

**Computer Science Department Technical Report
Cognitive Systems Laboratory
University of California
Los Angeles, CA 90024-1596**

ON THE LOGIC OF INFLUENCE DIAGRAMS

**Dan Geiger
Judea Pearl**

**July 1988
CSD-880060**

ON THE LOGIC OF INFLUENCE DIAGRAMS *

Dan Geiger & Judea Pearl

Cognitive Systems Laboratory
Computer Science Department
University of California
Los Angeles, CA 90024

ABSTRACT

This paper explores the role of influence diagrams as a representation of conditional independence relationships. We show that influence diagrams offer polynomially *sound* and *complete* inference mechanisms for inferring conditional independence relationships from a given *causal* set of such relationships. As a consequence, *d-separation*, a graphical criterion for identifying independencies in an influence diagram, is shown to uncover more valid independencies than any other criterion. In addition, we employ the *Armstrong* property of conditional independence to show that the dependence relationships displayed by an influence diagram are inherently *consistent*, i.e. every influence diagram represents all conditional independencies embodied in some distribution and non other.

*This work was partially supported by the National Science Foundation Grant #IRI-8610155. "Graphoids: A Computer Representation for Dependencies and Relevance in Automated Reasoning (Computer Information Science)".

ON THE LOGIC OF INFLUENCE DIAGRAMS *

Dan Geiger & Judea Pearl

Cognitive Systems Laboratory
Computer Science Department
University of California
Los Angeles, CA 90024

ABSTRACT

This paper explores the role of influence diagrams as a representation of conditional independence relationships. We show that influence diagrams offer polynomially *sound* and *complete* inference mechanisms for inferring conditional independence relationships from a given *causal* set of such relationships. As a consequence, *d-separation*, a graphical criterion for identifying independencies in an influence diagram, is shown to uncover more valid independencies than any other criterion. In addition, we employ the *Armstrong* property of conditional independence to show that the dependence relationships displayed by an influence diagram are inherently *consistent*, i.e. every influence diagram represents all conditional independencies embodied in some distribution and non other.

INTRODUCTION AND SUMMARY

Networks employing directed graphs have a long and rich tradition, starting with the geneticist Wright (1921). He developed a method called *path analysis* [Wright, 1934] which later on, became an established representation of causal models in economics [Wold, 1964], sociology [Kenny, 1979], [Blalock, 1971] and psychology [Duncan, 1975]. *Influence diagrams* represent another component in this tradition [Howard and Matheson, 1981], [Shachter, 1988] and [Smith, 1987]. These were developed for decision analysis and contain both event nodes and action nodes (our definition excludes action nodes). *Recursive models* is the name given to such networks by statisticians seeking meaningful and effective decompositions of contingency tables [Lauritzen, 1982], [Wermuth & Lauritzen, 1983], [Kiiveri et al, 1984]. *Bayesian Belief Networks* (or *Causal Networks*) is the name adopted by Pearl for describing networks that perform evidential reasoning ([Pearl, 1986a]). This paper establishes a clear semantics for these networks that might explain their wide usage for representing probabilistic knowledge and as metaphors for decision and evidential reasoning.

Influence diagrams are viewed as an economical scheme for representing conditional independence relationships. The nodes of an influence diagram represent variables in some domain and its topology is specified by a list of conditional independence judgements elicited from an expert in this domain. The list designates parents to each variable v by asserting that v is conditionally independent of all its predecessors, given its parents (in some total order of the variables). This input list implies many additional conditional independencies that can be read off the diagram. For example, the diagram asserts that, given its parents, v is also conditionally independent of all its

*This work was partially supported by the National Science Foundation Grant #IRI-8610155. "Graphoids: A Computer Representation for Dependencies and Relevance in Automated Reasoning (Computer Information Science)".

non-descendants [Howard and Matheson, 1981]. Additionally, if S is a set of nodes containing v 's parents, v 's children and the parents of those children, then v is independent of all other variables in the system, given those in S [Pearl, 1986b]. These assertions are examples of many *valid consequences* of the input list i.e., assertions that hold for every probability distribution that satisfies the conditional independencies specified by the input. If one ventures to perform topological transformations (e.g., arc reversal or node removal [Shachter, 1988]) on the diagram, caution must be exercised to ensure that each transformation does not introduce extraneous, invalid independencies, and/or that the number of valid independencies which become obscured by the transformation is kept at a minimum. Thus, in order to decide which transformations are admissible, one should have a simple graphical criterion for deciding which conditional independence statement is valid and which is not.

This paper deals with the following questions:

1. What are the valid consequences of the input list ?
2. What are the valid consequences of the input list that can be read off the diagram ?
3. Are the two sets identical?

The answers obtained are as follows

1. A statement is a valid consequence of the input set if and only if it can be derived from it using Dawid's [1979] axioms. Denoting by $I(X, Z, Y)$ the statement: "*the variables in X are conditionally independent of those in Y , given those in Z* ", we may express these axioms as follows:

$$\text{Symmetry} \quad I(X, Z, Y) \Rightarrow I(Y, Z, X) \quad (1.a)$$

$$\text{Decomposition} \quad I(X, Z, Y \cup W) \Rightarrow I(X, Z, Y) \ \& \ I(X, Z, W) \quad (1.b)$$

$$\text{Weak Union} \quad I(X, Z, Y \cup W) \Rightarrow I(X, Z \cup W, Y) \quad (1.c)$$

$$\text{Contraction} \quad I(X, Z \cup Y, W) \ \& \ I(X, Z, Y) \Rightarrow I(X, Z, Y \cup W) \quad (1.d)$$

2. Every statement that can be read off the diagram using the d -separation criterion is a valid consequence of the input list [Pearl & Verma, 1987].

The d -separation condition is defined as follows [Pearl, 1986b]: For any three disjoint subsets X, Y, Z of nodes in the diagram, Z is said to d -separate X from Y , denoted $I(X, Z, Y)_D$, if there is no path from a node in X to a node in Y along which: 1. every node that delivers an arrow is outside Z , and 2. every node with converging arrows either is in Z or has a descendant in Z (the definition is elaborated in the next section).

3. The two sets are identical, namely, a statement is valid IF AND ONLY IF it is graphically-validated under d -separation in the diagram.

The first result establishes the *decidability* of verifying whether an arbitrary statement is a valid consequence of the input set. The second result renders the d-separation criterion a *polynomially sound* inference rule, i.e., it runs in polynomial time and certifies only valid statements. The third renders the d-separation criterion a *polynomially complete* inference rule, i.e., the diagram constitutes a sound and complete inference mechanism that identifies, in polynomial time, each and every valid consequence in the system.

The results above are true only for *causal* input sets i.e., those that recursively specify the relation of each variable to its predecessors in some (chronological) order. The general problem of verifying whether a given conditional independence statement logically follows from an arbitrary set of such statements, may be undecidable. Its decidability hinges on whether Dawid's axioms are complete i.e., whether these axioms are powerful enough to derive all valid consequences of an arbitrary input list. The completeness problem is treated in [Geiger & Pearl, 1988] and completeness results for specialized subsets of probabilistic dependencies have been obtained. All axioms encountered so far are derivable from Dawid's axioms, which suggests that they are indeed complete. Result-1 can be viewed as yet another completeness result for the special case in which the input statements form a causal set. This means that applying axioms (1.a) through (1.d) on a causal input list is guaranteed to generate all valid consequences and non other. Interestingly, result-2 above holds for any statements that obey Dawid's axioms, not necessarily probabilistic conditional independencies. Thus, influence diagrams can serve as polynomially sound inference mechanisms (possibly not complete) for a variety of dependence relationships, e.g., partial correlations and qualitative database dependencies.

In the preceding sections we refer to influence diagrams conjuncted with d-separation as DAGs (Directed Acyclic Graphs).

SOUNDNESS AND COMPLETENESS

The definition of d-separation is best motivated by regarding DAGs as a representation of causal relationships. Designating a node for every variable and assigning a link between every cause to each of its direct consequences provides a graphical representation of a causal hierarchy. For example, the propositions "It is raining" (α), "the pavement is wet" (β) and "John slipped on the pavement" (γ) are well represented by a three node chain, from α through β to γ ; it indicates that either rain or wet pavement could cause slipping, yet wet pavement is designated as the *direct cause*; rain could cause someone to slip if it wets the pavement, but not if the pavement is covered. Moreover, knowing the condition of the pavement renders "slipping" and "raining" independent, and this is represented graphically by a d-separation condition, $I(\alpha, \gamma | \beta)_D$, showing node α and γ separated from each other by node β . Assume that "broken pipe" (δ) is considered another direct cause for wet pavement, as in figure 1. An induced dependency exists between the two events that may cause the pavement to get wet: "rain" and "broken pipe". Although they appear connected in Figure 1, these propositions are marginally independent and become dependent once we learn that the pavement is wet or that someone broke his leg. An increase in our belief in either cause would decrease our belief in the other as it would "explain away" the observation. The following definition of d-separation permits us to graphically identify such induced dependencies from the DAG (d connoted "directional").

Definition: If X , Y , and Z are three disjoint subsets of nodes in a DAG D , then Z is said to *d-separate* X from Y , denoted $I(X, Y | Z)_D$, iff there is no bi-directed path from a node in X to a node in Y along which every node which delivers an arrow is outside Z and every node with converging arrows either is or has a descendant in Z . A path satisfying the conditions above is said to be *active*, otherwise it is said to be *blocked* (by Z). Whenever a statement $I(X, Y | Z)_D$ holds in a DAG D , the predicate $I(X, Y | Z)$ is said to be *graphically-verified* (or an *independency*), otherwise it is *graphically-falsified* by D (or a *dependency*).

In figure 2, for example, $X = \{2\}$ and $Y = \{3\}$ are d -separated by $Z = \{1\}$; the path $2 \leftarrow 1 \rightarrow 3$ is blocked by $1 \in Z$ while the path $2 \rightarrow 4 \leftarrow 3$ is blocked because 4 and all its descendants are outside Z . Thus $I(2, 1, 3)$ is graphically-verified by D . However, X and Y are not d -separated by $Z' = \{1, 5\}$ because the path $2 \rightarrow 4 \leftarrow 3$ is rendered active: learning the value of the consequence 5, renders its causes 2 and 3 dependent, like opening a pathway along the converging arrows at 4. Consequently, $I(2, \{1, 5\}, 3)$ does not hold and is therefore graphically-falsified by D .




Figure 1




Figure 2

Definition: If X , Y , and Z are three disjoint subsets of variables of a distribution P , then X and Y are said to be conditionally independent given Z , denoted $I(X, Z, Y)_P$ iff $P(X | Z, Y) = P(X | Z)$ for all possible values of X , Y and Z . $I(X, Z, Y)_P$ is called a (*conditional independence*) *statement*. A conditional independence statement σ *logically follows* from a set Σ of such statements if σ holds in every distribution that obeys Σ , in which case we also say that σ is a *valid consequence* of Σ .

Ideally, to employ a DAG D as a graphical representation for dependencies of some distribution P we would like to require that for every three disjoint sets of variables in P (and nodes in D) the following equivalence would hold

$$I(X, Z, Y)_D \text{ iff } I(X, Z, Y)_P \quad (2)$$

This would provide a clear graphical representation of all variables that are conditionally independent. When equation (2) holds, D is said to be a *perfect map* of P . Unfortunately, this requirement is often too strong because there are many distributions that have no perfect map in DAGs. The spectrum of probabilistic dependencies is in fact so rich that it cannot be cast into any representation scheme that uses polynomial amount of storage ([Verma, 1987]). Geiger [1987] provides a graphical representation based on a collection of graphs (Multi-DAGs) that is powerful enough to perfectly represent an arbitrary distribution, however, as shown by Verma, it requires, on the average, an exponential number of DAGs. Being unable to provide perfect maps at a reasonable cost we compromise on graphs that represent each and every dependency of P , acknowledging the fact that some independencies would escape representation.

Definition: A DAG D is said to be an *I-map* of P if for every three disjoint subsets X , Y and Z of variables the following holds:

$$I(X, Z, Y)_D \Rightarrow I(X, Z, Y)_P$$

The natural requirement for these I-maps is that the number of undisplayed independencies be minimized.

The task of finding a DAG which is a minimal I-map of a given distribution P was solved in [Pearl & Verma, 1987]. Their algorithm consists of the following steps: assign a total ordering d to the variables of P . For each variable i of P , identify a minimal set of predecessors S_i that renders i independent of all its other predecessors (in the ordering of the first step). Assign a direct link from every variable in S_i to i . The resulting DAG is an I-map of P , and is minimal in the sense that no edge can be deleted without destroying its I-mapness. The input list L for this construction consists of n conditional independence statements, one for each variable, all of the form $I(i, S_i, U_{(i)} - S_i)$ where $U_{(i)}$ is the set of predecessors of i and S_i is a subset of $U_{(i)}$ that renders i conditionally independent of all its other predecessors. This set of conditional independence statements is called a *causal* input list and is said to *define* the DAG D . The term "causal" input list is derived from the following analogy: Suppose we order the variables chronologically, such that a cause always precedes its effect. Then a causal input list selects from all potential causes of an affect i a minimal subset that is sufficient to explain i , thus rendering all other preceding events superfluous. This selected subset of variables are considered the *direct causes* of i and therefore each is connected to it by a direct link.

Clearly, the constructed DAG represents more independencies than those listed in the input, namely, all those that are graphically verified by the d -separation criterion. [Pearl & Verma, 1987] analysis guarantees that all graphically-verified statements are indeed valid in P i.e., the DAG is an I-map of P . However, this paper shows that the constructed DAG has an additional property; it graphically-verifies every conditional independence statement that logically follows from L (i.e. holds in every distribution that obeys L). Hence, we cannot hope to improve the d -separation criterion to display more independencies, because all valid consequences of L (which defines D) are already captured by d -separation.

The three theorems below formalize the above discussion.

Theorem 1 (soundness) [Pearl & Verma, 1987]: Let D be a DAG defined by a causal input list L . Then, every graphically-verified statement is a valid consequence of L .

Theorem 2 [Pearl & Verma, 1987]: Let D be a DAG defined by a causal input list L . Then, the set of graphically-verified statements is exactly the closure of L under axioms (1.a) through (1.d).

Theorem 3 (completeness): Let D be a DAG defined by a causal input list L . Then, every valid consequence of L is graphically-verified by D (equivalently, every graphically-falsified statement in D is not a valid consequence of L).

The first two theorems are more general than the third in the sense that they hold for every dependence relationship that obeys axioms (1.a) through (1.d), not necessarily those based on probabilistic conditional independence (proofs can be found in [Verma, 1986]). Among these dependence relationships are partial correlations ([Pearl & Paz, 1986]) and qualitative dependencies ([Fagin, 1977], [Shafer et al, 1987]) which can readily be shown to obey axioms (1). Thus, for example, the transformation of arc-reversal and node removal ([Howard & Matheson, 1981]) can be shown valid by purely graphical consideration, simply showing that every statement verified in the transformed is also graphically-verified in the original graph. Theorem 1 guarantees that a DAG displays only valid statements. Theorem 2 guarantees that a DAG displays all statements that are derivable from L via axioms (1). The third theorem, which is the main contribution of this paper, assures that a DAG displays all statements that logically follow from L i.e., the axioms in (1) are complete, capable of deriving all valid consequences of a causal input list. Moreover, since a statement in a diagram can be verified in polynomial time, theorem 3 provides a complete polynomial inference mechanism for deriving all independency statements that are implied by a causal input set.

Theorem 3 is proven in the appendix by actually constructing a distribution P_σ that satisfies all conditional independencies in L and violates any statement σ graphically-falsified by D . This distribution precludes σ from being a valid consequence of L and therefore, since the construction can be repeated for every graphically-falsified statement, none of these statements is a valid consequence of L .

We conclude by showing how these theorems can be employed as an inference mechanism. Assume an expert has identified the following conditional independencies between variables denoted 1 through 5:

$$L = \{ I(2, 1, \emptyset), I(3, 1, 2), I(4, 23, 1), I(5, 4, 123) \}$$

(the first statement in L is trivial). We raise two questions. First, what is the set of all valid consequences of L ? Second, in particular, is $I(3, 124, 5)$ a valid consequence of L ? For general input lists the answer for such questions may be undecidable but, since L is a causal list, it defines a DAG that graphically verifies each and every valid consequence of L . The DAG D is the one shown in figure 2. Therefore, the DAG constitutes a dense representation of all valid consequences of L . To answer the second question, we simply observe that $I(3, 124, 5)$ is graphically-verified in D . A graph-based algorithm for another subclass of statements, called *fixed context* statements, is given in [Geiger & Pearl, 1988]. In that paper, results analogous to theorem 1 through 3 are proven for Markov-fields; a representation scheme based on undirected graphs ([Isham, 1981], [Lauritzen, 1982]).

EXTENSIONS AND ELABORATIONS

Theorem 3 can be restated to assert that for every DAG D and any dependency σ there exist a probability distribution P_σ that satisfies D 's causal input set L and the dependency σ . By theorem 2, P_σ must satisfy all graphically-verified statements as well because they are all derivable from L by Dawid's axioms. Thus, theorems 2 and 3 guarantee the existence of a distribution P_σ that satisfies all graphically verified statements and a single arbitrary-chosen dependency. The question answered in this section is the existence of a distribution P that satisfies all independencies of D and all its dependencies not merely a single dependency. We show that such a distribution exists, which legitimizes the use of DAGs as a representation scheme for probabilistic dependencies; a model builder who uses the language of DAGs to express dependencies is guarded from inconsistencies.

The construction of P is based on the Armstrong property of conditional independence.

Definition: Conditional independence is an *Armstrong relation* in a class of distributions \mathbf{P} if there exists an operation \otimes that maps finite sequences of distributions of \mathbf{P} into a distribution of \mathbf{P} , such that if σ is a conditional independence statement and if $P_i, i=1..n$ are distributions in \mathbf{P} , then σ holds for $\otimes(P_i | i=1..n)$ iff σ holds for each P_i .

The notion of Armstrong relation is borrowed from database theory ([Fagin 1980]). We concentrate on two families of distributions \mathbf{P} : All distributions, denoted PD and strictly positive distributions, denoted PD^+ . Conditional independence can be shown to be an Armstrong relation in both families. The construction of the operation \otimes is given below, however the proof is omitted and can be found in ([Geiger & Pearl, 1988]).

Theorem 4 ([Geiger & Pearl, 1988]): Conditional independence is an Armstrong relation in PD and in PD^+ .

We shall construct the operation \otimes for conditional independence using a binary operation \otimes' such that if $P = P_1 \otimes' P_2$ then for every conditional independency statement σ we get

$$\otimes' P_i \text{ obeys } \sigma \text{ iff } P_1 \text{ obeys } \sigma \text{ and } P_2 \text{ obeys } \sigma. \quad (5)$$

The operation \otimes is recursively defined in terms of \otimes' as follows:

$$\otimes \{ P_i \mid i=1..n \} = ((P_1 \otimes' P_2) \otimes' P_3) \otimes' \dots P_n.$$

Clearly, if \otimes' satisfies equation (5), then \otimes satisfies the the requirement of an Armstrong relation, i.e.

$$P \text{ obeys } \sigma \text{ iff } \forall_i P_i \text{ obeys } \sigma.$$

Therefore, it suffices to show that \otimes' satisfies (5).

Let P_1 and P_2 be two distributions sharing the variables x_1, \dots, x_n . Let A_1, \dots, A_n be the domains of x_1, \dots, x_n in P_1 and let an instantiation of these variables be $\alpha_1, \dots, \alpha_n$. Similarly, let B_1, \dots, B_n be the domains of x_1, \dots, x_n in P_2 and β_1, \dots, β_n an instantiation of these variables. Let the domain of $P = P_1 \otimes' P_2$ be the product domain $A_1 B_1, \dots, A_n B_n$ and denote an instantiation of the variables of P by $\alpha_1 \beta_1, \dots, \alpha_n \beta_n$. Define $P_1 \otimes' P_2$ by the following equation:

$$P(\alpha_1 \beta_1, \alpha_2 \beta_2, \dots, \alpha_n \beta_n) = P_1(\alpha_1, \alpha_2, \dots, \alpha_n) \cdot P_2(\beta_1, \beta_2, \dots, \beta_n).$$

The proof that P satisfies the condition of theorem 4 uses only the definition of conditional independence and can be found in [Geiger & Pearl 1988]. The adequacy of this construction for PD^+ is due to the fact that \otimes produces a strictly positive distribution whenever the input distributions are strictly positive.

Theorem 5: For every DAG D there exists a distribution P such that for every three disjoint sets of variables X, Y and Z the following holds:

$$I(X, Z, Y)_D \text{ iff } I(X, Z, Y)_P$$

Proof: Let $P = \otimes \{ P_\sigma \mid \sigma \text{ is a dependency in a DAG } D \}$ where P_σ is a distribution obeying all independencies of D and a dependency σ . By theorem 3, a distribution P_σ always exists. P satisfies the requirement of theorem 5 because it obeys only statements that hold in every P_σ and these are exactly the ones verified by D . \square

The construction presented in the proof of theorem 5 leads to a rather complex distribution, where the domain of each variable is unrestricted. It still does not guarantee that a set of dependencies and independencies represented by DAGs is realizable in a more limited class of distributions such as normal or those defined on binary variables. We conjecture that these two classes of distributions are sufficiently rich to permit the consistency of DAG representation.

ACKNOWLEDGMENT

We thank Ron Fagin for pointing out the usefulness of the notion of Armstrong relation. Thomas Verma and Azaria Paz provided many useful discussions on the properties of dependency models.

REFERENCES

- [Blalock, 1971] Blalock, H.M., *Causal Models in the Social Sciences*, London: Macmillan, 1971.
- [Dawid, 1979] A.P. Dawid, "Conditional Independence in Statistical Theory," *J.R. Statist. Soc. B*, Vol. 41, No. 1, 1979, pp. 1-31.
- [Duncan, 1975] O.D. Duncan, *Introduction to Structural Equation Models*, New York: Academic Press, 1975.
- [Fagin, 1977] R. Fagin, "Multivalued Dependencies and a New Form for Relational Databases," *ACM Transactions on Database Systems*, Vol. 2, No. 3, 262-278, September 1977.
- [Fagin, 1980] R. Fagin, "Horn clauses and database dependencies", Proc. ACM SIGACT Symp. on Theory of Computing, Los Angeles, April 1980.
- [Geiger, 1987] D. Geiger, "Towards the Formalization of Informational Dependencies," Based on the author's thesis. UCLA Cognitive Systems Laboratory, *Technical Report (R-102)*, December 1987.
- [Geiger & Pearl, 1988] D. Geiger & J. Pearl, "Logical and Algorithmic Properties of Conditional Independence," UCLA Cognitive Systems Laboratory, *Technical Report 870056 (R-97)*, February, 1988.
- [Howard & Matheson, 1981] R. A. Howard & J. E. Matheson, "Influence Diagrams," In, *Principles and Applications of Decision Analysis*, Menlo Park, Ca.: Strategic Decisions Group, 1981.
- [Isham, 1981] V. Isham, "An Introduction to Spatial Point Processes and Markov Random Fields," *International Statistical Review*, Vol. 49, 1981, pp. 21-43.
- [Kenny, 1979] D.A. Kenny, *Correlation and Causality*, John Wiley and Sons, 1979.
- [Kiiveri et al, 1984] H. Kiiveri, T.P. Speed & J.B. Carlin, "Recursive Causal Models," *Journal of Australian Math Society*, Vol. 36, 1984, pp. 30-52.
- [Lauritzen, 1982] S. L. Lauritzen, *Lectures on Contingency Tables*, 2nd Ed., University of Aalborg Press, Aalborg, Denmark, 1982.
- [Miller, 1964] K.S. Miller, *Multidimensional Gaussian distributions*, Wiley Siam series in applied math, 1964.
- [Pearl, 1986a] J. Pearl, "Fusion, Propagation and Structuring in Belief Networks," UCLA Computer Science Department *Technical Report 850022 (R-42)*; also, *Artificial Intelligence*, Vol. 29, No. 3, September 1986(b), pp. 241-288.
- [Pearl, 1986b] J. Pearl, "Markov and Bayes networks: A comparison of two Graphical Representations of Probabilistic Knowledge", UCLA Cognitive System Laboratory, Technical Report TR-46, October 1986.

- [Pearl & Paz, 1986] J. Pearl & A. Paz, "GRAPHOIDS: a Graph-based Logic for Reasoning about Relevance Relations," UCLA Computer Science Department *Technical Report 850038 (R-53)*, October 1985; also, *Proceedings, ECAI-86*, Brighton, U.K., June 1986.
- [Pearl & Verma, 1987] J. Pearl, T. Verma. "The logic of representing dependencies by directed acyclic graphs" , Proc. AAA-I, Seattle Washington, July 1987.
- [Shachter, 1987] Shachter, "Probabilistic Inference and Influence Diagrams", to appear in *Operations Research*, 1988.
- [Shafer et al, 1987] G. Shafer, P.P. Shenoy, & K. Mellouli, "Propagating Belief Functions in Qualitative Markov Trees," The University of Kansas, Lawrence School of Business, working paper no. 190, June 1987. To appear in *International Journal of Approximate Reasoning*.
- [Smith, 1987] J.Q. Smith, "Influence diagrams for Bayesian Decision Analysis", Technical report #100, department of Statistics, University of Warwick, Coventry, England, January 1987.
- [Verma, 1986] T. S. Verma, "Causal Networks: Semantics and Expressiveness," UCLA Cognitive Systems Laboratory *Technical Report R-65*, Los Angeles, California, in preparation, 1986.
- [Verma, 1987] T. S. Verma, "Some Mathematical Properties of Dependency Models" UCLA Cognitive Systems Laboratory, Los Angeles, California, in preparation, 1987.
- [Wermuth & Lauritzen, 1983] N. Wermuth & S.L. Lauritzen, "Graphical and Recursive Models for Contingency Tables," *Biometrika*, Vol. 70, 1983, pp. 537-552.
- [Wold, 1964] H. Wold, *Econometric Model Building*, Amsterdam: North-Holland Publishing Co., 1964.
- [Wright, 1934] S. Wright, "The Method of Path Coefficients," *Ann. Math. Statist.*, Vol. 5, 1934, pp. 161-215.

APPENDIX

Theorem 3 (completeness): Let D be a DAG defined by a causal input list L . Then, every valid consequence of L is graphically-verified by D .

Proof: Let $\sigma = I(X, Z, Y)$ be an arbitrary graphically-falsified statement in D . We construct a distribution P_σ that satisfies all conditional independencies in the input list L and violates σ . This distribution precludes σ from being a valid consequence of L and therefore, every valid consequence of L must be graphically-verified in D .

From the definition of d -separation, there must exist an active path between an element α in X and an element β in Y that is not d -separated by Z (i.e. a bi-directed path). Ensuring that P_σ violates the conditional independency $I(\alpha, Z, \beta)$, denoted σ' , guarantees that also σ is violated because any distribution that renders X and Y conditionally independent must render each of their individual variables independent as well (axiom (1.b)).

P_σ is defined in terms of a simplified DAG D_σ . This DAG is constructed by removing as many links as possible from D such that σ remains unverified in D_σ . This process clearly preserves all previously verified statements but caution is exercised not to remove links that would render σ graphically-verified in D_σ . We will conclude the proof by constructing a distribution P_σ which satisfies all graphically-verified statements of D_σ (hence also those of D) and violates σ' .

Let q be an active path (by Z) between α and β with a minimum number of head-to-head nodes (i.e. nodes with converging arrows) denoted, left to right, h_1, h_2, \dots, h_k . Let z_i be the closest (wrt path length) descendent of h_i in Z and let p_i be the directed path from h_i to z_i (if $h_i \in Z$ then $z_i = h_i$). We define D_σ to be a subgraph of D containing only the links that form the paths p_i 's and the path q . We make two claims about the topology of the resulting DAG. First, the paths p_i are all distinct. Second, for any i , h_i is the only node shared by p_i and q . The resulting DAG is depicted in figure 3 (note that some nodes, including nodes of Z , might become isolated in D_σ).




Figure 3

Proof of claim 1: Assume, by contradiction, that there are two paths p_i and p_j ($i < j$) with a common node γ . Under this assumption, we find an active path between α and β that has less head-to-head nodes than q , contradicting the minimality of the latter. If γ is neither h_i nor h_j then the path $(\alpha, h_i, \gamma, h_j, \beta)$ is an active path (by Z); Each of its head-to-head nodes is or has a descendent in Z because it is either γ or a head-to-head node of q . Every other node lies either on the active path q and therefore is outside Z or lies on p_i (p_j) in which case, since it has a descendent γ , it must also be outside Z . The resulting path contradicts the minimality of q since both h_i and h_j are no longer head-to-head nodes while γ is the only newly introduced head-to-head node. If $\gamma = h_j$ then, using similar arguments, the path $(\alpha, h_i, \gamma, \beta)$ (see figure 4), which has less head-to-head nodes than q , can readily be shown active (the case $\gamma = h_i$ is similar).




Figure 4

Proof of claim 2: Assume p_i and q have in common a node γ other than h_i and assume w.l.o.g that it lies between h_i and β . This node is not a head-to-head node on q because p_i is distinct from all other p_j 's. The node γ cannot belong to Z because otherwise q would not have been active. Thus, the path $(\alpha, h_i, \gamma, \beta)$ must be an active path which contradicts the minimality of q (figure 5).




Figure 5

In the following discussion we call a path containing no head-to-head node a *regular path*. Let P_σ be a normal distribution with the following covariance matrix Γ

$$\Gamma = (\rho_{ij}) \quad \rho_{ij} = \begin{cases} 0 & \text{There exists no regular path from node } i \text{ to node } j \\ \rho^l & \text{There exists a regular path of length } l \text{ between node } i \text{ and node } j \end{cases} \quad (3)$$

Since D_σ is singly connected there exist at most one path between any two nodes. Any value of ρ satisfying $0 < \rho < 1$ would render Γ positive definite and therefore a valid covariance matrix. We claim that this distribution satisfies all independencies of D_σ and violates $I(\alpha, Z, \beta)$. To evaluate $I(\alpha, Z, \beta)$ we first form the projection of P_σ on the variables α, β and Z . Since P_σ is normal, this projection is also normal and its covariance matrix is a submatrix Γ' of Γ that corresponds to the variables α, β and Z . The statement $I(\alpha, Z, \beta)$ holds in P_σ iff $\det(\Gamma'_{\alpha\beta}) = 0$ where $\Gamma'_{\alpha\beta}$ is a submatrix of Γ' obtained by removing the α -th line and the β -th column (Miller, 1964). Both Γ' and $\Gamma'_{\alpha\beta}$ are given below. The matrix Γ' is a tri-diagonal matrix whose off main-diagonal elements are integer powers of ρ or zeros. The columns and lines of Γ' correspond to the following order of variables: $\alpha, z_1, \dots, z_k, \beta$ and then all other variables of Z (see fig 3), thus, for example, the term ρ^{i_1} located at $(1,2)$ in Γ' is the correlation factor between α and z_1 because these variables are the first two in the above order. The integer i_1 is the length of the path between α and z_1 , i_2 is the length of the path between z_1 and z_2 and so on. The construction of $\Gamma'_{\alpha\beta}$ is based on the observation that the location (α, β) in Γ' is $(1, k+2)$ where k is the number of head-to-head nodes of q .

$$\Gamma' = \begin{bmatrix} 1 & \rho^{i_1} & 0 & 0 & 0 \\ \rho^{i_1} & 1 & \rho^{i_2} & 0 & 0 \\ 0 & \rho^{i_2} & 1 & \rho^{i_3} & 0 \\ 0 & 0 & \rho^{i_3} & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \quad \Gamma'_{\alpha\beta} = \begin{bmatrix} \rho^{i_1} & 1 & \rho^{i_2} & 0 \\ 0 & \rho^{i_2} & 1 & \rho^{i_3} \\ 0 & 0 & \rho^{i_3} & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

(These matrices are given for the case of two head-to-head nodes and a single additional isolated node of Z , their general form is obvious). Clearly, $\det(\Gamma'_{\alpha\beta}) = \rho^k$ ($k \geq 0$) and therefore choosing $\rho \neq 0$ guarantees that $I(\alpha, Z, \beta)$ does not hold in P_σ .

It remains to show that every graphically-verified statement in D_σ is satisfied by P_σ . We assign a total order d on the nodes of D_σ consistent with the partial order determined by D_σ . We show that the n statements that form the causal input list that defines D_σ are satisfied in P_σ . Theorem 1 ensures that all other graphically verified statements are valid consequences of this input list and therefore would all be satisfied in P_σ . In what follows we use the tag of a node as its name. Let $I(i, S_i, U_{(i)}-S_i)$ be a statement of L where S_i are the parents of i and $U_{(i)}$ are all the variables preceding i in the order d . By the topology of D_σ , S_i contains no more than two nodes.

Assume S_i is empty. This implies that i is not connected via a regular path to either of its predecessors. Hence, by the construction of Γ , $\rho_{ij} = 0$ for every $j \in U_{(i)}$, and therefore the statements $I(i, \emptyset, j)$ hold in P_σ . However, in normal distributions, the correlation between single variables determines the dependency between the sets containing these variables because the following axiom holds.

(Composition-Decomposition)

$$I(X, Z, Y \cup W) \iff I(X, Z, Y) \& I(X, Z, W) \quad (4)$$

Accordingly, $I(i, \emptyset, U_{(i)})$ holds in P_σ and this statement is exactly equal to $I(i, S_i, U_{(i)}-S_i)$ since S_i is empty.

Assume S_i consists of a single node h . In light of axiom (4) it is enough to show that for every $j \in U_{(i)}-h$ the statement $I(i, h, j)$ holds in P_σ . If exists a path from j to i , it must pass through h . Therefore, by definition (3) of Γ , since h is the only parent of i , the equality $\rho_{ij} = \rho_{ih} \rho_{hj}$ must be satisfied. This equality is a necessary and sufficient condition under which $I(i, h, j)$ holds in any normal distribution in which ρ_{ij} are the correlation factors (in particular, P_σ).

Assume $S_i = \{g, h\}$ (see figure 6). Again, it is enough to show that $I(i, \{g, h\}, j)$ holds in P_σ for every $j \in U_{(i)}-S_i$. Construct the covariance matrix for the variables g, h, i and j (the columns of the matrix correspond to this order). By equation (3), $\rho_{gi} = \rho$, $\rho_{hi} = \rho$ and $\rho_{gh} = 0$. The resulting matrix is given below.

$$\Gamma = \begin{bmatrix} 1 & 0 & \rho & \rho_{gj} \\ 0 & 1 & \rho & \rho_{hj} \\ \rho & \rho & 1 & \rho_{ij} \\ \rho_{gj} & \rho_{hj} & \rho_{ij} & 1 \end{bmatrix}$$




Figure 6

The statement $I(i, \{g, h\}, j)$ holds in this distribution iff the submatrix Γ_{ij} ' is singular, i.e.

$$\det \begin{bmatrix} 1 & 0 & \rho \\ 0 & 1 & \rho \\ \rho_{hj} & \rho_{hj} & \rho_{ij} \end{bmatrix} = 0$$

([Miller, 1964]) or equivalently, $\rho_{ij} = (\rho_{hj} + \rho_{hj}) \cdot \rho$.

The latter equality, however, holds for all possible selections of a node j ; If j is not connected to i via a regular path, i.e. $\rho_{ij} = 0$, then it is not connected through a regular path to either of i 's parents and therefore both ρ_{hj} and ρ_{hj} are zero. If j is connected through a regular path of length l to g (similarly when connected to h) then it is connected to its son i with a path of length $l+1$ and is not connected to i 's other parent, in which case $\rho_{ij} = \rho^{l+1}$, $\rho_{hj} = \rho^l$, $\rho_{hj} = 0$ and therefore the above equality holds. \square