JMLR: Workshop and Conference Proceedings vol 52, 462-473, 2016 PGM 2016

Decisions and Dependence in Influence Diagrams

Ross D. Shachter SHACHTER@STANFORD.EDU
Department of Management Science and Engineering
Stanford University
Stanford, CA 94305-4026

Abstract

The concept of dependence among variables in a Bayesian belief network is well understood, but what does it mean in an influence diagram where some of those variables are decisions? There are three quite different answers to this question that take the familiar concepts for uncertain variables and extend them to decisions. First is responsiveness, whether the choice for a decision affects another variable. Second is materiality, whether observing other variables before making a decision affects the choice for the decision and thus improves its quality. Third is the usual notion of dependence, assuming that all of the decisions are made optimally given the information available at the time of the decisions. There are some subtleties involved, but all three types of decision dependence can be quite useful for understanding a decision model.

Keywords: Influence diagrams; dependence; responsiveness; materiality; requisite sets; policy diagrams.

1. Introduction

Students learning how to construct even simple influence diagrams can be overwhelmed by the question "What does that decision depend on?" and soon find themselves drawing arcs into that decision node from almost every other node in the diagram. Of course, the parents of the decision represent the variables that will have been observed by the time that decision is made, and the value nodes in the model represent the criteria for making that decision. But the question remains, what does that decision depend on?

The concept of dependence among the variables in a Bayesian belief network has been studied and characterized well, but not so in an influence diagram where some of the variables are decisions. In this paper we present three complementary answers to the question, building on the familiar concepts of dependence among uncertain variables in belief networks.

First is *responsiveness*, whether the choice for a decision affects other variables. Second is *materiality*, whether observing additional variables before making a decision affects the choices made for the decision and thereby improves its quality. Closely related to materiality is a *requisite* set of variables, sufficient to determine an optimal choice. Third is the usual framework of independence applied to the belief network obtained by replacing each decision by its optimal policy as a function of variables that will be observed by the time of the decision. Together these three types of decision dependence help us better understand the decision problem represented by an influence diagram.

We will show in this paper how to answer the three questions, but the answers are not as straightforward as they might first appear. For example, consider the different influence diagrams shown in Figure 4. For which of these might the value V or the choice for $D2$ depend on $D1$? In 4a1, 4b1, 4c1, 4d1, 4e1, and 4f1, $D1$ is an ancestor of V, but, of these, only in 4d1 and 4e1 is there a path from $D1$ to V along conditional arcs, the condition proposed by Nilsson and Jensen (1998). We will see that $D1$ might affect V and the choice for $D2$ not only in those diagrams but also in 4a1. However, in 4b1, 4c1, and 4f1, it could not possibly affect either V or the choice for $D2$. Although we can recognize these properties for the latest decision in an influence diagram, such as $D2$, to recognize these properties for earlier decisions, such as $D1$, we need to construct belief networks with optimal policies replacing decisions, as shown in 4a2, 4b2, 4c2, 4d2, 4e2, and 4f2.

In the following sections: we explain our notation and framework; we consider responsiveness; we examine materiality and requisite sets; we generalize the framework of dependence to models with decisions; we provide a number of examples to illustrate these concepts; and we offer our conclusions and opportunities for future research.

2. Bayesian Belief Networks, Influence Diagrams, and Policy Diagrams

In this section we define Bayesian belief networks, influence diagrams, and policy diagrams, and introduce their associated notation.

A *Bayesian belief network* (BBN) graph is comprised of nodes, corresponding to *uncertain variables* $\mathbf{U}$ and *value variables* $\mathbf{V}$ [1], and directed *conditional* arcs without cycles. We denote single nodes by upper-case letters, such as X, and sets of nodes by bold letters, such as $\mathbf{X}$, and often refer interchangeably between nodes and their corresponding variables. We represent uncertain variables as ovals and value variables as rounded rectangles, as in the BBN shown in Figure 1b. If there is an (S, X) arc, then S is a *parent* of X, denoted $S \in Pa(X)$, and X is a *child* of S. If there is a directed path from S to D', then S is an *ancestor* of D' and D' is a *descendant* of S.

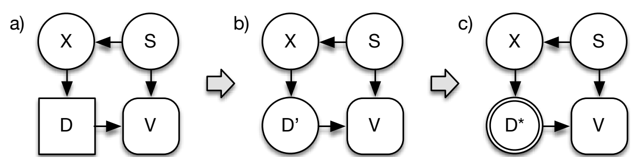

Figure 1: ID and corresponding PPD and DPD with imperfect observation X of state S

The structure of the BBN represents that X must be independent of its ancestors given its parents. This is generalized by *d-separation*, that $\mathbf{X}$ must be independent of $\mathbf{Y}$ given $\mathbf{Z}$ if there is no active trail between $(\mathbf{X} \setminus \mathbf{Z})$ and $(\mathbf{Y} \setminus \mathbf{Z})$, where an *active trail given* $\mathbf{Z}$ is a sequence of adjacent nodes such that each intermediate node has both arcs directed toward it if and only if it is in $\mathbf{Z}$. In that case, we say that $\mathbf{X}$ is *d-separated* from $\mathbf{Y}$ by $\mathbf{Z}$ (Pearl, 1988).

Some of the variables in the BBN can be *deterministic* variables, known with certainty when their parents are known. Deterministic variable X is said to be *functionally determined by* $\mathbf{Z}$ when each parent of X is either in $\mathbf{Z}$ or itself functionally determined by $\mathbf{Z}$. We represent deterministic variables as double ovals or double rounded rectangles, as in the BBN shown in Figure 1c, where D^* is functionally determined by X. The set of variables functionally determined by $\mathbf{Z}$ is denoted by $F(\mathbf{Z})$, and d-separation can be generalized to *D-separation*: $\mathbf{X}$ is *D-separated* from $\mathbf{Y}$ by $\mathbf{Z}$ if $(\mathbf{X} \setminus F(\mathbf{Z}))$ is D-separated from $(\mathbf{Y} \setminus F(\mathbf{Z}))$ by $F(\mathbf{Z})$ (Geiger et al., 1990a).

1. Expected utility is usually represented by deterministic value nodes and only in influence diagrams, but we are including uncertain value variables in our BBNs for several reasons. First, we can test them for independence at the graphical and numerical levels. Second, the expectation of value variables yields deterministic utility functions and we can assume that value variables are binary without loss of generality. Finally, we will be converting influence diagrams that include these value nodes into BBNs where they will continue to be value nodes.

Finally, to *fully specify* the BBN, we need to have a finite set of states for each variable and a probability distribution P. Probability distribution P is *consistent with* the BBN if

$$P\{\mathbf{U} \cup \mathbf{V}\} = \prod_{X \in \mathbf{U} \cup \mathbf{V}} P\{X|Pa(X)\},$$

and $P\{X|Pa(X)\}$ is a deterministic distribution for any deterministic variables $X \in \mathbf{U} \cup \mathbf{V}$. The following is a seminal result that we state as a proposition here because it is fundamental to the paper (Geiger et al., 1990a).

Proposition 1 (D-Separation is Sound and Complete) *Variables* $\mathbf{X}$ *and* $\mathbf{Y}$ *are D-separated by* $\mathbf{Z}$ *in a BBN graph if and only if* $\mathbf{X}$ *and* $\mathbf{Y}$ *are independent given* $\mathbf{Z}$ *for all probability distributions* P *consistent with the BBN.*

An *influence diagram (ID)* graph is comprised of nodes, corresponding to $\mathbf{U}$, $\mathbf{V}$, and *decision variables* $\mathbf{D}$, and directed arcs without cycles. The arcs into decisions are *informational* rather than conditional, and indicate that the parents of the decision will be observed by the *decision maker (DM)* at the time DM will choose an optimal alternative to maximize the expected sum of the value variables. We represent decision variables as rectangles, as in the ID shown in Figure 1a, where uncertain X will be observed before decision D. To fully specify the ID we need a finite set of alternatives for the decisions, states for the other variables, and probability distribution P for $P\{\mathbf{U} \cup \mathbf{V}|\mathbf{D}\}$ consistent with the ID.

There are several standard assumptions we make about IDs throughout this paper.

- We assume that the decisions are *totally ordered*, that is, there is a single directed path containing all of the decision nodes, and *no-forgetting*, that any observations and choices made in the past will be known for subsequent decisions. These ensure that there is no ambiguity about the information available to DM at the time of any given decision.
- We assume that there are no directed cycles. As a result, there can be no descendants of decision D known at the time the choice for D is made. (There are circumstances where this might be desirable, such as deterministic constraints on the sets of alternatives, but they could be represented instead with penalty value nodes.)
- We assume that there is at least one value node and that none of them have children. If one did, it could be replaced by an uncertain node with an equivalent value node child.

A *policy diagram* corresponds to an ID where each decision node is replaced by an uncertain *policy* node (Cooper, 1988; Shachter and Kenley, 1989). The policy diagram is *determinisitic (DPD)* if the policy nodes are deterministic and *probabilistic (PPD)* otherwise. The parents of each policy node initially include all of the explicit decision parents and implicit (no-forgetting) nodes that will be observed before the decision. For efficiency and understanding, we want to prune as many parents of the policy nodes as possible, while maintaining a sufficient set to represent optimal decisions.

We replace the decision node D in the ID with probabilistic node D' in the PPD and with deterministic node D^* in the DPD, as in the diagrams shown in Figure 1. As a convention, we assume that a PPD is not fully specified and our focus is on its graph, but that in a DPD the underlying states and probability distribution have been fully specified and all of the optimal deterministic policies have been computed. Thus, a DPD represents the optimal choices for DM but those are yet to be determined in a PPD. Surprisingly, it is necessary to have the PPD represent the policies as non-deterministic nodes. To understand why, consider the diagrams shown in Figure 1. In the DPD shown in 1c, where the optimal choice $D*$ is determined by X and therefore X D-separates $D*$ from V, we can not recognize any possible direct effect of the choice of $D*$ on V given X. On the other hand, in the PPD shown in 1b where non-deterministic D' is not D-separated from V by X, we can recognize whether V might depend on the choice of D' given X.

Computing the optimal policies is not the focus of this paper. It can be computationally challenging but it is well understood and there are many methods that could be used for it (Shachter, 1986; Cooper, 1988; Shenoy and Shafer, 1988; Tatman and Shachter, 1990; Shachter and Peot, 1993; Jensen et al., 1994; Bhattacharjya and Shachter, 2007; Luque and Diez, 2010; Nilsson and Lauritzen, 2013).

Multiple alternatives for a decision can be optimal given the information available. In that case, we assume that ties are broken by some arbitrary, but consistently applied, order of the alternatives.

3. Responsiveness

A variable X (not observable at the time of a decision D) is *responsive* to D if its value is affected by the choice for D (Heckerman and Shachter, 1995), that is, if there is a positive probability that a different choice for D would have led to a different value for X. While this definition is powerful for thinking about causality, it is subtle and complex. For example, consider the IDs shown in Figure 2. Suppose DM will *Call* "heads" or "tails" for a coin flip, will *Win* if the call is correct, and believes that either *Coin* face is equally likely and unaffected by *Call*. Therefore, $P\{Win|Call\} = P\{Win\} = 0.5$. Although all three diagrams are valid for this situation, the diagram in 2c does not adequately represent that *Win* is responsive to *Call*. That is, while *Win* and *Call* are independent, if DM had chosen a different *Call* with the same *Coin* result (because *Coin* is not responsive to *Call*) then *Win* would have had a different realization.

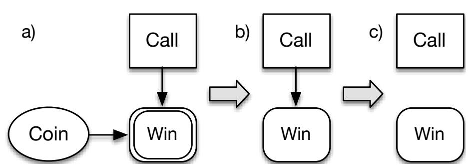

Figure 2: Influence diagrams for the decision to call "heads" or "tails" for a coin flip

To resolve dilemmas like this, we introduce a more limited notion of responsiveness. An uncertain or value variable X, not observable before a decision D is chosen, is *strongly responsive* to D if $P\{X|D, Pa(D)\}$ varies with the choice of D, that is, X depends on D given $Pa(D)$ probabilistically. In the coin flip example, *Win* is responsive to *Call* but *Win* is not strongly responsive to *Call*, because choosing a different *Call* changes whether DM wins but not the probability that DM wins.

Because strong responsiveness is defined in terms of dependence, we can use D-separation to characterize when it is possible. We can do so for any policy in a PPD, but only for the latest decision in an ID. Theorem 2 and Corollary 3 together show that X *might be* strongly responsive to the *latest* decision D in an ID (or any policy D' in a PPD) if and only if it is a descendant of D (or D'). For example, S can not be strongly responsive to D in the ID shown in Figure 1a, *Coin* can not be strongly responsive to *Call* in Figure 2a, nor can *Win* be strongly responsive to *Call* in Figure 2c.

Theorem 2 (Strong Responsiveness in the PPD) *Given any policy variable D' and any uncertain or value variable X in a PPD graph, the following statements are equivalent:*
1. X is a descendant of D';
2. X is not D-separated from D' by $Pa(D')$; and
3. there is some probability distribution consistent with the PPD for which X is strongly responsive to D, that is, D' and X are dependent given $Pa(D')$.

Proof (1) and (2) are equivalent because the active trail can only descend but without directed cycles there is nothing to block it. (2) and (3) are equivalent because of Proposition 1. ∎

Corollary 3 (Strong Responsiveness in the ID) *Given any decision D and any uncertain or value variable X in an ID graph:*
1. if D is the latest decision, there is some probability distribution consistent with the ID for which X is strongly responsive to D if and only if X is a descendant of D;
2. there is some probability distribution consistent with the ID for which X is strongly responsive to D if there is a path from D to X along conditional arcs; and
3. if X is strongly responsive to D then X must be a descendant of D.

We might be able to recognize when decision $D1$ made before $D2$ can not affect DM's optimal choice for $D2$, and, as a result, some descendants of $D2$ might not be strongly responsive to $D1$ even though they are descendants of $D1$ in the ID. In that case, we could prune $D1'$ from $Pa(D2')$ in the PPD, and the probability distribution in Theorem 2 for the original unpruned PPD would not correspond to an optimal policy. For example, in the ID shown in Figure 4f1, $D1$ is an ancestor of V but V can not be strongly responsive to $D1$ because $D1$ can not affect DM's choice for $D2$.

In the next section we will see how to recognize some observations that can not affect DM's choice for a decision and how to prune them from the PPD.

4. Materiality and Requisite Sets

Materiality arises when the observation of variables $\mathbf{Y}$ instead of variables $\mathbf{X}$, $\mathbf{X} \subset \mathbf{Y}$, before making decision D affects DM's choice and thereby improves the quality of the decision. That is, DM's optimal choice given $\mathbf{Y}$ is sometimes different from DM's optimal choice given $\mathbf{X}$. Otherwise, we can simplify DM's decision problem by considering only $\mathbf{X}$ instead of $\mathbf{Y}$.

More formally, given an ID with latest decision D (or any policy D' in a PPD), we use the corresponding fully specified DPD to compare sets of observations $\mathbf{X}$ and $\mathbf{Y}$, $\mathbf{X} \subset \mathbf{Y} \subseteq Pa(D^*)$, before DM makes the choice for D. $\mathbf{Y}$ is *material for D relative to* $\mathbf{X}$ if $\mathbf{Y}$ is not independent of the corresponding optimal policy D^* given $\mathbf{X}$ in the fully specified DPD. Put more simply, this recognizes whether observing $\mathbf{Y}$ instead of $\mathbf{X}$ *will* change any of DM's optimal decisions. When it will *not*, variables $(\mathbf{Y} \setminus \mathbf{X})$ have been called *numerically redundant* (Faguiouli and Zaffalon, 1998).

Requisite sets are a graphical property in the ID, PPD, or DPD closely related to materiality (Shachter, 1988, 1998, 1999; Nielsen and Jensen, 1999; Luque and Diez, 2010; Nilsson and Lauritzen, 2013). They are minimal subsets $\mathbf{X}$ of all observations $\mathbf{Y}$ available at the time of decision D that are sufficient for DM to make the optimal choice for D. We can simplify DM's decision problem by considering only $\mathbf{X}$ instead of $\mathbf{Y}$. This not only reduces the dimensions of the optimal policy but can also change which variables might be strongly responsive to and material for earlier decisions. For example, in the ID shown in Figure 4b1, $D1$, X, and S will be observed before $D2$ but S is requisite for $D2$ and, as a result, V is not strongly responsive to $D1$.

More formally, given an ID with latest decision D (or any policy D' in a PPD) and all observed variables $\mathbf{Y}$ available to DM before before making the choice for D, $\mathbf{X} \subseteq \mathbf{Y}$ is *requisite* for D if $\mathbf{X}$ is a minimal set such that $\mathbf{X} \cup D$ D-separates $\mathbf{Y}$ from the value node descendants of D. The value node descendants of D are DM's decision criteria that might be strongly responsive to the choice for D. As a result, when $\mathbf{X}$ is requisite for D, observing $\mathbf{Y}$ instead of $\mathbf{X}$ *could not possibly* change any of DM's choices for D given $\mathbf{X}$ for any probability distribution consistent with the ID. In that case, variables $(\mathbf{Y} \setminus \mathbf{X})$ have been called *structurally redundant* (Faguiouli and Zaffalon, 1998).

The following theorem and corollary capture the connection between materiality and requisite sets, extending earlier results (Shachter, 1988, 1999; Nielsen and Jensen, 1999).

Theorem 4 (Materiality and Requisite Sets in the PPD) *Given any decision D in an ID, the corresponding PPD and DPD with policies D' and D^* corresponding to D, and any $\mathbf{X}$ requisite for D' where $\mathbf{X} \subset \mathbf{Y} \subseteq Pa(D') = Pa(D^*)$, then $\mathbf{Y}$ can not be material for D relative to $\mathbf{X}$.*

Proof Because $\mathbf{X}$ is requisite for D', every value node descendant of D' is independent of $\mathbf{Y}$ given $\mathbf{X} \cup D'$, and any other value nodes to be summed are independent of D' by Theorem 2. Thus, given any instance $\mathbf{y}$ of $\mathbf{Y}$ in which the elements corresponding to $\mathbf{X}$ take on the instance $\mathbf{x}$, the optimal policy D^* given $\mathbf{x}$ must also be optimal given $\mathbf{y}$. Therefore, D^* must be independent of $\mathbf{Y}$ given $\mathbf{X}$, and $\mathbf{Y}$ can not be material for D relative to $\mathbf{X}$. ■

Corollary 5 (Materiality and Requisite Sets in the ID) *Given the latest decision D in an ID, the corresponding DPD with policy D^* corresponding to decision D, and any $\mathbf{X}$ requisite for D where $\mathbf{X} \subset \mathbf{Y} \subseteq Pa(D^*)$, then $\mathbf{Y}$ can not be material for D relative to $\mathbf{X}$.*

In practice, materiality is often used to consider the expansion of the observation set, rather than contracting it, including some variables that might only be observed with some effort or expense. For example, the value of clairvoyance on an uncertain variable S before decision D will be zero unless the observations with S are material for D relative to the observations without it (Howard, 1966). To implement the DPD test for materiality in that case requires determining an expanded optimal policy for D^* (and subsequent decisions) as a function of $S \cup Pa(D^*)$, assuming that they would actually be observed. On the other hand, the PPD test for requisite set $\mathbf{X}$ is simply performed on the graph. For example, in the ID shown in Figure 1, if we could observe S before D then S would be requisite for D instead of X. The theorem above shows how the simpler requisite test on the graph might recognize when materiality is impossible for an expanded $Pa(D^*)$ in the ID or PPD, without determining any policies or even specifying a distribution (Poh and Horvitz, 1996).

To determine the *requisite PPD*, the PPD with requisite sets, requires a single backward pass through the PPD, as described by the following theorem and algorithm.

Theorem 6 (Requisite PPD) *Given any two decision nodes $D1$ and $D2$ in an ID, where $D1$ precedes $D2$, and the corresponding PPD with policies $D1'$ and $D2'$ corresponding to $D1$ and $D2$, eliminating non-requisite nodes for $D1$ from $Pa(D1')$ does not change the requisite set for $D2'$.*

Proof The issue is whether removing the arc from some X to $D1'$ can change the requisite set for $D2'$. However, both X and $D1'$ were originally in $Pa(D2')$ by no-forgetting, so the presence of an arc between them can not affect whether they are included in the requisite set for $D2'$. ■

Algorithm 7 (Requisite PPD) *To obtain the requisite PPD from any PPD, consider each policy node D' in reverse order, find a requisite $\mathbf{X}$ for D', a minimal $\mathbf{X} \subseteq Pa(D')$ such that $\mathbf{X} \cup D'$ D-separates $Pa(D')$ from the value node descendants of D', and then set $Pa(D') \leftarrow \mathbf{X}$ (with a corresponding change to $Pa(D*)$ to obtain the requisite DPD).*

The minimal requisite set for each decision or policy node is usually unique and can be found efficiently (Geiger et al., 1990b; Shachter, 1998, 1999; Nielsen and Jensen, 1999). However, when there are deterministic nodes, there might be multiple minimal requisite sets. For example, consider the ID shown in Figure 3 where $D1$ is observed at the time of $D2$. Although X by itself is requisite for $D2$, $D1$ by itself is also requisite for $D2$ because X is functionally determined by $D1$. In this case, the diagram in 3b is preferred to the one in 3c because the alternative set for $D1$ is never smaller than the set of states of X it maps into, and therefore it could generate a larger policy space.

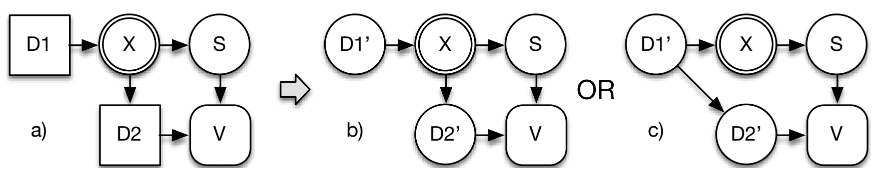

Figure 3: Requisite sets might not be unique when there are deterministic nodes in the ID

The requisite PPD now allows us to refine our test for strong responsiveness, building on Theorems 2 and 6. For example, in the ID shown in Figure 4c, $D1$ is not requisite for $D2$ and thus V can not be strongly responsive to $D1$.

Corollary 8 (Strong Responsiveness and Requisite PPD) *Given any decision D in an ID and corresponding requisite PPD graph with policy D' corresponding to D, if any uncertain or value variable X is strongly responsive to D then X must be a descendant of D' in the requisite PPD.*

5. Independence in the Policy Diagrams

Once the requisite PPD and/or fully specified DPD have been constructed for a given ID, we have seen how we can use established concepts and techniques to confirm any independence statements among the variables. Usually this check takes the form of D-separation in the requisite PPD graph or actual verification of independence in the fully specified DPD, and can be used to check for possible strong responsiveness, requisite sets, and materiality.

We can also use the PPD and DPD to answer queries from three different perspectives.

The first perspective explores the prospective implications of DM's optimal policy. For example, DM might want to understand the effect of the optimal policy on arbitrary variables, such as separate value variables when there are multiple value nodes. In general, the optimal policies are deterministic functions of the requisite sets of variables, some of which can be uncertain, so the

Table 1: Possible Strong Responsiveness and Requisite Sets for the Example IDs

Fig.	*Decn.*	*Poss. Strong Responsiveness*	*Requisite Set*	*Decn.*	*Poss. Strong Responsiveness*	*Requisite Set*
1	D	V	X			
3	$D1$	$\{X, S, D2, V\}$	$\emptyset$	$D2$	V	$D1$ or X
4a	$D1$	$\{X, D2, V\}$	$\emptyset$	$D2$	V	$\{D1, X\}$
4b	$D1$	X	$\emptyset$	$D2$	V	S
4c	$D1$	X	$\emptyset$	$D2$	V	$\emptyset$
4d	$D1$	$\{X, S, D2, V\}$	$\emptyset$	$D2$	V	$D1$
4e	$D1$	$\{X, S, D2, V\}$	$\emptyset$	$D2$	V	X
4f	$D1$	X	$\emptyset$	$D2$	V	$\emptyset$
5	$D1$	all but $\{S1, D1\}$	$S1$	$D3$	$\{V3, S4\}$	$S3$
	$D2$	$\{V2, S3, D3, V3, S4\}$	$S2$			
6	$D1$	all but $\{A, B, C, D1\}$	B	$D3$	$\{V4, K, V3\}$	F
	$D2$	$\{I, D4, L, V2\}$	E	$D4$	$\{L, V2\}$	$\{G, D2\}$

optimal choices that will be made following those policies can therefore be uncertain, too (Nilsson and Jensen, 1998). For example, consider the ID shown in Figure 1a. DM might not know in advance which choice DM will make for D as it could depend on the realization of X. If DM could also observe S before making decision D then S would be requisite for D and D would depend instead on the realization of S. At any point in time DM can think about the probability of future decisions and their requisite observations using either the DPD at the numerical level or the PPD at the graphical level.

The second perspective assumes DM has designed and fielded an agent and can observe some of its actions. For example, if DM observed that the agent took a particular action, DM could infer something about its earlier actions and observations, as well as its actions and observations going forward.

The third perspective models the actions of an independent third party, another person or an agent that DM did not design. If we assess our beliefs about the other's ID, then we can predict its behavior and response to new observations as in the case above of the agent DM designed, albeit with considerably more uncertainty. This is the situation when we model a multi-player game (Koller and Milch, 2003) and it allows us to predict the effects of our actions on another player's behavior, or perhaps better understand which observations might be requisite for that player's response.

6. Examples

In this section we consider a number of example IDs to see which variables are independent of the decisions and which are not, in the sense that we have developed those concepts in earlier sections. The variables that might be strongly responsive to the decisions and the requisite sets for the decisions are summarized in Table 1.

Consider the ID shown in Figure 1, along with its corresponding requisite PPD and DPD. Note that D is the latest decision, V is its only descendant, and X is not D-separated from V by D.

Therefore, V might be strongly responsive to D, S can not be strongly responsive to D, and X is requisite for D, leading to an optimal policy $D*$ as a function of X. Based on the PPD, D' might depend on S but S is independent of D' given X. If S were to be observed before D, it would be requisite for D instead of X, the optimal policy $D*$ would be a function of S instead of X, and $\{S, X\}$ could not be material for D relative to S.

Consider the ID shown in Figure 3, along with its two possible requisite PPDs. $D2$ is the latest decision and V is its only descendant. There is an earlier decision $D1$ and it has descendants X, $D2$, S, and V in both PPDs. Therefore, V might be strongly responsive to $D2$, and X, S, $D2$, and V might be strongly responsive to $D1$. $D1$ and X will be observed before the choice for $D2$, but each is D-separated from V by $D2$ and the other, so $D1$ and X are each requisite for $D2$ but $\{D1, X\}$ is not. Therefore, $\{D1, X\}$ is not material for $D2$ relative to either X or $D1$. S is independent of $D2'$ given either $D1'$ or X in either PPD.

Consider the IDs shown in Figure 4, along with their requisite PPDs. In all of them, $D2$ is the latest decision and V is its only descendant, so V might be strongly responsive to $D2$. In the PPD shown in 4a2, neither $D1'$ nor X can be D-separated from V by either $D2'$, $\{D2', D1'\}$ or $\{D2', X\}$, so $\{D1, X\}$ is requisite for $D2$, and X, $D2$, and V might be strongly responsive to $D1$. In the PPD shown in 4b2, both $D1'$ and X are D-separated from V by $\{D2', S\}$, so S is requisite for $D2$, $\{D1, X, S\}$ is not material for $D2$ relative to S, and only X could be strongly responsive to $D1$. In the PPD shown in 4c2, $D1'$ is D-separated from V by $D2'$, so $\emptyset$ is requisite for $D2$, and $D1$ is not material for $D2$ relative to $\emptyset$. If X were to be observed before $D2$, we would obtain the ID in 4a1. For the ID shown in 4d1, if X were to be observed before $D2$ we would obtain the ID in 4e1, and X instead of $D1$ would be requisite for $D2$. For the IDs shown in 4a1, 4c1 or 4f1, if S were to be observed before $D1$ or $D2$, we would have a PPD similar to 4b2, so $\{D1, X, S\}$ could not be material for $D2$ relative to S in any of those IDs. In the PPDs shown in Figure 4, S and $D2'$ are independent in 4c2 or 4f2, and they are independent given $D1'$ in 4c2, 4d2, and 4f2.

In the Markov Decision Process example (Tatman and Shachter, 1990) ID and corresponding requisite PPD shown in Figure 5, the no-forgetting arcs implicit in the ID, such as $(S1, D3)$ and $(D1, D3)$, are coincidentally eliminated in the requisite PPD. $\{S1, D1', S2, D2'\}$ is D-Separated

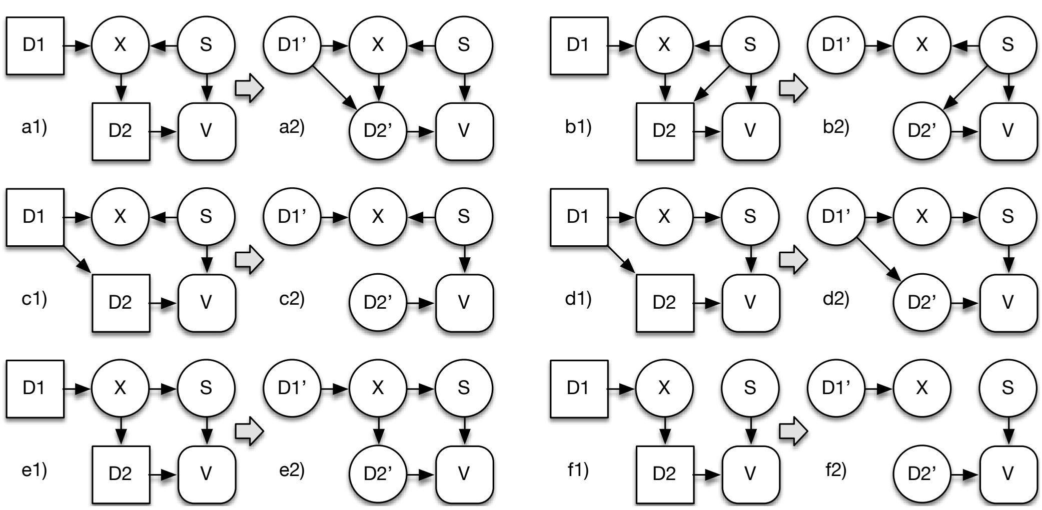

Figure 4: Influence diagram variations and their corresponding requisite PPDs

from $V3$ by $\{D3', S3\}$, so $S3$ is requisite for $D3$, and $\{S1, D1, S2, D2, S3\}$ can not be material for $D3$ relative to $S3$. Similarly, $\{S1, D1'\}$ is D-Separated from $V3$ by $\{D2', S2\}$, so $S2$ is requisite for $D2$. $\{S1, D1', V1\}$ must be independent of $\{D3', V3, S4\}$ given either $S2$ or $S3$.

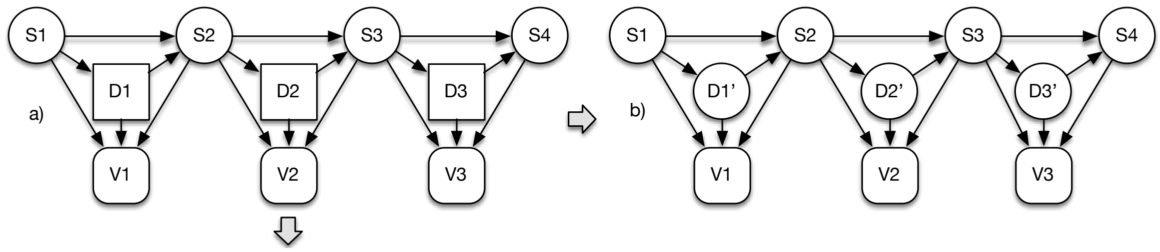

Figure 5: Markov Decision Process ID and its corresponding requisite PPD

In the ID example (Jensen et al., 1994) and corresponding requisite PPD shown in Figure 6 there are four decisions, $D1$, $D2$, $D3$, and $D4$, and four value nodes, $V1$, $V2$, $V3$, and $V4$. $\{D4', G, D2'\}$ D-separates $V2$ from $\{B, D1', E, F, D3'\}$ so $\{G, D2\}$ is requisite for $D4$. If I were also observed before $D4$ then it would be requisite for $D4$, and $\{B, D1, E, F, G, D2, D3, I\}$ could not be material for $D4$ relative to I. F is requisite for $D3$ but if H were also observed before $D3$ then H would be requisite for $D3$. E is requisite for $D2$ but if G were also observed before $D2$ then G would be requisite for $D2$. B is requisite for $D1$ but if $\{A, B, C\}$ were observed before $D1$ then $\{B, C\}$ would be requisite for $D1$. $D3'$ and $\{D2', D4'\}$ might be dependent, but they would be independent given either $\{C, D\}$, F, or E.

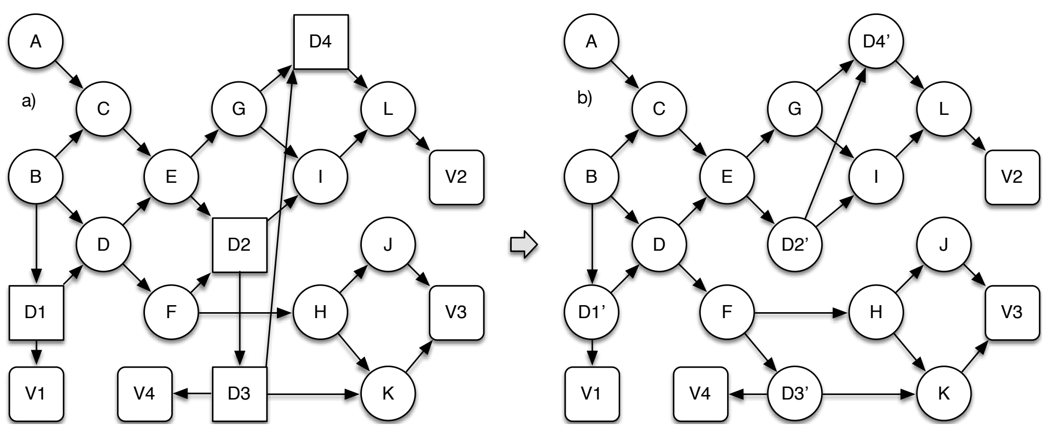

Figure 6: Jensen et al. (1994) ID example and its corresponding requisite PPD

7. Conclusions

Independence among uncertain variables in a Bayesian belief network is well understood, and we have extended that understanding to include decision variables in influence diagrams. Our focus has been on three types of relationships: variables affected by, or *strongly responsive* to, the choice for a decision; variables observed before a decision is made that affect the optimal choice, or are *material* for the decision relative to a subset of those observations; and arbitrary independence relationships in a *policy diagram*, a belief network obtained by substituting optimal *policy variables* for the decisions in the influence diagram.

In the original influence diagram we can only recognize which variables might be strongly responsive to the latest decision and a minimal, or *requisite*, set of observations sufficient to determine the optimal choice for the latest decision.

From the influence diagram we construct a probabilisitic policy diagram with probabilistic policy nodes replacing decision nodes. Its graph allows us to recognize which variables might be strongly responsive to each of the decisions, requisite sets for each of the decisions, and independence relations using D-separation. We use the probabilisitic policy diagram to recognize as much independence as possible at the graphical level.

Finally, we construct a fully specified deterministic policy diagram, including all states and probabilities, and optimal policies determined as functions of the requisite sets of variables for each decision. In that deterministic policy diagram we can test for materiality and other independence queries at the numerical level.

This work could be extended in a number of ways. In particular, if we did not assume no-forgetting then we could apply limited information IDs (Lauritzen and Nilsson, 2001; Nilsson and Lauritzen, 2013) or if we assumed some monotonicity of the value distributions we might be able to recognize even smaller requisite sets (Luque and Diez, 2010).

Acknowledgments

We thank Gordon Hazen for posing the question this paper addresses and the anonymous reviewers for their helpful suggestions.

References

D. Bhattacharjya and R. D. Shachter. Evaluating influence diagrams with decision circuits. In *Proceedings of the 23rd Conference on Uncertainty in AI*, pages 9–16, 2007.

G. Cooper. A method for using belief networks as influence diagrams. In *Proceedings of the 4th Conference on Uncertainty in AI*, pages 55 – 63, 1988.

E. Faguiouli and M. Zaffalon. A note about redundancy in influence diagrams. *International Journal of Approximate Reasoning*, 19:231–246, 1998.

D. Geiger, T. Verma, and J. Pearl. Identifying independence in Bayesian networks. *Networks*, 20 (5):507–534, 1990a.

D. Geiger, T. Verma, and J. Pearl. d-separation: From theorems to algorithms. In *Proceedings of the 5th Conference on Uncertainty in AI*, pages 139–148, 1990b.

D. Heckerman and R. D. Shachter. Decision-Theoretic Foundations for Causal Reasoning. *Journal of Artificial Intelligence Research*, 3:405–430, 1995.

R. A. Howard. Information Value Theory. *IEEE Transactions on Systems Science and Cybernetics*, 2(1):22–26, 1966.

F. Jensen, F. V. Jensen, and S. L. Dittmer. From Influence Diagrams to Junction Trees. In *Proceedings of the 10th Conference on Uncertainty in AI*, pages 367–373. Morgan Kaufmann, 1994.

D. Koller and B. Milch. Multi-Agent Influence Diagrams for Representing and Solving Games. *Games and Economic Behavior*, 45(1):181–221, 2003.

S. L. Lauritzen and D. Nilsson. Representing and Solving Decision Problems with Limited Information. *Management Science*, 47(9):1235–1251, 2001.

M. Luque and F. J. Diez. Variable elimination for influence diagrams with super value nodes. *International Journal of Approximate Reasoning*, 51(6):615–631, 2010.

T. D. Nielsen and F. V. Jensen. Well-Defined Decision Scenarios. In K. B. Laskey and H. Prade, editors, *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, pages 502–511. Morgan Kaufmann, 1999.

D. Nilsson and F. V. Jensen. Probabilities of Future Decisions. In *IPMU*, pages 161–171. Kluwer Academic Publishers, 1998.

D. Nilsson and S. L. Lauritzen. Evaluating Influence Diagrams using LIMIDs. In *Proceedings of the 14th Conference on Uncertainty in AI*, 2013.

J. Pearl. *Probabilistic Reasoning in Intelligent Systems*. Elsevier, 1988.

K. L. Poh and E. Horvitz. A Graph-Theoretic Analysis of Information Value. In *Proceedings of the 12th Conference on Uncertainty in AI*, pages 427–435. Morgan Kaufmann, 1996.

R. D. Shachter. Evaluating Influence Diagrams. *Operations Research*, 34(6):871–882, 1986.

R. D. Shachter. Probabilistic Inference and Influence Diagrams. *Operations Research*, 36(4):589–604, 1988.

R. D. Shachter. Bayes-ball: Rational pastime (for determining irrelevance and requisite information in belief networks and influence diagrams). *Proceedings of the 14th Conference on Uncertainty in AI*, pages 480–487, 1998.

R. D. Shachter. Efficient value of information computation. In *Proceedings of the 15th Conference on Uncertainty in AI*, pages 594–601, 1999.

R. D. Shachter and C. R. Kenley. Gaussian Influence Diagrams. *Management Science*, 35(May): 527–550, 1989.

R. D. Shachter and M. A. Peot. Decision Making Using Probabilistic Inference Methods. In *Proceedings of the 9th Conference on Uncertainty in AI*, pages 383–390, 1993.

P. Shenoy and G. Shafer. Axioms for probability and belief-function propagation. In *Proceedings of the 4th Conference on Uncertainty in AI*, pages 169–198, 1988.

J. Tatman and R. D. Shachter. Dynamic programming and influence diagrams. *IEEE Transactions on Systems, Man and Cybernetics*, 20(2):365–379, 1990.