

Causal Models, Value of Intervention, and Search for Opportunities

Tsai-Ching Lu and Marek J. Druzdzel

Decision Systems Laboratory
School of Information Sciences and Intelligent Systems Program
University of Pittsburgh, Pittsburgh, PA 15260
ching,marek@sis.pitt.edu

Summary. While algorithms for influence diagrams allow for computing the optimal setting for decision variables, they offer no guidance in generation of decision variables, arguably a critical stage of decision making. A decision maker confronted with a complex system may not know which variables to best manipulate to achieve a desired objective. We introduce the concept of *search for opportunities* which amounts to identifying the set of decision variables and computing their optimal settings, given an objective expressed by a utility function. Search for opportunities is built on value of intervention in causal models.

Key words: causal models, value of intervention, search for opportunities, causal Bayesian networks, and structural equation models.

1 Introduction

Influence diagrams [4] are popular tools for representing decision problems under uncertainty and identifying optimal strategies. The key problem with using influence diagrams is that we need to specify beforehand all decision alternatives and their consequences explicitly. In complex systems, this may result in a cumbersome, if not a totally unmanageable, modeling process. Ideally, a modeling language should support the prediction of the effects of actions that were not considered in model's construction [1, 8, 9]. This allows us to search for the best actions to be taken to achieve a set of objectives, a concept that we refer to as *search for opportunities*.

The problem of search for opportunities is related to the problem of *information gathering* [10]. In information gathering, a decision maker tries to decide which information to acquire to reduce the uncertainty over a model, and consequently to improve the quality of the decision at hand. The means of acquiring information is constrained to observations that modify the decision maker's belief over the states of a system. Search for opportunities, in

contrast, seeks to apply intervening actions that alter the trajectory of the system toward those outcomes that are preferable to the decision maker.

In decision analysis, the primary tool for information gathering is *value of information* [3]. Value of information is defined as the upper bound on what a decision maker should be willing to pay in employing a clairvoyant to reveal the outcome of a chance variable. Similarly, the concept of *value of control* has been introduced and defined as the upper bound on what a decision maker should be willing to pay a wizard for setting a chance variable into a preferred state. Both value of information and value of control are defined with respect to a decision problem [6]. However, to our knowledge, value of control computation has only been applied to chance nodes with no predecessors in influence diagrams encoded in Howard canonical form [6]. Since influence diagrams may describe probabilistic rather than causal relations, there is no guarantee that converting a chance node with predecessors or a chance node without predecessors in diagrams that are not in Howard canonical form into a decision node will correctly model the effects of control.

In this chapter, we discuss the problem of search for opportunities, where a decision maker seeks for creative decision options in order to achieve a given objective. The basis of the search is a causal model of the system that is subject to the decision. Causal models based on structural equations support causal reasoning and, in particular, prediction of the effects of actions [9, 12]. A causal model consists of a self-contained set of simultaneous structural equations, each of which represents a causal mechanism active in the modeled system. Causal models support prediction of the effects of actions by replacing those mechanisms that are impacted by actions with new mechanisms, possibly not contemplated during model's construction, and leaving the rest intact. The problem of search for opportunities, in this formulation, amounts to searching for variables that were not originally contemplated as decision variables and intervening into mechanisms governing these variables in order to affect the outcomes. Therefore, search for opportunities leads to discovery of *novel* actions that help to achieve decision objectives.

To address the problem of search for opportunities, we introduce the concept of *value of intervention*. The value of intervention, related to the value of control, arises from considering jointly the economic factors and effects of actions in causal models. It can be considered as a generalization of the value of control since the intervention operates at the level of mechanism in causal models, but the control operates at the level of variables in influence diagrams. The value of intervention computation is also applicable to influence diagrams in canonical form [2], an extension of Howard Canonical Form that supports causal reasoning, but not to influence diagrams that do not represent causal relations.

The remainder of this chapter is organized as follows. Section 2 gives an overview of causal models. Section 3 introduces augmented models for describing decision problems at hand. Section 4 discusses the concept of value of intervention. Section 5 shows how value of intervention can be used to solve

the problem of search for opportunities. Section 6 describes the modeling of non-intervening actions. Finally, we discuss the direction of our future work.

2 Causal Models

A causal model consists of variables and causal relations among them, modeled by a set of simultaneous *structural equations*, where each equation represents a distinct mechanism active in the system.¹ More formally, we denote the set of variables appearing in an equation e as $\text{Vars}(e)$, and in a set of equations \mathbf{E} as $\text{Vars}(\mathbf{E}) \equiv \cup_{e \in \mathbf{E}} \text{Vars}(e)$. A *causal model* $M = \langle \mathbf{X}, \mathbf{E} \rangle$ consists of a *self-contained* set of simultaneous structural equations \mathbf{E} over a set of variables $\mathbf{X} \equiv \text{Vars}(\mathbf{E})$. Each structural equation $e \in \mathbf{E}$ is generally written in its implicit form $e(X_1, X_2, \dots, X_n) = 0$ where $X_i \in \text{Vars}(e)$. We say that a variable $X_i \in \mathbf{X}$ is *exogenous* to M if its value is determined by factors outside the system, i.e., if there exists a structural equation $e \in \mathbf{E}$, $e(X_i) = 0$, and *endogenous* otherwise. In other words, the set of variables \mathbf{X} consists of two disjoint sets \mathbf{U} and \mathbf{V} of exogenous and endogenous variables respectively. Therefore, a causal model is sometimes denoted as $M = \langle \mathbf{U}, \mathbf{V}, \mathbf{E} \rangle$. Let $D(X_i)$ be the domain of a variable X_i , and $D(\mathbf{X}) = D(X_1) \times \dots \times D(X_n)$ be the domain of the set of variables $\mathbf{X} = \{X_1, \dots, X_n\}$. Given $\mathbf{u} \in D(\mathbf{U})$, the solutions for endogenous variables $\mathbf{Y} \subseteq \mathbf{V}$, denoted as $\mathbf{Y}_M(\mathbf{u})$ or $\mathbf{Y}(\mathbf{u})$, in a causal model M can always be determined uniquely. The pair $\langle M, \mathbf{u} \rangle$ is called a *causal world*, or simply world. Given a probability distribution $\text{Pr}(\mathbf{u})$ defined over $D(\mathbf{U})$, the pair $\langle M, \text{Pr}(\mathbf{u}) \rangle$ is called a *probabilistic* causal model where for each $Y \in \mathbf{V}$, $\text{Pr}(Y = y) \triangleq \sum_{\{\mathbf{u} | Y(\mathbf{u}) = y\}} \text{Pr}(\mathbf{u})$. Simon [11] developed an algorithm that explicates the asymmetries among variables in a causal model M and represents them as a *causal graph* $G(M)$. A causal model M is *recursive* if the associated $G(M)$ is a directed acyclic graph, where each node corresponds to a variable, and each family (a node with its parents in $G(M)$) a structural equation [1]. In other words, each structural equation $e(X_1, \dots, X_n) = 0$ is expressed in its explicit functional form $X_i = f_{X_i}(X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_n)$ and is depicted graphically as a family with arcs from nodes representing arguments of f_{X_i} (i.e., $X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_n$) to X_i . In the sequel, the term “causal model” refers to a recursive causal model, in which each equation is indexed by the dependent variable of its explicit form.

Example 1. Consider a model for the operational status of a command center (CC). CC depends on the status of communications (C) and radar (R). Radar depends on the antenna structure (A) and the power supplied by the generator (G). Communications rely on the power supplied by the generator. The generator relies on fuel supply (F) to generate power. Each of the variables has state *operational* or *damaged*. We assume that for each of the above relations there is a corresponding exogenous variable, denoted as U_{cc} ,

¹ Please see [9] or [12] for general introductions to causal models.

U_r , U_c , and U_g , that for each relation summarizes those factors that are outside the model. F and A are themselves exogenous variables. We assume that all exogenous variables $U = \{U_{cc}, U_r, U_c, U_g, F, A\}$ are independent. The set of structural equations, representing the domain of our interest, and its corresponding causal graph are shown in Fig. 1. We have included an explicit graphical representation of variables U_{cc}, U_r, U_c, U_g for the sake of clarity of explanation. In practice, these variables are modeled implicitly by error terms in the corresponding equations and we will omit them in the sequel of this chapter for the sake of clarity.

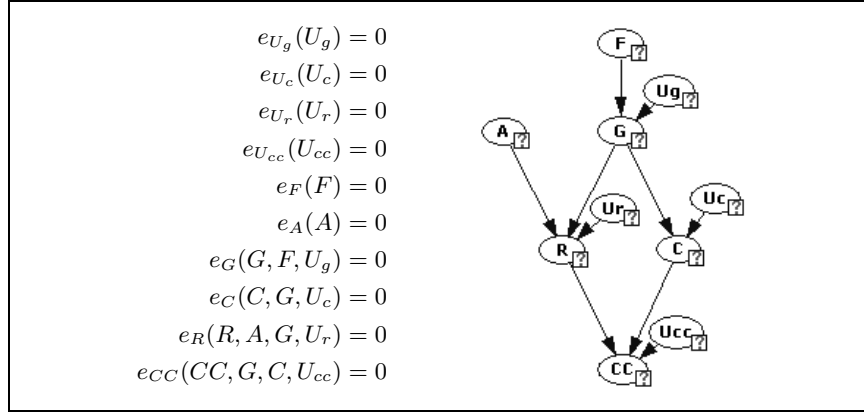


Fig. 1. Causal model and its corresponding causal graph for modeling the operational status of a command center.

2.1 Actions and Effects

Given a causal model $M = \langle \mathbf{U}, \mathbf{V}, \mathbf{E} \rangle$, an *atomic action*, $do(X = x)$, sets an endogenous variable $X \in \mathbf{V}$ to the value $x \in D(X)$ and transforms M to the modified model $M_x = \langle \mathbf{U}, \mathbf{V}, \mathbf{E}_x \rangle$, where $\mathbf{E}_x = \{e_Y | Y \in \mathbf{V} \setminus \{X\}\} \cup \{X = x\}$. The effect of action $do(X = x)$ is given by M_x . The *action operator* $do(\cdot)$ replaces e_X with its argument, a (probability) function that (probabilistically) assigns X with the value $x \in D(X)$, to derive the modified model M_x and consequently its corresponding effects. In general, the intervention in the argument of $do(\cdot)$ operator may result in non-recursive models. For example, when applying a $do(e'_X(X, Y) = 0)$ to a recursive causal model $M = \langle \{X, Y\}, \{e_X(X) = 0, e_Y(Y, X) = 0\} \rangle$, we have a non-recursive causal model $\langle \{X, Y\}, \{e'_X(X, Y) = 0, e_Y(Y, X) = 0\} \rangle$. However, in this chapter we restrict our analysis to actions that result in recursive models. Pearl and Robins [9] formalized three types of actions in this class: *Conditional action*, denoted as $do(X = x \downarrow_{x=g(\mathbf{z})})$, sets $X \in \mathbf{V}$ to the value $x = g(\mathbf{z})$ whenever

\mathbf{Z} attain values \mathbf{z} , where $g : \mathbf{D}(\mathbf{Z}) \rightarrow \mathbf{D}(X)$ and \mathbf{Z} are non-descendants of X in $G(M)$. *Stochastic action*, denoted as $do(X = x \downarrow_{\text{Pr}^*(x)})$, sets $X \in \mathbf{V}$ to the value x with probability $\text{Pr}^*(x)$ where $\text{Pr}^*(x)$ is specified externally. We may also have *stochastic policy*, denoted as $do(X = x \downarrow_{\text{Pr}^*(x|\mathbf{z})})$, that sets $X = x$ with probability $\text{Pr}^*(x|\mathbf{z})$ whenever \mathbf{Z} attain values \mathbf{z} where \mathbf{Z} are non-descendants of X in $G(M)$ and $\text{Pr}^*(x|\mathbf{z})$ is set externally. In the sequel, we follow Pearl's notation [9] where \hat{x} is the abbreviation of X attaining value $x \in \mathbf{D}(X)$ due to an intervention.

Given a world $\langle M, u \rangle$, the *potential response* of $Y \in \mathbf{V}$ to action $do(\cdot)$ on variable $X \in \mathbf{V}$, denoted as $Y_{M_x}(\mathbf{u})$ or $Y_x(\mathbf{u})$, is the solution for Y of the set of equations \mathbf{E}_x of M_x . $Y_x(\mathbf{u})$ can also be interpreted as the *counterfactual* value that Y would obtain had X been x in the counterfactual world brought about by action $do(\cdot)$. Given a probabilistic causal model $\langle M, \text{Pr}(\mathbf{u}) \rangle$, the causal effect on Y of an atomic action $do(X = x)$ is given by $\langle M_x, \text{Pr}(\mathbf{u}) \rangle$ as $\text{Pr}(y|\hat{x}) \equiv \text{Pr}(Y_x = y) \triangleq \sum_{\{\mathbf{u} | Y_x(\mathbf{u}) = y\}} \text{Pr}(\mathbf{u})$. The causal effect on Y of a conditional action $do(X = x \downarrow_{x=g(\mathbf{z})})$ is expressed as $\text{Pr}(y|\hat{x}) \downarrow_{x=g(\mathbf{z})} \equiv \text{Pr}(Y_x = y) \downarrow_{x=g(\mathbf{z})} \triangleq \sum_{\mathbf{z}} \text{Pr}(y|\hat{x}, \mathbf{z}) \downarrow_{x=g(\mathbf{z})} \text{Pr}(\mathbf{z})$. The causal effect on Y of a stochastic action $do(X = x \downarrow_{\text{Pr}^*(x)})$ is expressed as $\text{Pr}(y|\hat{x}) \downarrow_{\text{Pr}^*(x)} \equiv \text{Pr}(Y_x = y) \downarrow_{\text{Pr}^*(x)} \triangleq \sum_x \text{Pr}(y|\hat{x}) \text{Pr}^*(x)$. The causal effect on Y of a stochastic policy $do(X = x \downarrow_{\text{Pr}^*(x|\mathbf{z})})$ is expressed as $\text{Pr}(y|\hat{x}) \downarrow_{\text{Pr}^*(x|\mathbf{z})} \equiv \text{Pr}(Y_x = y) \downarrow_{\text{Pr}^*(x|\mathbf{z})} \triangleq \sum_x \sum_{\mathbf{z}} \text{Pr}(y|\hat{x}, \mathbf{z}) \text{Pr}^*(x|\mathbf{z}) \text{Pr}(\mathbf{z})$.

Example 2. Suppose the model in Example 1 is an enemy's command center and one objective is to disrupt the enemy's communications. We can act on the communications C by, for example, jamming the signal with noise, and setting C to *damaged*. The modified causal model and its corresponding causal graph are shown in Fig. 2. Please note that the intervention makes the arcs coming into C inactive (arc $G \rightarrow C$).

3 Augmented Models

In order to describe a decision problem at hand, we propose to augment a probabilistic causal model $\langle M, \text{Pr}(\mathbf{u}) \rangle$ by characterizing its variables along three properties: *observability* (observable or unobservable), *manipulability* (manipulable or non-manipulable), and *focus* (focus or non-focus). Let \mathbf{X} denote all variables in M . A variable $X_i \in \mathbf{X}$ is *observable*, denoted as $X_i.o$, if it represents an entity that can be *measured* directly; *unobservable*, denoted as $X_i.\bar{o}$, otherwise. A variable $X_i \in \mathbf{X}$ is *manipulable*, denoted as $X_i.m$, if it represents an entity that can be *manipulated* directly; *non-manipulable*, denoted as $X_i.\bar{m}$, otherwise. We assume that a manipulable variable is always observable, i.e., we assume that we can always observe the effect of our manipulation. A variable $X_i \in \mathbf{X}$ is a *focus* variable, denoted as $X_i.f$, if it represents an objective; *non-focus*, denoted as $X_i.\bar{f}$, otherwise.

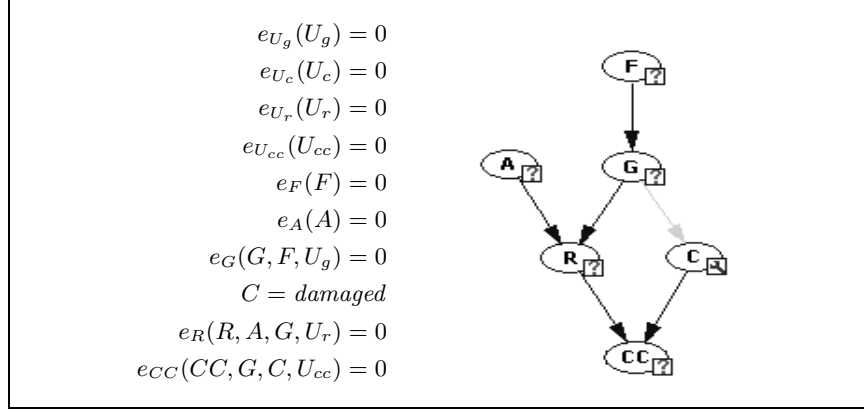


Fig. 2. The modified causal model and its corresponding causal graph after atomic action $do(C = \text{damaged})$.

The goal of this work is to build a system that suggests decisions. At any stage of working with the system, there may be variables on which users have decided to intervene, whether based on the system's suggestions or the user's prior choices. We represent such decisions by augmenting the model with a set of *decision* variables \mathbf{D} along with their corresponding settings. The domain of each $D_i \in \mathbf{D}$ consists of the choices of setting $X_i.m$, the augmented manipulable variable, to a value $x_i \in \mathbf{D}(X_i)$, denoted as x'_i , and a special state *idle* representing the force of nature [9]. Let \mathbf{PA}_i denote the set of parents of X_i in $G(M)$, i.e., $\mathbf{PA}_i = \text{Vars}(e_{X_i}) \setminus \{X_i\}$. We augment the equation $e_{X_i}(X_i, \mathbf{PA}_i) = 0$ to $e'_{X_i}(X_i, \mathbf{PA}'_i) = 0$, where $\mathbf{PA}'_i = \mathbf{PA}_i \cup \mathbf{Z} \cup \{D_i\}$ and $\mathbf{Z} \subset \mathbf{X}$, a set of non-descendants of X_i in $G(M)$ brought about by interventions. We define the augmented equation $e'_{X_i}(X_i, \mathbf{PA}'_i) = 0$ as

$$e'_{X_i}(X_i, \mathbf{PA}'_i) \triangleq \begin{cases} e_{X_i}^*(X_i, \mathbf{PA}_i) = 0 & \text{if } D_i = x'_i \\ e_{X_i}(X_i, \mathbf{PA}_i) = 0 & \text{if } D_i = \text{idle} \end{cases}, \quad (1)$$

where the form of $e_{X_i}^*(X_i, \mathbf{PA}'_i) = 0$ depends on the type of intervention (see Table 1). To represent *concurrent actions* on X_i and X_j , in addition to D_i and D_j and corresponding augmentations on e_{X_i} and e_{X_j} , we add an additional decision variable, denoted as D_{ij} , to represent the concurrency. The domain of D_{ij} is $\mathbf{D}(D_{ij}) = \mathbf{D}(D_i) \times \mathbf{D}(D_j)$. We add additional *projection* equations $e_{D_i} : \mathbf{D}(D_{ij}) \rightarrow \mathbf{D}(D_i)$ and $e_{D_j} : \mathbf{D}(D_{ij}) \rightarrow \mathbf{D}(D_j)$ such that $D_i = d_i$ and $D_j = d_j$ for each $d_{ij} \in \mathbf{D}(D_{ij})$.

Finally, to represent our preferences over the given set of objectives and decisions, we augment the model by a set of *utility* variables \mathbf{UT} along with their utility functions \mathbf{U} . Each utility function $U_i \in \mathbf{U}$ can only have focus or decision variables as its arguments. Formally, we can define an augmented model as follows.

Table 1. The form of $e_{X_i}^*(X_i, \mathbf{PA}'_i)$ characterized with respect to different types of interventions: atomic action, conditional action, stochastic action, and stochastic policy.

Type	$e_{X_i}^*(X_i, \mathbf{PA}'_i) = 0$
Atomic	$\Pr(x_i \mathbf{pa}'_i) = \begin{cases} 1 & \text{if } x'_i = x_i, \\ 0 & \text{otherwise.} \end{cases}$
Conditional	$\Pr(x_i \mathbf{pa}'_i) = \begin{cases} 1 & \text{if } x'_i = x_i \text{ and } x_i = g(\mathbf{z}), \\ 0 & \text{otherwise.} \end{cases}$
Stochastic	$\Pr(x_i \mathbf{pa}'_i) = \begin{cases} \Pr^*(x_i) & \text{if } x'_i = x_i, \\ 0 & \text{otherwise.} \end{cases}$
Policy	$\Pr(x_i \mathbf{pa}'_i) = \begin{cases} \Pr^*(x_i \mathbf{z}) & \text{if } x'_i = x_i \text{ and } x_i = g(\mathbf{z}), \\ 0 & \text{otherwise.} \end{cases}$

Definition 1 (Augmented Model).

An augmented model for a decision problem is

$M_A = \langle \langle M, \Pr(\mathbf{u}) \rangle, C(\mathbf{X}), \langle \mathbf{D}, \mathbf{E}' \rangle, \langle \mathbf{UT}, \mathbf{U} \rangle \rangle$ where:

1. $\langle M, \Pr(\mathbf{u}) \rangle$ is a probabilistic causal model.
2. $C(\mathbf{X})$ is a characterization of observability, manipulability, and focus for each $X_i \in \mathbf{X}$.
3. $\langle \mathbf{D}, \mathbf{E}' \rangle$ is a set of decision variables \mathbf{D} and the modified equation \mathbf{E}' with respect to the decisions.
4. $\langle \mathbf{UT}, \mathbf{U} \rangle$ is a set of utility variables \mathbf{UT} and its corresponding utility functions \mathbf{U} over a set of focus variables, characterized by $C(\mathbf{X})$, and a subset of decision variables in \mathbf{D} .

Example 3. Suppose that variables F , A , G , R , C , and CC are manipulable and CC is the only focus variable in Example 1. We add a utility node, *Utility*, with utility function $U(CC)$, to represent our preference over the states of CC . The corresponding causal graph is shown in Fig. 3 (a). Suppose we have a decision option of manipulating the communications C with no direct influence on *Utility*. We then have a causal graph with D_c as a decision variable shown in Fig. 3 (b).

We emphasize that specifying a variable manipulable, which merely acknowledges the possibility of interventions, is not the same as designating a decision variable in an influence diagram, which requires the explicit specifications of a decision variable along with its consequences. Declaring a variable manipulable allows the algorithm of search for opportunities to explore possible interventions that might not be foreseen when the model is constructed. Furthermore, we require neither a manipulable variable being intervened upon, nor an observable variable being observed. It is the task of search for opportunities and information gathering to determine which variable one should intervene or observe and in what order. In other words, we propose to relax




Fig. 3. (a) Example model augmented with utility over variable CC . (b) The model in (a) augmented with the atomic action D_c .

not only the assumption of a fixed sequence of intervening actions and observations in influence diagrams, but also the assumption of a fixed operation over a variable. For example, a manipulable variable may be intervened, observed, or unknown, depending on different decision sequences generated by search for opportunities and information gathering. Only when a manipulable variable is augmented by the decision variable and its augmented equation as in the form of Equation 1, we commit ourselves to intervene the manipulable variable with one of the policies specified by the decision variable.

In this chapter, we constrain ourselves in applying augmented models on decision scenarios that contain only intervening actions. We treat an augmented model as an influence diagram and compute the optimal strategy and the maximum expected utility using algorithms for inference diagrams.

4 Value of Intervention

Suppose that we are considering an additional atomic intervention on an unaugmented manipulable variable X_k in an augmented model M_A . We augment M_A by adding a new decision variable D_k and modify e_{X_k} to e'_{X_k} as demonstrated in Equation 1. Let $\mathbf{D}' = \mathbf{D} \cup \{D_k\}$ be the new set of decision variables. We also augment utility function \mathbf{U} to \mathbf{U}' if the intervention directly influences \mathbf{U} . Most interventions come at a certain cost and the cost of intervention can be incorporated by augmenting \mathbf{U} . Let M'_A denote the newly augmented model, $\pi(M'_A)$ its optimal strategy, and $\text{MEU}(M_A)$ and $\text{MEU}(M'_A)$ the maximum expected utility yielded by the optimal strategies $\pi(M_A)$ and $\pi(M'_A)$ respectively. We define the *value of intervention* on X_k as

$$\text{VOINT}(D_k = d_k^*) \triangleq \text{MEU}(M'_A) - \text{MEU}(M_A), \quad (2)$$

where $d_k^* \in \mathbf{D}(D_k)$ is yielded in $\pi(M'_A)$ by the optimal policy of D_k . Note that value of atomic intervention can account for the concept of the value of control in influence diagrams. Since augmented models support prediction of the effect of actions, we are not constrained to only atomic intervention (control) on those nodes with no predecessors as in the case of value of control in influence diagrams. To compute value of intervention for a conditional, stochastic, or stochastic policy action, we simply substitute the action of interest for the atomic intervention in augmenting model M_A and perform analysis using Equation 2.

Theorem 1. *Let M_A be an augmented model and X_k be an unaugmented manipulable variable in M_A . If M'_A is the augmented model of M_A that considers an intervention on X_k that has no direct impacts on utility functions \mathbf{U} , then $\text{MEU}(M'_A) \geq \text{MEU}(M_A)$.*

Proof. When evaluating $\pi(M_A)$, we can decompose the joint probability distribution of M_A according to $G(M_A)$. When considering an additional intervention on X_k that has no direct impacts on \mathbf{U} , we augment M_A into M'_A by modifying $\Pr(x_k|\mathbf{pa}_k)$ to $\Pr(x_k|\mathbf{pa}'_k)$. Now, when evaluating $\pi(M'_A)$, we also decompose the joint probability distribution of M'_A according to the $G(M'_A)$. Notice that X_k participates in $\text{MEU}(M_A)$ as

$$\text{MEU}(M_A) = \dots \sum_{x_k \in \mathbf{D}(X_k)} \Pr(x_k|\mathbf{pa}_k) \dots \mathbf{U},$$

and in $\text{MEU}(M'_A)$ as

$$\text{MEU}(M'_A) = \dots \max_{d_k \in \mathbf{D}(D_k)} \sum_{x_k \in \mathbf{D}(X_k)} \Pr(x_k|\mathbf{pa}'_k) \dots \mathbf{U},$$

where $d_k \in \mathbf{pa}'_k$. Since $\Pr(x_k|\mathbf{pa}_k)$ is also represented in $\Pr(x_k|\mathbf{pa}'_k)$ as $d_k = \text{idle}$, $\text{MEU}(M'_A) \geq \text{MEU}(M_A)$ according to the maximization operator.

Consider a simple model M_A which consists of one variable X_k with probability distribution $\Pr(x_k)$, where $x_k \in \mathbf{D}(X_k)$, and a utility variable UT with utility function $\mathbf{U}(x_k)$. We have $\text{MEU}(M_A) = \text{EU}(M_A) = \sum_{x_k \in \mathbf{D}(X_k)} \Pr(x_k)\mathbf{U}(x_k)$. Consider an additional stochastic intervention on X_k that has no direct impact on \mathbf{U} . We augment M_A to M'_A with a new decision variable D_k with domain $\mathbf{D}(D_k) = \mathbf{D}(X_k) \cup \{\text{idle}\}$, and modify the probability distribution of X_k to

$$\Pr(x_k|d_k) = \begin{cases} \Pr^*(x_k) & \text{if } D_k = x'_k \text{ and } x_k = x'_k, \\ 0 & \text{if } D_k = x'_k \text{ and } x_k \neq x'_k, \\ \Pr(x_k) & \text{if } D_k = \text{idle}. \end{cases}$$

We have

$$\text{MEU}(M'_A) = \max_{d_k \in \mathbf{D}(D_k)} \sum_{x_k \in \mathbf{D}(X_k)} \Pr(x_k|d_k)U(x_k)$$

and the optimal value of setting D_k ,

$$d_k^* = \arg \max_{d_k \in \mathbf{D}(D_k)} \sum_{x_k \in \mathbf{D}(X_k)} \Pr(x_k|d_k)U(x_k).$$

It shows that d_k^* is taken on one of x'_k only if

$$\sum_{x_k \in \mathbf{D}(X_k)} \Pr^*(x_k)U(x_k) > \sum_{x_k \in \mathbf{D}(X_k)} \Pr(x_k)U(x_k).$$

In other words, it suggests not to act if the stochastic intervention under consideration does no better than the nature. By the same token, d_k^* can always take on the state *idle* for other types of interventions that have no direct impacts on utility functions U , if it will not do better than the nature.

Next, consider an augmented model with variables $D_A \rightarrow A \rightarrow B \rightarrow U$. We have

$$\text{MEU}(M_A) = \max_{d_A \in \mathbf{D}(D_A)} \sum_{a \in \mathbf{D}(A)} \Pr(a|d_A) \sum_{b \in \mathbf{D}(B)} \Pr(b|a)U(b).$$

Consider an additional intervention on B with no direct impacts on U and augment M_A to M'_A correspondingly. We have

$$\text{MEU}(M'_A) = \max_{d_A \in \mathbf{D}(D_A)} \sum_{a \in \mathbf{D}(A)} \Pr(a|d_A) \max_{d_B \in \mathbf{D}(D_B)} \sum_{b \in \mathbf{D}(B)} \Pr(b|a, d_B)U(b).$$

We see again that D_B will take on the state other than *idle* only if

$$\sum_{b \in \mathbf{D}(B)} \Pr(b|a, d_B)U(b) > \sum_{b \in \mathbf{D}(B)} \Pr(b|a)U(b).$$

5 Search for Opportunities

Search for opportunities refers to the problem of identifying novel interventions that can improve the outcomes. Ideally, we should consider all possible novel interventions on all unaugmented manipulable variables simultaneously, along with existing decisions, to find the optimal strategy and the maximum expected utility for the model. In a complex system, however, such analysis can easily challenge our modeling and computational capabilities. For example, even if we constrain ourselves to considering only atomic interventions on manipulable variables in the command center example, theoretically, we need

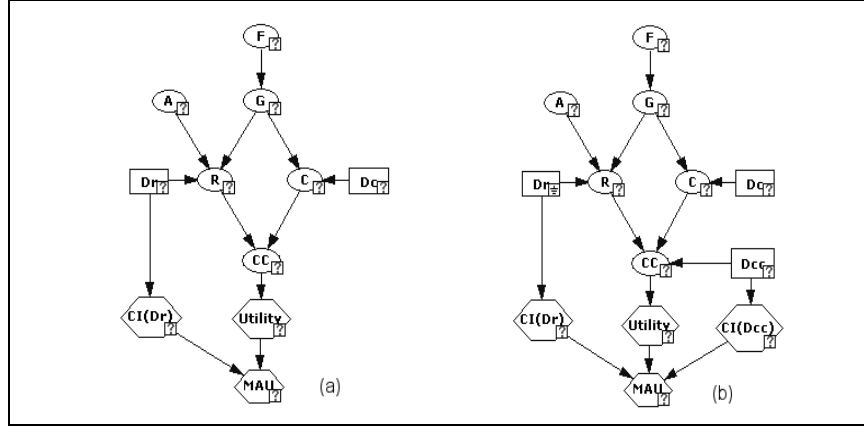


Fig. 4. (a) The model augmented with a possible atomic intervention on R , the cost of intervention CI_{D_R} , and the multi-attribute utility function MAU . (b) The optimal atomic intervention on R is instantiated by setting D_R to d_R^* and the model is augmented with a possible intervention on CC and the cost of intervention $CI_{D_{CC}}$ in the myopic approach of search for opportunities.

to elicit utilities and to evaluate strategies for 3^6 combinations of all possible atomic interventions. In general, if we have n manipulable variables with m states for each, we will have $(m + 1)^n$ combinations of utilities and strategies, including one extra dimension for the force of nature.

To simplify the problem of modeling, we assume that all novel interventions under consideration have no direct impact on utility functions except by the cost of intervention. We elicit the cost of intervention CI_k for each novel intervention on X_k . Let \mathbf{CI} denote all the costs of interventions under consideration. We assume that our multi-attribute utility function over the existing individual utility functions \mathbf{U} and the costs of interventions is *decomposable*, i.e., there exists a multi-attribute utility function MAU that takes as arguments each $U_i \in \mathbf{U}$ and each $CI_k \in \mathbf{CI}$ and combine them with a functional form (such as a simple linear or a multiplicative form). In the case of a linearly additive MAU , we need to elicit $n \times (m + 1)$ numbers for \mathbf{CI} and at most $n + 1$ numbers for the weights in MAU (if the units of cost are the same, this number can be significantly smaller). Next, we approximate the optimal strategy computation by the *myopic* (greedy) search that considers one intervention at a time and selects the one with the maximum value of intervention to perform. We act according to the selected intervention and perform the myopic search again to select the next intervention to act until there is no intervention that can improve the maximum expected utility over our predefined threshold. Fig. 5 outlines this procedure.

Example 4. Given the model in Fig. 3 (b), suppose that we use the myopic approach to identify that R is the next variable to intervene by an atomic

Procedure MyopicSearchForOpportunities

Input: An augmented model M_A , a threshold δ^* for the increase of expected utility, and Cl_i for each $X_i \in X^m$, the set of unaugmented manipulable variables in M_A .

Output: A sequence of interventions D on the subset of X^m .

```

1.  $D := \emptyset$ ;  $M^* := M_A$ .
2.  $Update(M^*)$ ;  $\mu^* := MEU(M^*)$ ;  $found := \text{false}$ .
3. while  $X^m \neq \emptyset$  and  $found = \text{false}$ 
4.    $M := M^*$ ;  $\mu := \mu^*$ ;  $\delta := \delta^*$ .
5.   for each  $X_i \in X^m$ 
6.      $M' := Augment(M^*, do(X_i), Cl_i)$ .
7.      $Update(M')$ ;  $\mu' := MEU(M')$ .
8.      $\Delta\mu := \mu' - \mu^*$ . /*  $VOINT(D_i)$  */
9.     if  $\Delta\mu > \delta$  then
10.       $\delta := \Delta\mu$ ;  $M := M'$ ;  $\mu := \mu'$ ;
11.       $X := X_i$ ;  $d := \pi_i(M')$ ;
12.       $found := \text{true}$ .
13.   end for each
14.   if  $found = \text{true}$  then
15.      $M^* := Instantiate(M, X, d)$ ;  $\mu^* := \mu$ ;
16.      $X^m := X^m \setminus \{X\}$ ;  $D := D \cup \{(X, d)\}$ ;
17.      $found := \text{false}$ ,
18.   else  $found := \text{true}$ .
19. end while.
20. return  $D$ .

```

Fig. 5. Myopic approach to search for opportunities. $Update(M)$ computes the optimal strategy and maximum expected utility for a model M . $Augment(M, do(X_i), Cl_i)$ denotes the operation of augmenting the model M with an intervention on X_i with the cost of intervention Cl_i . $Instantiate(M, X, d)$ denotes the operation of setting the value of D_X to d in M .

intervention. The model is augmented with an intervention on R as shown in Fig. 4 (a). After intervening on R by setting D_R to d_R^* , suppose that we identify that CC is the next variable to act on by an atomic intervention. We have the model augmented as shown in Fig. 4 (b).

The procedure in Fig. 5 can be applied by a robot to find out the next most effective action. It can also be a useful extension of a modeling environment, which is how we plan to apply it. As illustrated in Fig. 6, we present users with a list of ranked values of interventions, generated by the code Lines 5-13 in Fig. 5.² Users may take the suggestion from the myopic search to perform the intervention at the top of the list, or select any other intervention from the list to alter the generation of the decision sequences. Once the user has entered the

² As far as utility and cost of intervention are concerned, we use a simple linearly additive form of MAU function.

intervention into the system, the system performs the myopic search again to update the ranked list of possible interventions. This interactive environment allows the users also to perform “what if” analysis in generating decision sequences.




Fig. 6. A ranked list of value of interventions on unaugmented manipulable variables.

6 Non-intervening Action

So far, we have constrained ourselves in applying augmented models on decision sequences containing only intervening actions. In many decision problems, we bring in a new mechanism into the model only when we consider a *non-intervening* action. For example, if we model the relations between heart disease (HD) and blood pressure (BP) as mechanisms $f_{HD}(HD) = 0$ and $f_{BP}(HD, BP) = 0$, we have the causal graph $HD \rightarrow BP$. An example of non-intervening action would be the decision of measuring blood pressure (MBP), which brings the variable blood pressure reading (BPR) and the mechanism describing how the blood pressure is measured, $f_{BPR}(BP, BPR, MBP) = 0$, into consideration. We have the causal graph of the augmented model as $HD \rightarrow BP \rightarrow BPR \leftarrow MBP$. We can also represent the cost of measuring blood pressure (CO_{MBP}) as a value function of MBP , i.e., $U(MBP)$. We extend the causal graph into $HD \rightarrow BP \rightarrow BPR \leftarrow MBP \rightarrow CO_{MBP}$.

Example 5. Give the model in Fig. 4, suppose that we consider a non-intervening action (D_A) that assesses the operational status of antenna (A_S) with the cost of observation CO_{D_A} . The causal graph of the augmented model is shown in Fig. 7.




Fig. 7. The model in Fig. 4 augmented with a non-intervening action D_A and the cost of observation CO_{D_A} .

7 Discussion

We present augmented causal models that support decision making with a flexible set of interventions. We introduce the concept of search for opportunities which uses the computation of value of intervention to search for intervening actions on manipulable variables in an augmented model. The result is a greedy search algorithm that suggests the best action to perform.

The concept of value of intervention has also been proposed for causal discovery in active learning [7, 13]. Since the focus of that work is on discovering the true model structure, the value of intervention is defined over all possible models. Our approach, on the other hand, assumes the availability of the true model and uses the value of intervention to advise the next intervention to perform to achieve a desired objective.

Jensen and Vomlelová [5] introduced *unconstrained influence diagrams* that address decision problems where the order of decisions and observations is not determined, but partial temporal ordering of decisions and observations is specified. In our framework, we address decision problems where the choice of a variable being observed or intervened is not even determined. We are currently extending our approach to address decision sequences that contain mixed interventions and observations where previous observations can be invalidated by intervening actions.

Acknowledgments

Our research was supported by the Air Force Office of Scientific Research under grant F49620-00-1-0112. The value of intervention computation and

search for opportunities are implemented in GeNIe and SMILE, available at <http://www2.sis.pitt.edu/~genie>.

References

1. Marek J. Druzdzel and Herbert A. Simon. Causality in Bayesian belief networks. In *Proceedings of the Ninth Annual Conference on Uncertainty in Artificial Intelligence (UAI-93)*, pages 3–11, San Francisco, CA, 1993. Morgan Kaufmann Publishers.
2. David Heckerman and R. Shachter. Decision-theoretic foundations for causal reasoning. *Journal of Artificial Intelligence Research*, 3:405–430, 1995.
3. Ronald A. Howard. Information value theory. *IEEE Transactions on Systems Science and Cybernetics*, SSC-2(1):22–26, August 1966.
4. Ronald A. Howard and James E. Matheson. Influence diagrams. In Ronald A. Howard and James E. Matheson, editors, *The Principles and Applications of Decision Analysis*, pages 719–762. Strategic Decisions Group, Menlo Park, CA, 1981.
5. Finn V. Jensen and Marta Vomlelová. Unconstrained influence diagrams. In *Proceedings of the Eighteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-2002)*, pages 234–241, San Francisco, CA, 2002. Morgan Kaufmann Publishers.
6. James E. Matheson. Using influence diagrams to value information and control. In *Influence Diagrams, Belief Nets and Decision Analysis*, chapter 2, pages 25–48. John Wiley & Sons Ltd., 1990.
7. Kevin Murphy. Active learning of causal Bayes net structure. Technical report, U.C. Berkeley, March 2001.
8. Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, Inc., San Mateo, CA, 1988.
9. Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, UK, 2000.
10. Stuart J. Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, Englewood Cliffs, NJ, 1995.
11. Herbert A. Simon. Causal ordering and identifiability. In William C. Hood and Tjalling C. Koopmans, editors, *Studies in Econometric Method. Cowles Commission for Research in Economics. Monograph No. 14*, chapter III, pages 49–74. John Wiley & Sons, Inc., New York, NY, 1953.
12. Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, Prediction, and Search*. The MIT Press, Cambridge, MA, second edition, 2000.
13. Simon Tong and Daphne Koller. Active learning for structure in Bayesian networks. In *International Joint Conference on Artificial Intelligence*, pages 863–869, Seattle, Washington, August 2001.

Index

action operator, 5
atomic action, 5
augmented models, 7

causal graph, 3
causal model, 3
concurrent actions, 6
conditional action, 5

information gathering, 2

manipulability, 5

observability, 5

potential response, 5
probabilistic causal model, 3

search of opportunities, 11
stochastic action, 5
stochastic policy, 5

value of control, 2
value of information, 2
value of intervention, 2, 9