

Decision analysis networks

Francisco Javier Díez

Dept. Artificial Intelligence. UNED. Madrid, Spain

Manuel Luque

Dept. Artificial Intelligence. UNED. Madrid, Spain

Caroline Leonore König

Dept. Artificial Intelligence. UNED. Madrid, Spain

Abstract

In this paper we introduce a new type of probabilistic graphical model, called decision analysis networks (DANs). Like influence diagrams (IDs), DANs are much more compact and easier to build than decision trees, and are able to represent conditional independencies. However, IDs are unable to represent many real-world problems, as they require a total ordering of the decisions and have difficulties to represent structural constraints, while DANs can easily represent those kinds of problems. We also show that DANs compare favorably with other formalisms proposed recently for modeling asymmetric decision problems.

1 Introduction

The two formalisms most widely used for the representation and analysis of decision problems are decision trees (DT) and influence diagrams (IDs) (Howard and Matheson, 1984). DTs have the advantage of almost absolute flexibility, but 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). Even in cases with only a few chance variables, this preprocessing of probabilities is a difficult task. On the other hand, IDs have the advantages of being very compact, representing conditional independence, and using direct probabilities, and the drawback that they can only represent symmetric decision problems. However, most real-world problems are asymmetric: there is *structural asymmetry* when the value taken on by a variable restricts the domain of other variables, and there is *order asymmetry* when several ordering of the decisions are possible, for instance, when the decisions about

what tests to perform can be taken in any order (Jensen et al., 2006; Bielza et al., 2011).

Several formalisms have been proposed for representing and solving asymmetric decision problems—see Section 4. However, as far as we can tell, none of them has been used to build any real-world application, which may be a sign that none of them is completely satisfactory. For this reason, we present a new formalism, called *decision analysis networks* (DANs), which, in our opinion, can represent asymmetric problems more naturally.

The rest of the paper is structured as follows. First, we introduce the *n*-test problem, which will serve us to illustrate the properties of DANs and to compare different formalisms. Section 2 presents the definition of DANs, Section 3 explains how to convert any DAN into a DT, Section 4 compares DANs with other formalisms, Section 5 discusses some lines for future research, and Section 6 contains the conclusions.

Example 1. The *n*-test problem consists in deciding how to treat a patient that may suffer from a certain disease *D*. After an initial exam-

ination of the symptoms, the doctor may order one or several of n available tests, each one having a cost. It is assumed that each test can be performed only once and that the result of each test will be known immediately or, at least, before making a decision about other tests. In its simplest version, we assume that there is only one symptom, S , and all the variables are dichotomic, i.e., that the disease and the symptom are either present or absent, and the result of each test is either positive or negative.

An instance of this problem is the diagnosis of diabetes based on two tests: blood and urine, where BT (UT) represents the decision of whether performing the blood (urine) test and by B (U) represents the result of the test (Demirer and Shenoy, 2006).

2 Definition of a DAN

2.1 Graph and variables of a DAN

The set of variables of a DAN, \mathbf{V} , is formed by three disjoint subsets, \mathbf{C} (chance variables), \mathbf{D} (decisions), and \mathbf{U} (utilities): $\mathbf{V} = \mathbf{C} \cup \mathbf{D} \cup \mathbf{U}$. The first represent real-world properties that are not under the direct control of the decision maker, while the utilities represent his/her preferences. We denote the variables by capital letters and their values by lower-case letters. A bold upper-case letter denotes a set of variables and a bold lower-case letter denotes a configuration of them.

The graph of a DAN is an acyclic directed graph (ADG) such that each node represents a variable; for this reason we will speak of nodes and variables indifferently. We use the terms *link* and *arc* as synonyms. In this paper we will assume that utility nodes cannot have children.

The DAN in Figure 1 contains four chance variables, drawn with circles, three decisions, drawn with rectangles, and three utilities, drawn with diamonds.¹

It is possible to define a list of stages for a DAN, representing different phases of the decision process, such that each node belongs to

¹This DAN, encoded in the ProbModelXML format, can be found at www.cisiad.uned.es/ProbModelXML/examples and opened with OpenMarkov, an open-source tool available at www.openmarkov.org.




Figure 1: A DAN for the diabetes problem.

one stage. However, in this paper, due to space limits, we describe only one-stage DANs.

2.2 Restrictions

A *restriction* associated to a link $X \rightarrow Y$, where X and Y are chance variables or decisions, 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 x . If X and Y are discrete, the restrictions associated with this link can be represented by a compatibility table with a column for each value of x and a row for each value of y ; the cell (x, y) contains a 1 when x and y are compatible and 0 when there is a restriction.² For example, Table 1 means that when the decision about blood test (BT) is not to perform it, the result (B) is neither positive nor negative.

	BT	$+bt$	$-bt$
B	$+b$	1	0
	$-b$	1	0

Table 1: Compatibility table for the link $BT \rightarrow B$.

2.3 Potentials

A potential ψ is a function that maps each configuration \mathbf{x} of a set of variables \mathbf{X} onto a value

²In OpenMarkov, cells corresponding to compatible pairs are colored in green and those corresponding to restrictions are colored in red. Arcs with restrictions are marked with a short perpendicular double line. The compatibility table and the revelation conditions (see below) associated to a link can be accessed from its contextual menu.

of the set $\mathbb{R} \cup \{-\}$, i.e., $\psi(\mathbf{x})$ is either a real number or “-”. Each chance variable Y whose parents in the graph are \mathbf{X} has an associated potential, denoted by $P(y|\mathbf{x})$, satisfying three conditions:

- $P(y|\mathbf{x}) \in [0, 1] \cup \{-\}$
- if there is a restriction (x_i, y) and $\mathbf{x}^{\setminus X_i} = x_i$ (i.e., the i -th value of configuration \mathbf{x} is the same as x_i in the restriction), then $P(y|\mathbf{x}) = -$;
- for each configuration \mathbf{x} of \mathbf{X} , if some of the values of $P(y|\mathbf{x})$ are real-numbers, then their sum is 1.

Table 2 shows the potential for variable B , assuming that the sensitivity of this test for disease D is 0.92 and its specificity is 0.97.

BT		$+bt$		$-bt$	
D		$+d$	$-d$	$+d$	$-d$
B	$+b$	0.92	0.03	-	-
	$-b$	0.08	0.97	-	-

Table 2: Potential $P(b|bt, d)$, associated to node B .

2.4 Representing the flow of information

In DANs there are three ways to indicate when a variable becomes observed: always observed variables, revelation arcs, and stages.

2.4.1 Always observed variables

We can declare some chance variables as *always observed*, which means that we do not need to take any action to know their value. For instance, in Figure 1 variable S is put in a double circle, because we always know whether the patient has the symptom, without performing any test.

2.4.2 Revelation arcs

Given a link $D \rightarrow Y$, where D is a decision and Y a chance variable, we can declare that some values of D *reveal* the value of Y ; we then say that $D \rightarrow Y$ is a *revelation arc*. For instance, the revelation arc $BT \rightarrow B$ in

Figure 1, drawn as a double-line arrow, indicates that when we decide to perform the test ($BT = +bt$), its result becomes known for future decisions.

2.5 Summary: The meaning of links

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

1. *Causal influence*: when Y is a chance node.
2. *Functional dependence*: when Y is a utility node.
3. *Temporal order*: when both X and Y are decisions.
4. *Revelation*: when X is a decision and Y is a chance node.
5. *Restriction*: when X and Y are chance or decision nodes.

The first three meanings are the same as in IDs. Revelation arcs in DANs substitute information arcs in IDs.

3 Equivalent decision tree

In this section we explain how to convert a DAN into a DT, with the main purpose of providing a clear semantics for DANs. 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 .

The initial step is to pick up the always-observed variables, order them arbitrarily (the order will not affect the evaluation of the tree) and organize them into a tree: the first variable will be the root, there will be a branch for each value of this variable, each branch will contain a node for the second variable, and so on. In the diabetes problem, there is only one always-observed variable, S , which will be the root of the tree. It has two branches, $+s$ and $-s$, as shown in Figure 2.

Then we consider the order of the decisions. For each pair of decisions, if there is a directed path from D_1 to D_2 , then D_1 must be closer to the root in the tree. For example, the links $BT \rightarrow Tr$ and $UT \rightarrow Tr$ mean that the decisions BT and UT must be made before Tr , but the DAN does not specify which of BT and UT must be made first. Therefore, the decision maker faces one additional decision, that




Figure 2: A decision tree for the DAN in Figure 1.

we denote by OD (order of decisions). For each branch of S , we add a decision node OD with two branches, labelled bt (which means that decision BT is made first) and ut . In the bt branch, we place a decision node BT , and in the ut branch a node UT .

If the value d of decision D reveals the value of some variables, these must appear immediately in the branch d . For example, given that $+bt$ reveals the value of B , we place in this branch a chance node B . For each branch of this node there are two remaining decisions, UT and Tr . As UT must be made first, we place a decision node UT in each branch. In the branch $-bt$ no variable is revealed.

Each of the three UT nodes has two branches, $+ut$ and $-ut$, which are expanded as in the case of BT . Next we place the last decision, Tr , which does not reveal the value of any variable, and finally we place in the tree the unobserved variables. In the upper scenario, the only unobserved variable is D . In the scenarios containing the $-bt$ branch shown in Figure 2, we should add also the variable B , but because of the restrictions in Table 1, no value of this variable is possible, which means that the variable B does not make sense in these scenarios, and therefore it must not appear in this part of the tree.

Finally, we place at each branch a leaf node, which represents a utility.

Each path from the root node to a leaf node defines a *scenario*. The probability of

each scenario is the product of the conditional probabilities involved in it, and its utility is the sum of the utility functions. For example, the probability of the upper scenario $(+s, bt, +bt, +b, +ut, +u, +tr, +d)$ is $P(+s|+d) \cdot P(+b|+d, +bt) \cdot P(+u|+d, +ut) \cdot P(+d)$ and its utility is $u_1(+bt) + u_2(+ut) + u_3(+tr, +d)$. Using the probabilities of the scenarios, it is possible to compute the conditional probability for each branch of the tree given the values of the variables at the left of that branch by normalizing them. This concludes the construction of the DT.

4 Comparison of different formalisms

In this section we examine seven formalisms for decision analysis. First we compare DANs with IDs, which are the standard probabilistic graphical model for decision analysis, and then (in Sec. 4.2) we analyze six formalisms for asymmetric decision problems.

4.1 DANs vs. IDs

There are many problems for which the DAN representation is much more simple than its ID counterpart. Let us consider, for instance, the one-test problem, in which the ID needs to introduce a dummy value *no-result* for the variable that represents the result of the test, which complicates the CPT for this variable, which will have 3 rows instead of 2. In the n -test prob-

lem, we need a decision node, T_1 , representing which test is performed in the first place, if any; therefore the domain of T_1 has $n + 1$ values. Then we need a chance variable, R_1 , for the result of the first test. Its probability table must contain two columns, $+d$ and $-d$, for each value of T_1 , which makes a total of $2(n + 1)$ columns and 3 rows. We need another decision node, T_2 , indicating which test is made in the second place. If we wish to indicate that the test performed first cannot be repeated, and that if the first decision is not to perform any test then the second one must also be not to test, we find a serious difficulty, because standard IDs do not admit restrictions. We might assign n values to T_2 , representing the $n - 1$ remaining test plus the possibility of no test, but the meaning of each value of T_2 would depend on the test performed in the first place. Therefore, the CPT for the result of the second test, R_2 , is difficult to build, and its numerical values are the same as some of those in the table for R_1 , which complicates the maintainability of the model and makes sensitivity analysis impossible, unless we introduce symbolic names for the parameters. In summary, the representation of the n -test problem with IDs is difficult for $n = 2$ and virtually unfeasible for $n > 2$.

This difficulty is due to one of the basic properties of IDs: they can only represent symmetric problems. In a few cases, asymmetric problems can be made symmetric by adding dummy states, such as *no-result*, but IDs are unsuitable for many asymmetric problems, such as the n -test problem, that can be easily represented with DANs.

In general, for each ID there is an equivalent one-stage DAN; this is the case for all the IDs we have found in the literature and in our practice. In the extremely rare cases in which there is no one-stage DAN for a certain ID, it is always possible to build an equivalent multi-stage DAN.

4.2 Comparison with other formalisms for asymmetric decision problems

In this section we compare briefly six formalisms for representing asymmetric decision

problems: influence diagrams with constraints (IDCs) (Smith et al., 1993), asymmetric influence diagrams (AIDs) (Nielsen and Jensen, 2000), sequential valuation networks (SVNs) (Shenoy, 2000; Demirer and Shenoy, 2006), unconstrained influence diagrams (UIDs) (Jensen and Vomlelová, 2002; Ahlmann-Ohlsen et al., 2009), sequential influence diagrams (SIDs) (Jensen et al., 2006), and decision analysis networks (DANs)—see also the comparison in (Bielza and Shenoy, 1999). Table 3 summarizes the main features of each method.

The first column indicates whether the method requires a total ordering of the decisions. The methods that impose this condition, such as IDCs, AIDs and SVNs, have to include decision nodes with as many states as the total number of decisions, as we explained in Section 4.1 when discussing why IDs cannot represent the n -test problem, and for the same reason, these formalisms are unsuitable for decision problems involving order asymmetries.

The second column means that IDs, IDCs, and UIDs need to add dummy states, such as *no-result*, to reflect the fact that the test result is not available when the test is not performed. This complicates the edition of the probability tables and the policies obtained.

The third column indicates which methods need to add labels to some arcs. In AIDs, the labels indicate which variable is revealed depending on which test has been performed. The problem is that a decision node representing the n tests will have n outgoing arcs, which complicates the representation when n is large. In the case of SVNs and SIDs, the labels indicate the sequence of decisions and observations, and again this makes the graph difficult to read for the n -test problem and for many others. In the case of a revelation arc $X \rightarrow Y$ in a DAN, the indication of which values of X reveal the value of Y can be seen as a label for that arc; however, this label only depends on X , not on any sequence of previous decisions and observations, and therefore the label is always very simple, independently of the complexity of the decision problem. In contrast, UIDs do not use labels because they assume that decisions are

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

Table 3: Main features of several methods for representing decision problems.

binary: Y is revealed if and only if X takes the value *true*. This is a limitation for non-binary decisions. For example, an echocardiogram can be transtoracic or transesophageal (decision X); the sensitivity and specificity of a finding, such as the calcification of a valve, depend significantly on the type of echocardiogram, but this difference cannot be modelled if the decision is modelled as a binary variable: do test / do not test.

Finally, the forth column shows which methods need information arcs, which are all except UIDs and DANs. Using these arcs forces the expert to indicate whether the variable is known or not to when making a decision, and complicates the graph when there are many always-observed variables and also when there are many sequences of decisions that can make a variable observed. In the n -test problem, we would need an arc from each result-of-test variable to each test decision. König (2012) has built an IDC, an AID, a UID, an SVN, and an SID for the diabetes problem. Seeing how cumbersome those diagrams are, even for a two-test problem, it is easy to understand why none of these methods is appropriate for the n -test problem.

In contrast, the graph of the UID for the diabetes problem is virtually the same as that of the DAN (see Figure 1), because none of these models require information arcs. This is a crucial advantage in many cases, such as the n -test problem, but turns out to be a problem in other situations, because both UIDs and one-stage DANs assume that the value of a variable is known immediately after making the decision

that reveals it. Multi-stage DANs overcome this difficulty, but UIDs do not.

In the same way, if any of two different decisions can reveal the value of a variable, this cannot be represented with an UID, either, because UIDs assume that a variable is known only when all its parent decisions have been made. (In contrast, in a DAN it suffices that one of the decisions is made.) This means that there are some problems that can be represented with IDs, but not with UIDs, which is a drawback of UIDs with respect to all the other formalisms listed in Table 3.

Finally, another drawback of UIDs, due to the lack of restrictions, is that they often require dummy states, which make the representation more obscure and less efficient.

A more detailed comparison of these formalisms for different decision problems can be found in (König, 2012).

5 Future work

In the paper that presented influence diagrams (Howard and Matheson, 1984) by the first time, the only evaluation method proposed was to build an equivalent decision tree. Later, several researchers developed more efficient methods, including algorithms for dealing with continuous variables. Similarly, in this paper we have introduced DANs and explained how to convert them into decision trees, which, given the power of today’s computers, is sufficient for many real-world problems. In a future paper we will present more efficient algorithms. As an advance, we can say that symmetric DANs can be evaluated by variable elimination and arc

reversal, as in the case of IDs; structural asymmetries can be tackled by adding dummy nodes in the evaluation; order asymmetries can be solved by decomposing the problem into a set of completely ordered subproblems, as in (Demirer and Shenoy, 2006) or by building an S-DAG, as in (Ahlmann-Ohlsen et al., 2009; Luque et al., 2010). DANs containing numeric variables might be solved by adapting some of the algorithms for IDs (Cobb and Shenoy, 2008; Bielza et al., 2011).

Another pending task is to complete the implementation of DANs in `OpenMarkov`,³ an open-source tool for probabilistic graphical models. Currently this tool can be used to edit DANs, but it cannot yet evaluate them.

6 Conclusion

In this paper, we have proposed a new type of probabilistic graphical model, called *decision analysis networks* (DANs), and have explained how to convert each DAN into a decision tree. The main purpose of this conversion is to give a clear semantics to our model (in other models presented in the literature, the semantics is described only by an example). Any other algorithm should prove that it returns the same expected utility and optimal strategy as the evaluation of the tree.

A second reason for this conversion is that many decision analysts will understand DANs much better if they can examine the equivalent decision tree—in medicine, for example, the standard analysis tool are decision trees, while influence diagrams are almost unknown (Pauker and Wong, 2005), let alone the recent methods for asymmetric decision problems discussed in Section 4.2. In fact, the conversion of IDs into DTs was one of the explanation facilities proposed in (Lacave et al., 2007).

As explained in Section 4.1, for every influence diagram there is an equivalent DAN, but the converse is not true: there are many cases that are difficult if not impossible to represent with IDs, but can be modelled easily with DANs, such as the n -tests problem. Even

for problems that can be modelled with IDs, there is usually a DAN being more simple and more intuitive than the ID, because it does not need dummy states; an example is the one-test problem. The main reason why IDs have troubles to represent order asymmetry is the use of information arcs, as explained in that section, and therefore the other formalisms proposed recently for representing asymmetric decision problems face the same difficulty. The only exception are unconstrained influence diagrams (UIDs), which share with DANs the advantage of not using information arcs. However, unlike the rest of the models, UIDs have the drawback that for some IDs there is no equivalent UID. Another difference between UIDs and DANs is that the former do not use restrictions. It would be possible to extend UIDs with restrictions, stages, and labels for indicating which values of the decision reveal the value of a chance variable, but there would still be a subtle difference in the semantics of both types of models: in a DAN, every chance variable is revealed as soon as we decide to observe it, while in UIDs its value is not known until we make all the decisions about the actions that may reveal it.

Another difference between DANs and the rest of the models for asymmetric problems is that these have been illustrated with toy problems, and apparently they have not been used to build any real-world application. In contrast, we arrived at DANs when trying to solve an instance of the n -test problem: the mediastinal staging of non-small cell lung cancer. There are several tests available (CT-scan, PET, TBNA, EBUS, EUS, mediastinoscopy, etc.) and the experts do not agree on the optimal sequence of them. Clearly, similar problems are typical in other domains, such as troubleshooting different types of devices.

Additionally, it seems that none of those models has been implemented, with the exception of a program for evaluating UIDs that is not available outside the research group that used it for some experiments (Ahlmann-Ohlsen et al., 2009). In contrast, it is possible to edit DANs in `OpenMarkov` and store them in the Prob-

³www.openmarkov.org.

ModelXML format,⁴ and in the near future this tool will have algorithms for evaluating them.

Given the advantages of DANs over decision trees and influence diagrams and the fact that we have found so far no problem based on the non-forgetting assumption—which is also implicit in decision trees—that cannot be represented with DANs, we conjecture that this formalism will play a significant role as a tool for decision analysis.

Acknowledgments

We thank the anonymous reviewers of the PGM-2012 workshop for many useful comments.

This work has been supported by grants TIN2006-11152 and TIN2009-09158 of the Spanish Ministry of Science and Technology.

References

- [Ahlmann-Ohlsen et al.2009] K. S. Ahlmann-Ohlsen, F. V. Jensen, T. D. Nielsen, O. Pedersen, and M. Vomlelová. 2009. A comparison of two approaches for solving unconstrained influence diagrams. *International Journal of Approximate Reasoning*, 50:153–173.
- [Bielza and Shenoy1999] C. Bielza and P. P. Shenoy. 1999. A comparison of graphical techniques for asymmetric decision problems. *Management Science*, 45:1552–1569.
- [Bielza et al.2011] C. Bielza, M. Gómez, and P. P. Shenoy. 2011. A review of representation issues and modelling challenges with influence diagrams. *Omega*, 39:227–241.
- [Cobb and Shenoy2008] B. R. Cobb and P. P. Shenoy. 2008. Decision making with hybrid and influence diagrams using mixtures of truncated exponentials. *European Journal of Operational Research*, 186:261–275.
- [Demirer and Shenoy2006] R. Demirer and P. P. Shenoy. 2006. Sequential valuation asymmetric decision problems. *European Journal of Operational Research*, 169:286–309.
- [Howard and Matheson1984] R. A. Howard and J. E. Matheson. 1984. Influence diagrams. In R. A. Howard and J. E. Matheson, editors, *Readings on the Principles and Applications of Decision Analysis*, pages 719–762. Strategic Decisions Group, Menlo Park, CA.
- [Jensen and Vomlelová2002] F. V. Jensen and M. Vomlelová. 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.
- [Jensen et al.2006] F. V. Jensen, T. D. Nielsen, and P. Shenoy. 2006. Sequential influence diagrams: A unified asymmetry framework. *The International Journal of Approximate Reasoning (IJAR)*, 42:101–118.
- [König2012] C. König. 2012. Representing decision problems with Decision Analysis Networks. Master's thesis, Dept. Artificial Intelligence, UNED, Madrid, Spain.
- [Lacave et al.2007] C. Lacave, M. Luque, and F. J. Díez. 2007. Explanation of Bayesian networks and influence diagrams in Elvira. *IEEE Transactions on Systems, Man and Cybernetics—Part B: Cybernetics*, 37:952–965.
- [Luque et al.2010] M. Luque, T. D. Nielsen, and F. V. Jensen. 2010. An anytime algorithm for evaluating unconstrained influence diagrams. Technical Report CISIAD-10-05, Dept. Artificial Intelligence, UNED, Madrid, Spain.
- [Nielsen and Jensen2000] T. D. Nielsen and F. Jensen. 2000. Representing and solving asymmetric Bayesian decision problems. In *Proceedings of the 16th Annual Conference on Uncertainty in Artificial Intelligence (UAI-2000)*, pages 416–425, San Francisco, CA. Morgan Kaufmann.
- [Pauker and Wong2005] S. Pauker and J. Wong. 2005. The influence of influence diagrams in medicine. *Decision Analysis*, 2:238–244.
- [Shenoy2000] P. P. Shenoy. 2000. Valuation network representation and solution of asymmetric decision problems. *European Journal of Operational Research*, 121:579–608.
- [Smith et al.1993] J. E. Smith, S. Holtzman, and J. E. Matheson. 1993. Structuring conditional relationships in influence diagrams. *Operations Research*, 41:280–297.
- [Howard and Matheson1984] R. A. Howard and J. E. Matheson. 1984. Influence diagrams. In R. A. Howard and J. E. Matheson, editors, *Readings on*

⁴www.cisiad.uned.es/ProbModelXML.