN. For instance, the arcs pointing at the node *Net_QALE* mean that is the combination of *Survivors_QALE*, *MED_Survival* and *Immediate_Survival*.

2.2 PROBABILITIES AND UTILITIES

The quantitative part of the ID consists of a set of probability and utility potentials. For example, for each chance node C we must give a conditional probability potential $p(C|pa(C))$ for each configuration of its parents, $pa(C)$. Then, the table for $p(C|pa(C))$ requires $|dom(C)|\cdot \prod_{X\in pa(C)} |dom(X)|$ numbers, but given the restriction that $\sum_c P(c|pa(C)) = 1$, only some of them are independent.

Given that the parameters of MEDIASTINET are not known with precision we attached a probability distribution to each parameter. We identified the type of distribution of each parameter with the expert's help. For the probabilities (prevalence of the disease, the sensitivities and the specificities of tests) we assigned beta distributions. For the utilities (QALE of the survivors to the treatments) we assigned normal distributions.

In spite of the uncertainty of the parameters, the analysis of the optimal strategies requires to focus on a particular model, called *reference case*, in which all the parameters are assumed to be known with certainty. We have assumed that the reference case of MEDIASTINET takes the mean of each numerical parameter as the value in the reference model.

2.3 COST-EFFECTIVENESS AND NET HEALTH BENEFIT

The version of MEDIASTINET presented above does not include the economic cost of the diagnostic tests and the treatments. However, in medical decision making, costs cannot be ignored. Including the economic cost turns the above problem into a multiobjective prob-

lem with two attributes: the effectiveness, measured in clinical unit, which we want to maximize, and the economic cost, measured in monetary units, which we want to minimize.

One approach to solve the above problem is based on the concept of *net health benefit* (Stinnett and Mullahy, 1998), defined as follows:

$$NHB = E - C/\lambda = E - \lambda^* C, \quad (1)$$

where E is the effectiveness, C is the cost, λ, sometimes called *willingness to pay*, is used here to convert the effectiveness into a monetary scale, and $\lambda^* = 1/\lambda$. The value of λ depends on each decision maker, it is assumed to be positive, but it is usually unknown.

Other possible solution in the framework of IDs would be to use multi-currency IDs (Nielsen et al., 2007). However, this approach would require to specify two parameters, α_1 and α_2, which act as weights of the efectiveness and the economic cost. We have instead preferred to use the approach based on the NHB, besides other reasons (see Section 5), because it only requires one parameter, λ, which has been included explicitly in the ID.

In our model, we identified the effectiveness with the QALE, whose unit is the *quality-adjusted life year* (QALY) (Weinstein and Statson, 1977).

Nevertheless, instead of performing the analysis based on the incremental cost-effectiveness ratios (ICERs) (Gold et al., 1996), which is the standard method, we will apply an equivalent approach: the maximization of the net health benefit, defined in Equation 1. Its integration in MEDIASTINET is as follows (see Figure 2):

- The cost, C, is represented by the sum node *Total_Economic_Cost*, whose parents represent the economic costs of tests and treatments.
- The effectiveness, E, is depicted by *Total_QALE*, explained in Section 2.1.
- The parameter λ^*, the inverse of λ, is represented by *C2E* (cost to effectiveness).
- *Weighted_Economic_Cost* is a product node standing for $\lambda^* C$.
- *Net_Health_Benefit* represents the NHB (Equation 1).

With regards to the utilities, the economic costs have been attached to normal distributions, and parameter λ was characterized by a log-normal distribution.

If we make $\lambda^* = 0$, the evaluation of the ID returns the strategy that maximizes the effectiveness, without taking into account the economic costs. The medical doctor participating in this study was very interested in knowing this strategy, which turns out to be different from the one obtained with the value of $\lambda = 30,000$ €/QALY, used as a reference point for the Spanish public health system (Sacristán et al., 2002).

This justifies why in our ID we have used λ^* as a parameter in Equation 1 instead of λ: because when looking for the maximum-effectiveness strategy (without caring about the costs), it suffices to make $\lambda^* = 0$, In contrast, making $\lambda = +\infty$ would present computational problems.

3 OPTIMAL STRATEGIES FOR THE REFERENCE CASE

3.1 COMPUTATION AND REPRESENTATION OF THE STRATEGIES

The object of decision analysis on a probabilistic decision problem, represented for example in a decision tree or an ID, is twofold: to determine an optimal strategy, and to compute the maximum expected utility (MEU).

We have computed two strategies for MEDIASTINET with two different criteria: the maximization of the effectiveness (disregarding costs) and the maximization of the net health benefit. They have been obtained by solving MEDIASTINET twice: one with $\lambda^* = 0$ (see Eq. 1), and the other one making $\lambda^* = 1/\lambda = 1/30,000$. Changing λ^* in MEDIASTINET only implies setting the utility node *C2E* to the value of λ^*.

The optimal strategy of an ID contains a policy for each decision. Policies are usually presented in the form of a *policy table*, containing a column for each configuration of informational predecessors of the decision. For example, Figure 3 displays optimal policy for *Decision_PET* of MEDIASTINET.

However, given that the size of the policy tables grows exponentially with the number of informational predecessors, we felt the need of presenting the optimal policy of each decision in the form of a *policy tree* (see Figure 4). A policy tree (PT) is similar to a decision tree (DT) (Raiffa and Schlaifer, 1961): it consists of chance and decision nodes, and arcs labeled with the states of the nodes. The ancestors of a decision node in the PT are informational predecessors in the ID. Leaves indicate the optimal decision in the corresponding scenario. In contrast with DTs, a PT only represent scenarios that are possible by following the optimal strategy. This reduces enormously the size of the representation and makes it more understand-

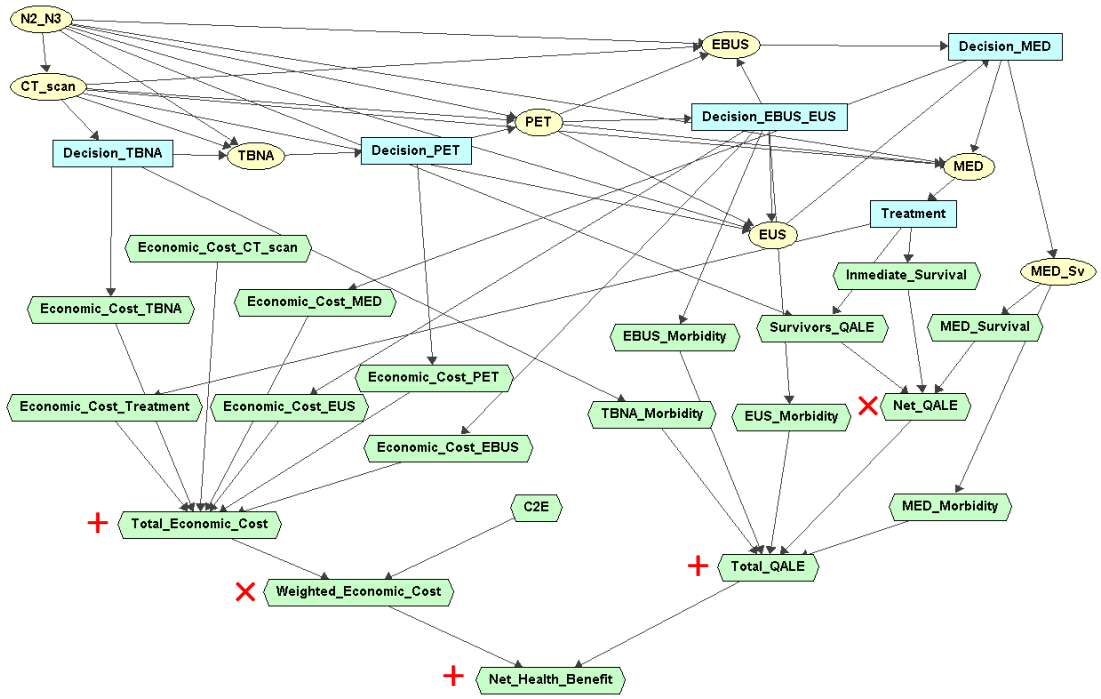

Figure 2: A new version of MEDIASTINET, including economic costs.

TBNA	positive	positive	positive	positive	negative	negative	negative	negative	no_result	no_result	no_result	no_result
Decision_TBNA	yes	yes	no	no	yes	yes	no	no	yes	yes	no	no
CT_scan	positive	negative	positive	negative	positive	negative	positive	negative	positive	negative	positive	negative
Decision_PET	no	no	yes	yes	no	no	yes	yes	yes	yes	no	no

Figure 3: Policy table for *Decision_PET*

able for the medical expert. For example, the policy table for decision *Treatment* of MEDIASTINET has 15,552 columns. In contrast, the PT of *Treatment* in MEDIASTINET when considering costs has 5 leaves (see Figure 4). That PT also represents the entire optimal strategy of the ID.

3.2 SUBJECTIVE EVALUATION OF MEDIASTINET'S STRATEGIES

After obtaining the two optimal strategies we have presented it to the expert to know his opinion about the policies obtained. He said that he would apply a slightly different strategy but he is not sure whether his decisions are better than those recommended by MEDIASTINET. However, the expert's recommendation and MEDIASTINET's agree in that the treatment must be selected depending on the result of the last test performed: If it is positive then apply chemotherapy, otherwise apply thoracotomy.

The expert concluded that the optimal strategies yielded by MEDIASTINET were very reasonable and "logic", and that the system was "quite intelligent."

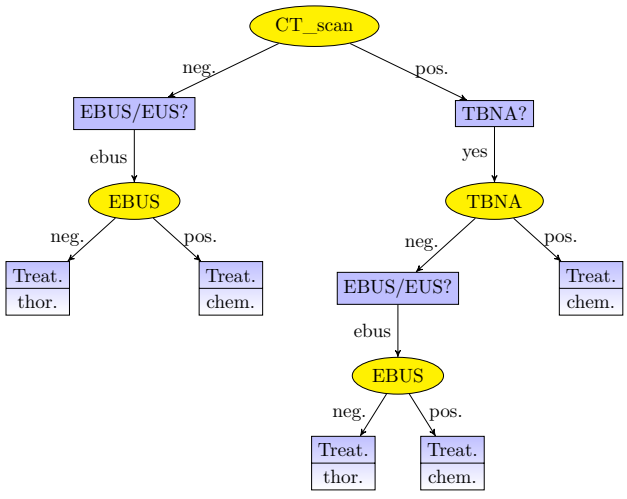

Figure 4: Optimal strategy for MEDIASTINET with economic costs.

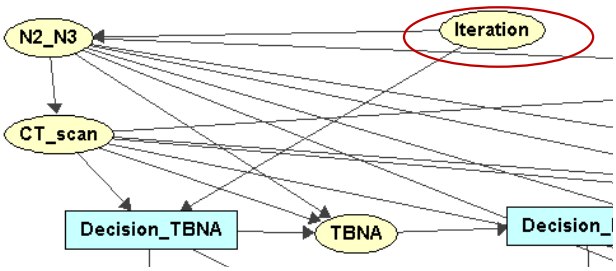

Figure 5: Part of the augmented ID of MEDIASTINET for performing SA of the prevalence of node *N2_N3*.

4 PROBABILISTIC SENSITIVITY ANALYSIS IN MEDIASTINET

After computing the optimal policies and the MEU for the reference case, we investigated whether the results depend on (are sensitive to) the uncertainty in the model. This post-hoc investigation is called *sensitivity analysis* (SA).

4.1 UNCERTAINTY ON THE NUMERICAL PARAMETERS OF MEDIASTINET

We have performed a SA that can be characterized as *quantitative*, *probabilistic*, *multi-one-way*.[5]

It is *quantitative* because we only consider variations in the numerical parameters and do not vary the structure of the ID. It is *probabilistic* because we have a probability distribution for each parameter.

It is *multi-one-way* because we consider the individual variations of a set of parameters, as for example in a tornado diagram.

Depending on the effects analyzed, *value SA* measures variations in the EU, and *decision SA* explores the changes in the optimal strategy.

For the SA we have built an *augmented ID* for each parameter. For example, Figure 5 shows the augmented ID for performing SA of the prevalence of *N2_N3*, identical to MEDIASTINET except that we have added the node *Iteration* and two arcs: one to the node affected by the parameter, *N2_N3*, and another to the first decision of the ID. Because of the non-forgetting hypothesis, this link implies that we will obtain the optimal strategy for each value of the parameter under study, and we can determine whether it is the same as the optimal strategy for the reference case. It also allows us to calculate the expected value of perfect information (Felli and Hazen, 1998).

[5]A complete definition of the characterizations of SA in IDs can be found in (Nielsen and Jensen, 2003).

All the chance and decision variables in MEDIASTINET are discrete. Each continuous distribution has been discretized by taking 100 points of an interval of the domain of the parameter. The intervals partitioned for normal and log-normal distributions of parameters μ and σ^2 have been $[\mu - k \cdot \sigma, \mu + k \cdot \sigma]$ and $[e^{\mu - k \cdot \sigma}, e^{\mu + k \cdot \sigma}]$ respectively, by using $k = 3.5$, which accumulates 99.953 % of the probability mass. We have taken the entire interval $[0, 1]$ when discretizing beta distributions.

4.2 RESULTS OF THE SA

We recorded three metrics of analysis:

- the thresholds of policy change, which define a set of intervals, contained in the domain of the parameter, where the optimal strategy is identical to the reference case,
- the expected value of perfect information (EVPI), very well-known in SA literature (Felli and Hazen, 1998), and
- the sensitivity of each decision to each parameter, which analyzes the probability of change in the optimal policies when the parameter varies.

4.2.1 Thresholds of policy change

Our analysis has shown that most of the parameters have a wide range of variation where the optimal policies do not change, and thus the optimal strategy is very robust. However, there are some exceptions, such as the sensitivity of CT scan. Its value in the reference model is 0.51, but its policy change thresholds are given by the interval $[0.41, 0.574]$. It means the value of sensitivity of CT scan in the reference model is not very far from the thresholds, and some policy might change if the value of the parameter varies.

4.2.2 EVPI

Most of the parameters present very small values of EVPI. The parameter with the highest EVPI is λ, as we expected. Thus, knowing its value with certainty would have a high impact on the MEU of the ID.

4.2.3 Sensitivity of each decision to each parameter

We have also observed that decisions are not sensitive to the variations of the parameters in most cases, which indicates that the optimal strategy is very robust.

The three parameters with highest probability of changing the optimal policies are: (1) the QALE of the

survivors to the thoracotomy when there is no metastasis; (2) the sensitivity of the TBNA when the result of CT scan is negative; and (3) λ. The only parameter that affects the policies of all the decisions is λ.

Finally, the decisions more affected by the variations on the parameters are *Decision_TBNA* and *Treatment.*

The main conclusion of the SA is that there are only two parameters that can have significant impact on the strategy: the QALE of the survivors to the thoracotomy when there is no metastasis and λ. Even though the former is the parameter with the highest impact on a decision (*Dec_TBNA*), the parameter that reflects to have more overall impact in the strategy is λ.[6]

5 RELATED WORK

Our model MEDIASTINET has several differences with the ID for the mediastinal stating of NSCLC built by Nease and Owens (1997):

1. MEDIASTINET assumes that a CT scan is always performed.
2. Four new laboratory tests have been included, namely *TBNA*, *PET*, *EBUS*, and *EUS*, as well as the decisions about whether to perform them.
3. MEDIASTINET considers the morbidities of the tests.
4. In MEDIASTINET the results of CT scan and PET influence the sensitivity and specificity of the other tests.
5. Palliative care is a possible treatment.
6. The economic costs of tests and treatments and λ (willingness to pay) are represented in MEDIASTINET.

As a result, MEDIASTINET is much bigger and more complex than the model of Nease and Owens (1997). For example, the decision table of the treatment has a domain of 72 columns in their model, while it contains 15,552 scenarios in MEDIASTINET.

Nease and Owens (1997) also built an ID that analyzes any arbitrary sequencing of CT scan and mediastinoscopy. In contrast, the order of decisions has been set in MEDIASTINET because the dependence relations of the test results in the problem are quite difficult to analyze in an ID with partial order and would need additional expert help. This would require a hard work of elicitation because the result of a test in our model can influence the sensitivity and specifity of future tests. For example, if the result of CT scan is positive then it also gives valuable information about where the doctor has to stick in the needle during the TBNA. However, that information is not available if the TBNA is performed before the CT scan.

We discarded the use of multi-currency IDs (Nielsen et al., 2007) for representing and solving the problem because that approach is a bit more difficult to be understood by a medical expert and there are no software tools with explanation capabilities for multi-currency IDs. In contrast, explanation capabilities of Elvira system for IDs (Lacave et al., 2007) have been quite useful while building and debugging the model with the expert's help.

Although quantitative SA has also been studied by Nielsen and Jensen (2003), the main preliminary steps in PSA in IDs can be found in (Felli and Hazen, 1998) and (Bielza et al., 2000). However, they do not consider the computation of the thresholds of policy change and the sensitivity of each decision to each parameter.

6 CONCLUSIONS AND FUTURE WORK

We have built an ID, MEDIASTINET, for the mediastinal staging of NSCLC.

From a medical point of view, there is a vivid debate among specialists about which technologies should be used for the mediastinal staging of NSCLC, and it is not possible to arrive at a consensus (Fritscher-Ravens et al., 2003; Schimmer et al., 2006). For this reason, MEDIASTINET is very useful as a decision analysis tool that combines objective data and subjective estimates and may show whether the discrepancies are due to differences in the numerical parameters used by each expert or to a wrong estimation of the consequences of each policy.

From the perspective of IDs, we have proposed a method for finding a tradeoff between cost and effectiveness, based on λ, the willingness to pay, which is also used in cost-effectiveness analyses. This parameter has been included explicitly in our model, which allows us to modify its value easily.

Additionally, we have performed a probabilistic sensitivity analysis that has studied three metrics, one of them is very well known (EVPI), and the others are new: the probability of change in the optimal strategy, and the intervals of the parameters where the optimal policies do not change. We have used for each param-

[6]The overall impact in the strategy is calculated as the average impact on each of the decisions of the ID.

eter the most appropriate distribution: beta, normal, or log-normal. We have proposed efficient methods for recording the three metrics when analyzing the variations of all the parameters on an ID of considerable size such as MEDIASTINET.

Finally, due to the interest of the expert in considering the possibility of having partial orderings of the decisions, unconstrained IDs (Jensen and Vomlelová, 2002) are a future research topic line. Decisions were totally ordered in MEDIASTINET because tests are not independent given that sensitivities and specificities can be influenced by other tests, as we explained above. A partial order would require a hard work of elicitation for every possible ordering.

The expert would also like to include in the model the possibility of repeating some decisions, which is known as *restaging*.

Acknowledgements

This work has been supported by the Department of Education of Madrid and the European Social Fund under a doctoral grant, and by the Spanish Ministry of Education of Science under grants TIN2006-11152 and TIN2009-09158.

References

Bielza, C., Ríos-Insua, S., Gómez, M., and del Pozo, J. A. F. (2000). Sensitivity analysis in Ictneo. *Lecture Notes in Statistics*, 152:317–334.

Felli, J. C. and Hazen, G. B. (1998). Sensitivity analysis and the expected value of perfect information. *Medical Decision Making*, 18(1):95–109.

Fritscher-Ravens, A., Bohuslavizki, K. H., Brandt, L., Bobrowski, C., Lund, C., Knofel, W. T., and Pforte, A. (2003). Mediastinal lymph node involvement in potentially resectable lung cancer. *Chest*, 123(2):442–451.

Gold, M. R., Siegel, J. E., Russell, L. B., and Weinstein, M. C. (1996). *Cost-Effectiveness in Health and Medicine*. Oxford University Press, New York.

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.

Jensen, F. V. and Vomlelová, M. (2002). Unconstrained influence diagrams. In *Proceedings of the Eighteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI'02)*, pages 234–241, San Francisco, CA. Morgan Kaufmann.

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.

Lloyd, C. and Silvestri, G. A. (2001). Mediastinal Staging of Non-Small-Cell Lung Cancer. *Cancer Control*, 8(4):311–317.

Nease, R. F. and Owens, K. D. K. (1997). Use of influence diagrams to structure medical decisions. *Medical Decision Making*, 17(3):263–275.

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. V. (1999). Welldefined decision scenarios. In *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence (UAI'99)*, pages 502–511, San Francisco, CA. Morgan Kaufmann.

Nielsen, T. D. and Jensen, F. V. (2003). Sensitivity analysis in influence diagrams. *IEEE Transactions on Systems, Man, and Cybernetics, Part A*, 33:223–234.

Raiffa, H. and Schlaifer, R. (1961). *Applied Statistical Decision Theory*. MIT press, Cambridge.

Sacristán, J. A., Oliva, J., del Llano, J., Prieto, L., and Pinto, J. (2002). ¿Qué es una tecnología sanitaria eficiente en España? *Gaceta Sanitaria*, 16:334–343. In Spanish.

Schimmer, C., Neukam, K., and Elert, O. (2006). Staging of non-small cell lung cancer: clinical value of positron emission tomography and mediastinoscopy. *Interact CardioVasc Thorac Surg*, 5(4):418–423.

Stinnett, A. A. and Mullahy, J. (1998). Net health benefits: A new framework for the analysis of uncertainty in cost-effectiveness analysis. *Medical Decision Making*, 18:S68–S80.

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

Weinstein, M. and Statson, W. (1977). Foundations of cost-effectiveness analysis for health and medical practices. *New England Journal of Medicine*, 296:716–721.