

ON THE LOGIC OF REPRESENTING DEPENDENCIES BY GRAPHS*

Judea Pearl
Computer Science Department
University of California, Los Angeles
Los Angeles, California 90024

and

Azaria Paz
Computer Science Department
Technion, Israel Institute of Technology
Haifa, Israel

ABSTRACT: We consider 3-place relations $I(x, z, y)$ where, x, y , and z are sets of propositional variables, and $I(x, z, y)$ stands for the statement: "Knowing z renders x independent of y ." We give sufficient conditions on I for the existence of a (minimal) graph G such that $I(x, z, y)$ can be validated by testing whether z separates x from y in G . These conditions define a GRAPHOID. The theory of graphoids uncovers the axiomatic basis of informational dependencies and ties it to vertex-separation conditions in graphs. The defining axioms can also be viewed as inference rules for deducing which propositions are relevant to each other, given a certain state of knowledge.

1. INTRODUCTION

Any system that reasons about knowledge and beliefs must make use of information about dependencies and relevancies. If we have acquired a body of knowledge z and now wish to assess the truth of proposition x , it is important to know whether it would be worthwhile to consult another proposition y , which is not in z . In other words, before we examine y , we need to know if its truth value can potentially generate new information relative to x , information not available from z . For example, in trying to predict whether I am going to be late for a meeting, it is normally a good idea to ask somebody on the street for the time. However, once I establish the precise time by listening to the radio, asking people for the time becomes superfluous and their responses would be irrelevant. Similarly, knowing the color of x 's car normally tells me nothing about the color of Y 's. However, if X were to tell me that he almost mistook Y 's car for his own, the two pieces of information become relevant to each other — whatever I learn about the color of X 's car will have bearing on what I believe the color of Y 's car to be. What logic would facilitate this type of reasoning?

In probability theory, the notion of relevance is given precise quantitative underpinning using the device of *conditional independence*. A variable x is said to be independent of y given the information z if $P(x, y | z) = P(x | z)P(y | z)$. However, it is rather unreasonable to expect people or machines to resort to numerical verification of equalities in order to extract relevance information. The ease and conviction with which people detect relevance relationships strongly suggest that such information is readily available from the organizational structure of human memory, not from numerical values assigned to its components. Accordingly, it would be interesting to explore how assertions about relevance can be inferred qualita-

tively from various models of memory and, in particular, whether the logic of such assertions can be characterized axiomatically.

Since models of human knowledge are often portrayed in terms of various associational networks (e.g. semantic networks [Woods 1975], constraint networks [Montanari 1974] inference nets [Duda, Hart and Nilsson 1976]), a natural starting point would be to examine what types of dependency relations can be captured by a network representation, in the sense that all assertions about dependencies (and independencies) in a given model be deducible from the topological properties of some network.

When we deal with a phenomenon where the notion of neighborhood or connectedness is explicit (e.g., family relations, electronic circuits, communication networks, etc.), we have no problem configuring a graph which represents the main features of the phenomenon. However, in modeling conceptual relations such as causation, association and relevance, it is often hard to distinguish direct neighbors from indirect neighbors; so, the task of constructing a graph representation then becomes more delicate. The notion of conditional independence in probability theory is a perfect example of such a relational structure. For a given probability distribution P and any three variables x, y, z , while it is fairly easy to verify whether knowing z renders x independent of y , P does not dictate which variables should be regarded as direct neighbors. Thus, many topologies might be used to display the dependencies embodied in P .

This paper studies the feasibility of devising graphical representations for dependency models in which the notion of neighborhood is not specified in advance. Rather, what is given explicitly is the relation of "in betweenness." In other words, we are given the means to test whether any given subset S of elements *intervenes in a relation between* elements x and y , but it remains up to us to decide how to connect the elements together in a graph that accounts for these interventions.

Section 1 uncovers the axiomatic basis of dependency models which are isomorphic to vertex separation in graphs. The axioms established can be used both for testing whether a given model lends itself to a complete graphical representation, and for inferring new dependencies from a given initial set. Section 2 examines dependency models called graphoids which may have no graph isomorphism yet possess an effective graphical representation; all their dependencies together with the highest possible number of indepen-

* This work was supported in part by the National Science Foundation, Grants DCR 83-13875 & 85-01234.

dependencies can be displayed by some graph. The theory of graphoids, of which probabilistic dependence is a special case, provides methods for constructing such optimal graphs.

2. GRAPH REPRESENTATION - SEMANTICS AND SYNTAX

2.1 What's in a Link? Suppose we have a collection U of interacting elements and we decide to represent their interactions by an undirected graph G in which the nodes correspond to individual elements of U . Naturally, we would like to display independence between elements by the lack of connectivity between their corresponding nodes in G and, conversely, dependent elements should correspond to connected nodes in G . This requirement alone, however, does not take full advantage of the expressive power of graph representation. It treats all connected components of G as equivalence classes and does not attribute any special significance to the topological configuration within each connected component of G .

Clearly, if graph topology is to convey meaning beyond its connectedness, a semantic distinction must be made between "direct connection" and "indirect connection" in the sense that arbitrarily adding a link between otherwise connected elements should correspond to a totally different state of dependency. This means that the model which governs our understanding of the interactions between the elements of U must also provide us with a criterion for testing 3-place, *conditional independence* relations of the form $I(x, z, y) =$ "x is independent of y conditioned on z." While a variety of interpretations might be given to the terms "independent" and "conditioned on," we shall see that some reasonable general constraints can be imposed on the relation $I(x, z, y)$ if we associate it with the intuitive statement: "knowing y would tell me nothing new about x, if I already know z."

Ideally, we would like to require that if the removal of some subset S of nodes from the graph G renders nodes x and y disconnected (written $\langle x \mid S \mid y \rangle_G$), then this separation should correspond to conditional independence between x and y given S , namely,

$$\langle x \mid S \mid y \rangle_G \implies I(x, S, y)$$

and, conversely,

$$I(x, S, y) \implies \langle x \mid S \mid y \rangle_G$$

This would provide a clear graphical representation for the notion that x does not affect y directly, that its influence is mediated by the variables in S . Unfortunately, we shall next see that these two requirements are too strong; there is often no way of using vertex separation in a graph to display *all* dependencies and independencies embodied in the common notion of information relevancy.

Let $U = \{\alpha, \beta, \dots\}$ be a finite set of elements (e.g. propositions, variables etc.) and let x, y , and z stand for three non intersecting subsets of elements in U . Let M be a model which assigns truth values to the 3-place predicate $I(x, z, y)$ or, in other words, M determines a subset I of triplets (x, y, z) for which the assertion "x is independent of y given z" is true.

Definition: An undirected graph G is a *dependency map* (D -map) of M if there is a one-to-one correspondence between the variables in U and the nodes of G , such that for all non-intersecting subsets,

x, y, z , of variables we have:

$$I(x, z, y)_M \implies \langle x \mid z \mid y \rangle_G \quad (1)$$

Similarly, G is an *Independency map* (I -map) of M if:

$$I(x, z, y)_M \Leftarrow \langle x \mid z \mid y \rangle_G \quad (2)$$

A D -map guarantees that vertices found to be connected are, indeed, dependent, but it may occasionally display dependent variables as separated vertices. An I -map works the opposite way: it guarantees that vertices found to be separated always correspond to genuinely independent variables but does not guarantee that all those shown to be connected are, in fact, dependent. Empty graphs are trivial D -maps, while complete graphs are trivial I -maps.

It is not hard to see that in many reasonable models of informational dependency no graph can be both a D -map and an I -map of M . For example, in models where $I(x, z, y)$ means "y is irrelevant to x once we learn z", we often find *nonmonotonic* behavior -- totally unrelated propositions can become relevant to each other upon learning new facts. For instance, whether it is cloudy or sunshine outside has nothing to do with the type of paper I am currently writing on. However, upon learning that I have difficulty reading my pencil marks, seeing the sun shining through the window makes me doubt the quality of the paper. Such a nonmonotonic model M , implying both $I(x, z_1, y)_M$ and $NOT-I(x, z_1 \cup z_2, y)_M$, cannot have a graph representation which is both an I -map and a D -map, because graph separation always satisfies $\langle x \mid z_1 \mid y \rangle_G \implies \langle x \mid z_1 \cup z_2 \mid y \rangle_G$ for any two subsets z_1 and z_2 of vertices. Thus, D -mapness forces G to display z_1 as a cutset separating x and y , while I -mapness prevents $z_1 \cup z_2$ from separating x and y . No graph can satisfy these two requirements simultaneously.

Being unable to provide graphical representations to some (e.g. nonmonotonic) interpretations of $I(x, z, y)$, raises the question of whether we can formally delineate the class of models which *do* lend themselves to graphical representation. This is accomplished in the following substitution by establishing an axiomatic characterization of the type of dependency relations which are isomorphic to vertex separation in graphs.

2.2 Axiomatic Characterization of Graph-Isomorph Dependencies

Definition: A dependency model M is said to be *graph-isomorph* if there exists a graph $G = (U, E)$ which is both an I -map and a D -map of M , i.e., for every three non-intersecting subsets x, y and z of U we have:

$$I(x, z, y)_M \iff \langle x \mid z \mid y \rangle_G \quad (3)$$

Theorem 1: A necessary and sufficient condition for a dependency model M to be graph-isomorph is that $I(x, z, y)_M$ satisfies the following five independent axioms (the subscript M dropped for clarity):

(symmetry)

$$I(x, z, y) \iff I(y, z, x) \quad (4.a)$$

(subset closure)

$$I(x, z, y \cup w) \implies I(x, z, y) \ \& \ I(x, z, w) \quad (4.b)$$

(intersection)

$$I(x, z \cup w, y) \& I(x, z \cup y, w) \implies I(x, z, y \cup w) \quad (4.c)$$

(strong union)

$$I(x, z, y) \implies I(x, z \cup w, y) \quad \forall w \subset U \quad (4.d)$$

(transitivity)

$$I(x, z, y) \implies I(x, z, \gamma) \text{ or } I(\gamma, z, y) \quad \forall \gamma \notin x \cup z \cup y \quad (4.e)$$

The axioms in (4) are clearly satisfied for vertex separation in graphs. (4.e) is the counter-positive form of connectedness transitivity, stating that, if x is connected to γ and γ is connected to y , then x must also be connected to y . (4.d) states that, if z is a vertex cutset separating x from y , then removing additional vertices w from the graph still leaves x and y separated. (4.c) claims that, if x is separated from w with y removed and, simultaneously, x is separated from y with w removed, then x must be separated from both y and w .

The logical independence of the five axioms can be demonstrated by letting U contain four elements and showing that it is always possible to contrive a subset I of triplets (from the subsets of U) which violates one axiom and satisfies the other four. The proof of Theorem 1 [Pearl and Paz 1985] also provides a simple method of constructing the unique graph G , which is isomorphic to I -- starting with a complete graph, we simply delete every edge (α, β) for which a triplet of the form (α, ζ, β) appears in I .

Having a complete characterization for vertex separation in graphs makes it easy to test whether a given model of dependency lends itself to graphical representation. In fact, it is easy to show that the unrestricted intuitive notion of informational relevancy will, in some context, violate each of the last three axioms. Axiom (4.d) is clearly violated in the non-monotonic example of the preceding subsection. Transitivity (4.e) is violated by that same example because reading difficulties may depend on both the paper quality and the ambient light; yet the latter two are independent of each other. (4.c) is violated in contexts where the propositions y and w logically constrain one another. For instance, if y stands for the proposition "The water temperature is above freezing," and w stands for "The water temperature is above 32°F," then, clearly, knowing the truth of either one of them renders the other superfluous. Yet, contrary to (4.c), this should not render both y and w irrelevant to a third proposition x , say, whether we will enjoy swimming in that water.

Having failed to provide isomorphic graphical representations for even the most elementary models of informational dependency, we settle for the following compromise: Instead of insisting on complete graph isomorphism, we will consider I -maps which may not be D -maps. However, succumbing to the fact that some independencies may escape representation, we will insist that their number be kept at a minimum or, in other words, that the graphs in those maps should contain no superfluous edges.

3. DEPENDENCY MODELS WITH MINIMAL I-MAPS

3.1 Formal Characterization

Definition: A graph G is a *minimal I-map* of dependency model M if no edge of G can be deleted without destroying its I -mapness.

We now define a class of dependency models which possess unique, easily constructed minimal I -maps.

Definition: A *graphoid* is a set I of triplets (x, z, y) where x, z, y are three non-intersecting subsets of elements drawn from a finite collection $U = \{\alpha, \beta, \dots\}$, having the following four properties. (We shall write $I(x, y, z)$ to state that the triplet (x, y, z) belongs to graphoid I .)

Symmetry

$$I(x, z, y) \iff I(y, z, x) \quad (5.a)$$

Subset Closure

$$I(x, z, y \cup w) \implies I(x, z, y) \& I(x, z, w) \quad (5.b)$$

Intersection

$$I(x, z \cup w, y) \& I(x, z \cup y, w) \implies I(x, z, y \cup w) \quad (5.c)$$

Union

$$I(x, z, y \cup w) \implies I(x, z \cup w, y) \quad (5.d)$$

Obviously, every graph-isomorphic dependency is a graphoid, but not vice-versa. The first three properties in (5) are identical to those in (4), while the transitivity requirement (4.e) is waived. Moreover, the union property (5.d) is weaker than (4.d) in that it severely restricts the conditions under which a cutset z can be enlarged by w . In the context of informational dependency, this restriction amounts to saying that learning new facts w will not help an irrelevant fact (y) become relevant if the learned facts (w) were, themselves, irrelevant to begin with.

Theorem 2: Every graphoid I has a unique edge-minimum I -map, $G_0 = (U, E_0)$, constructed by connecting *only* pairs (α, β) for which the triplet $(\alpha, U - \alpha - \beta, \beta)$ is not in I , i.e.,

$$(\alpha, \beta) \notin E_0 \quad \text{iff } I(\alpha, U - \alpha - \beta, \beta) \quad (6)$$

Definition: A *relevance sphere* $R_I(\alpha)$ of an element $\alpha \in U$ is any subset S of elements for which

$$I(\alpha, S, U - S - \alpha) \text{ and } \alpha \in S \quad (7)$$

Let $R_I^*(\alpha)$ stand for the set of all relevance spheres of α . A set is called a *relevance boundary* of α , denoted $B_I(\alpha)$, if it is in $R_I^*(\alpha)$ and if, in addition, none of its proper subsets is in $R_I^*(\alpha)$.

$B_I(\alpha)$ is to be interpreted as the smallest set that "shields" α from the influence of all other elements. Note that $R_I^*(\alpha)$ is non-empty because $I(x, z, \emptyset)$ guarantees that the set $S = U - \alpha$ satisfies (7).

Theorem 3: Every element $\alpha \in U$ in a graphoid I has a unique *relevance boundary* $B_I(\alpha)$. $B_I(\alpha)$ coincides with the set of vertices $B_{G_0}(\alpha)$ adjacent to α in the minimal graph G_0 .

Corollary 1: The set of relevance boundaries $B_I(\alpha)$ forms a *neighbor system*, i.e., a collection $B_I^* = \{B_I(\alpha) : \alpha \in U\}$ of subsets of U such that

- (i) $\alpha \notin B_I(\alpha)$, and
- (ii) $\alpha \in B_I(\beta) \quad \text{iff } \beta \in B_I(\alpha), \quad \alpha, \beta \in U$

Corollary 2: The edge-minimum I -map G_0 can be constructed by connecting each α to all members of its relevance boundary $B_I(\alpha)$.

The usefulness of Corollary 2 lies in the fact that in many cases it is the relevance boundaries $B_I(\alpha)$ that define the organizational structure of human memory. People find it natural to identify the immediate consequences and/or justifications of each action or event, and

these relationships constitute the neighborhood semantics for inference nets used in expert systems [Duda et al. 1976]. The fact that $B_I(\alpha)$ coincides with $B_{G_0}(\alpha)$ guarantees that many independencies can be validated by tests for graph separation at the knowledge level itself (Pearl, 1985).

3.2 An Illustration: To illustrate the role of these axioms consider a simple graphoid defined on a set of four integers $U = \{1, 2, 3, 4\}$. Let I be the set of twelve triplets listed below:

$$I = \{(1, 2, 3), (1, 3, 4), (2, 3, 4), \\ \{1, 2\}, 3, 4), (1, \{2, 3\}, 4), \\ (2, \{1, 3\}, 4), + \text{symmetrical images}\}$$

It is easy to see that I satisfies (5.a)-(5.d) and thus it has a unique minimal I -map G_0 , shown in Figure 1. This graph can be constructed either by deleting the edges (1, 4) and (2, 4) from the complete graph or by computing from I the relevance boundary of each element, i.e., $B_I(1) = \{2, 3\}$, $B_I(2) = \{1, 3\}$, $B_I(3) = \{1, 2, 4\}$, $B_I(4) = \{3\}$.




Figure 1: The Minimal I-Map, G_0 , of I

Suppose that the list contained only the last two triplets (and their symmetrical images):

$$I' = \{(1, \{2, 3\}, 4), (2, \{1, 3\}, 4) + \text{symmetrical images}\}$$

I' is clearly not a graphoid because the absence of the triplets (1, 3, 4) and (2, 3, 4) violates the intersection axiom (5.c). Indeed, if we try to construct G_0 by the usual criterion of edge deletion, the graph in Figure 1 ensues, but it is no longer an I -map of I' ; it shows 3 separating 1 from 4 while (1, 3, 4) is not in I' . In fact, the only I -maps of I' are the three graphs in Figure 2, and the edge-minimum graph is clearly not unique.




Figure 2: The Three I-Maps of I'

Now consider the list

$$I'' = \{(1, 2, 3), (1, 3, 4), (2, 3, 4), (\{1, 2\}, 3, 4), + \text{images}\}$$

I'' satisfies the first three axioms (5.a)-(5.c) but not the union axiom (5.d). Since no triplet of the form $(\alpha, U - \alpha - \beta, \beta)$ appears in I'' , the

only I -map for this list is the complete graph. However, the relevance boundaries of I'' do not form a neighbor set; e.g., $B_{I''}(4) = \{3\}$, $B_{I''}(2) = \{1, 3, 4\}$, so $2 \notin B_{I''}(4)$ while $4 \in B_{I''}(2)$.

3.3

An Example: Probabilistic Dependencies and Their Graphical Representation

Let $U = \{\alpha, \beta, \dots\}$ be a finite set of discrete-valued random variables characterized by a joint probability function $P(\cdot)$, and let x , y , and z stand for any three subsets of variables in U . We say that x and y are conditionally independent given z if

$$P(x, y | z) = P(x | z) P(y | z) \text{ when } P(z) > 0 \quad (8)$$

Eq.(8) is a terse notation for the assertion that for any instantiation z_k of the variables in z and for any instantiation x_i and y_j of x and y , we have

$$P(x=x_i \ \& \ y=y_j | z=z_k) = P(x=x_i | z=z_k) P(y=y_j | z=z_k) \quad (9)$$

The requirement $P(z) > 0$ guarantees that all the conditional probabilities are well defined, and we shall henceforth assume that $P > 0$ for any instantiation of the variables in U . This rules out logical and functional dependencies among the variables a case which would require special treatment.

We shall use the notation $(x \perp z \perp y)_P$ or simply $(x \perp z \perp y)$ to denote the independence of x and y given z . Thus,

$$(x \perp z \perp y)_P \text{ if } P(x, y | z) = P(x | z) P(y | z) \quad (10)$$

Note that $(x \perp z \perp y)$ implies the conditional independence of all pairs of variables $\alpha \in x$ and $\beta \in y$, but the converse is not necessarily true.

The relation $(x \perp z \perp y)$ satisfies the following properties [Lauritzen 1982]:

$$(x \perp z \perp y) \iff P(x | y, z) = P(x | z) \quad (11.a)$$

$$(x \perp z \perp y) \iff P(x, z | y) = P(x | z) P(z | y) \quad (11.b)$$

$$(x \perp z \perp y) \iff \exists f, g : P(x, y, z) = f(x, z) g(y, z) \quad (11.c)$$

$$(x \perp z \perp y) \iff P(x, y, z) = P(x | z) P(y, z) \quad (11.d)$$

$$(x \perp z \perp y) \implies (x \perp z, f(y) \perp y) \quad (12.a)$$

$$(x \perp z \perp y) \implies (f(x, z) \perp z \perp y) \quad (12.b)$$

These properties are based on the numeric representation of P and, therefore, would not be adequate for characterizing its graphical representation.

We now ask what *logical* conditions, void of any reference to numerical forms, should constrain the relationship $(x \perp z \perp y)$. The next set of properties constitute such a logical basis.

Theorem 4: Let x , y , and z be three non-intersecting subsets of variables from U , and let $(x \perp z \perp y)$ stand for the relation "x is independent of y, given z" in some probabilistic model P . The following five independent conditions must then hold:

Symmetry

$$(x \perp z \perp y) \iff (y \perp z \perp x) \quad (13.a)$$

Closure for subsets:

$$(x \perp z \perp y, w) \implies (x \perp z \perp y) \ \& \ (x \perp z \perp w) \quad (13.b)$$

Weak closure for intersection:

$$(x \perp z, w \perp y) \ \& \ (x \perp z, y \perp w) \implies (x \perp z \perp y, w) \quad (13.c)$$

Weak closure for union:

$$(x \perp z \perp y, w) \implies (x \perp z, w \perp y) \quad (13.d)$$

Contraction:

$$(x \perp z, y \perp w) \ \& \ (x \perp z \perp y) \implies (x \perp z \perp y, w) \quad (13.e)$$

The intuitive interpretation of Eqs. (13.c) through (13.e) follows. (13.c) states that if y does not affect x when w is held constant and if, simultaneously, w does not affect x when y is held constant, then neither w nor y can affect x . (13.d) states that learning an irrelevant fact (w) cannot help another irrelevant fact (y) become relevant. (13.e) can be interpreted to state that if we judge w to be irrelevant (to x) after learning some irrelevant facts y , then w must have been irrelevant before learning y . Together, the expansion and construction properties mean that learning irrelevant facts should not alter the relevance status of other propositions in the system; whatever was relevant remains relevant, and what was irrelevant remains irrelevant.

The proof of Theorem 1 can be derived by elementary means from the definition (8) and from the basic axioms of probability theory. The intersection property is the only one which requires the assumption $P(x) > 0$ and will not hold when the variables in U are constrained by logical dependencies. In such a case, Theorem 1 will still retain its validity if we regard each logical constraint as having some small probability ϵ of being violated and let $\epsilon \rightarrow 0$.

Obviously, probabilistic dependencies form a graphoid and, therefore, possess the graph properties of Theorems 2 and 3. In particular, we have:

Corollary 3: To every probability distribution P , there corresponds a unique minimal I -map $G_o = (U, E_o)$ constructed by the criterion

$$(\alpha, \beta) \in E_o \text{ iff } (\alpha \perp U - \alpha - \beta \perp \beta).$$

Equivalently, G_o can be constructed by connecting each variable α to the smallest set S of variables satisfying

$$P(\alpha | S) = P(\alpha | U - \alpha)$$

4. CONCLUSIONS

We have established an axiomatic characterization of dependency models which are representable by graphs, and we have identified two essential properties: weak closure for intersection (5.c), and weak closure for union (5.d). These two axioms enable us to construct an edge-minimum graph in which every cutset corresponds to a genuine independence condition, and these were chosen, therefore, as the formal definition of graphoid systems — a general model of informational dependency. Vertex separation in graphs, probabilistic independence and partial uncorrelatedness are special cases of graphoid systems where the two defining axioms are augmented with additional requirements.

The graphical properties associated with graphoid systems offer an effective inference mechanism for deducing, in any given state of knowledge, which propositional variables are relevant to each other. If we identify the relevance boundaries associated with each proposition in the system, and treat them as neighborhood relations defining a graph G_o , then we can correctly deduce irrelevance relationships by testing whether the set of currently known propositions constitutes a cutset in G_o .

REFERENCES

- Duda, R. O., Hart, P. E., and Nilsson, N. J., (1976), Subjective Bayesian Methods for Rule-Based Inference Systems, *Proceedings, 1976 NCC (AFIPS)*, 45, 1075-1082.
- Lauritzen, S.L. (1982), *Lectures on Contingency Tables*, 2nd Ed., U. of Aalborg Press, Aalborg, Denmark.
- Montanari, U. (1974), Networks of Constraints, *Information Science*, Vol. 7, pp. 95-132.
- Pearl, J. (1985), Fusion, Propagation and Structuring in Bayesian Networks, Technical Report CSD-850022, June 1985.
- Pearl, J. & Paz, A., (1985), GRAPHOIDS: a Graph-Based Logic for Reasoning about Relevance Relations, *Technical Report CSD-850038*, December 1985.
- Woods, W.A. (1975), "What's in a Link: Foundations for Semantic Networks" in Bobrow and Collins (eds.), *Representation and Understanding*, Academic Press, Inc.