A Polynomial Time Algorithm For Determining DAG Equivalence in the Presence of Latent Variables and Selection Bias

by Peter Spirtes (Department of Philosophy, Carnegie Mellon University, ps7z@andrew.cmu.edu)
and Thomas Richardson (Department of Statistics, University of Washington)

Following the terminology of Lauritzen et. al. (1990) say that a probability measure over a set of variables **V** satisfies the **local directed Markov property** for a directed acyclic graph (DAG) *G* with vertices **V** if and only if for every *W* in **V**, *W* is independent of the set of all its non-descendants conditional on the set of its parents. One natural question that arises with respect to DAGs is when two DAGs are "statistically equivalent". One interesting sense of "statistical equivalence" is "d-separation equivalence" (explained in more detail below.) In the case of DAGs, d-separation equivalence is also corresponds to a variety of other natural senses of statistical equivalence (such as representing the same set of distributions). Theorems characterizing d-separation equivalence for directed acyclic graphs and that can be used as the basis for polynomial time algorithms for checking d-separation equivalence were provided by Verma and Pearl (1990), and Frydenberg (1990). The question we will examine is how to extend these results to cases where a DAG may have latent (unmeasured) variables or selection bias (i.e. some of the variables in the DAG have been conditioned on.) D-separation equivalence is of interest in part because there are algorithms for constructing DAGs with latent variables and selection bias that are based on observed conditional independence relations. For this class of algorithms, it is impossible to determine which of two d-separation equivalent causal structures generated a given probability distribution, given only the set of conditional independence and dependence relations true of the observed distribution. We will describe a polynomial (in the number of vertices) time algorithm for determining when two DAGs which may have latent variables or selection bias are d-separation equivalent.

A DAG *G* **entails a conditional independence relation** if and only if it is true in every probability measure satisfying the local directed Markov property for *G*. (We place definitions and sets of variables in boldface.) Pearl, Geiger, and Verma (Pearl 1988) have shown that there is a graphical relation, d-separation, that holds among three disjoint sets of variable **A**, and **B**, and **C** in DAG *G* if and only if *G* entails that **A** is independent of **B** given **C**. A vertex Y is a **collider** on an undirected path U if U contains a subpath X $\rightarrow$Y $\leftarrow$ Z. Say that a vertex V on an undirected path *U* between X and Y is **active** on *U* given **Z** (**Z** not containing X and Y) if and only if either V is not a collider on *U* and not in **Z**, or V is a collider on *U* and is an ancestor of **Z**. For three disjoint sets of variables **A**, **B**, and **C**, **A** is **d-connected** to **B** given **C** in graph *G*, if and only if there is an undirected path from some member of **A** to a member of **B** such that every vertex on *U* is active given **C**; for three disjoint sets of variables **A**, **B**, and **C**, **A** is **d-separated** from **B** given **C** in graph *G*, if and only **A** is not d-connected to **B** given **C**.

Two DAGs are **d-separation equivalent** if and only if they have the same vertices, and they have the same d-separation relations. If two DAGs G_1 and G_2 are d-separation equivalent, then they entail the same set of conditional independence relations, and the set of distributions that satisfy the local directed Markov property for G_1 equals the set of distribution that satisfy the local directed Markov property for G_2. Theorems that provide the basis for polynomial time algorithms for testing d-separation equivalence for DAGs was given in Verma and Pearl (1990), for cyclic directed graphs in Richardson (1994), and for directed acyclic graphs with latent variables in Spirtes and Verma (1992).

DAGs are also used to represent causal processes. Under this interpretation, a directed edge from A to B means that A is a direct cause of B relative to the variables in the DAG. Suppose a causal process represented by DAG *G* generates some population with a given distribution $P(\mathbf{V})$ that satisfies the local directed Markov property for *G*. If some of the variables in **V** are unmeasured, and some have been conditioned on (due to those variables being causally related to the sampling mechanism) then the set of conditional independence relations entailed for the subset of measured variables in the subpopulation from which the sample is drawn is not necessarily equal to the set of conditional independence relations entailed by any DAG (without latent variables or selection bias). Assume then that the variables in **V** can be partitioned into **O** (observed), **L** (latent), and **S** (selected, or conditioned on.) In that case instead of observing $P(\mathbf{V})$, we may be able to observe only $P(\mathbf{O}|\mathbf{S} = \mathbf{1})$, that is the marginal distribution over the observed variables in the selected subpopulation. Let us call $P(\mathbf{O}|\mathbf{S} = \mathbf{1})$ the "observed" distribution. There

are algorithms which, under some plausible assumptions relating probability distributions to causal processes, are correct in the large sample limit, and that can construct a representation of the class of DAGs (that may have latent variables and variables conditioned on) that are compatible with the observed conditional independence relations. See Spirtes et al. 1993 for the latent variable case without selection bias.

For a given DAG G, and a partition of the variable set $\mathbf{V}$ of G into observed ($\mathbf{O}$), selection ($\mathbf{S}$), and latent ($\mathbf{L}$) variables, we will write $G(\mathbf{O},\mathbf{S},\mathbf{L})$. Let us now extend the definition of d-separation equivalence to the case where there may be latent variables and selection bias. Two directed graphs $G_1(\mathbf{O},\mathbf{L},\mathbf{S})$ and $G_2(\mathbf{O'},\mathbf{L'},\mathbf{S'})$ are **d-separation equivalent** if and only if $\mathbf{O} = \mathbf{O'}$, and for all $\mathbf{X}$, $\mathbf{Y}$ and $\mathbf{Z}$ included in $\mathbf{O}$, $\mathbf{X}$ and $\mathbf{Y}$ are d-separated given $\mathbf{Z} \cup \mathbf{S}$ in $G_1(\mathbf{O},\mathbf{L},\mathbf{S})$ if and only if $\mathbf{X}$ and $\mathbf{Y}$ are d-separated given $\mathbf{Z} \cup \mathbf{S'}$ in $G_2(\mathbf{O'},\mathbf{L'},\mathbf{S'})$. Intuitively, the conditional independence relations true in the observed distribution could have been generated either by the causal DAG $G_1(\mathbf{O},\mathbf{L},\mathbf{S})$ or by $G_2(\mathbf{O'},\mathbf{L'},\mathbf{S'})$. Information just about the observed conditional independence relations cannot distinguish any two DAGs which are d-separation equivalent.

In order to state necessary and sufficient conditions for d-separation equivalence, we will need the following concept. A mixed ancestral graph (MAG) is an extended graph consisting of a set of vertices $\mathbf{V}$, and a set of edges between vertices, where there may be the following kinds of edges: A $\leftrightarrow$ B, A o—o B, A o$\rightarrow$ B, A $\leftarrow$o B, A $\rightarrow$ B, or A $\leftarrow$ B. (A MAG may be considered a special case of a PAG that represents a single graph. See Richardson 1996.) We say that the A endpoint of an A $\rightarrow$ B edge is "–"; the A endpoint of an A $\leftrightarrow$ B, A $\leftarrow$o B, or. A $\leftarrow$ B edge is "<"; and the A endpoint of an A o—o B or A o$\rightarrow$ B edge is "o". The conventions for the B endpoints are analogous. A mixed ancestral graph for a set of graphs $\mathbf{G}$ each sharing the same set of observed variable $\mathbf{O}$ contains information about the ancestor relations in $\mathbf{G}$, namely only those ancestor relations common to all members of $\mathbf{G}$. In the following definition, which provides a semantics for MAGs we use "*" as a meta-symbol indicating the presence of any one of {o, –, >}, e.g. A *$\rightarrow$ B represents either A $\rightarrow$ B, or A $\leftrightarrow$ B.

Mixed Ancestral Graphs (MAGs)

A MAG represents directed acyclic graph $G(\mathbf{O},\mathbf{S},\mathbf{L})$ (in which case we write MAG($G(\mathbf{O},\mathbf{S},\mathbf{L})$) if:

(i) If A and B are in $\mathbf{O}$, there is an edge between A and B in MAG($G(\mathbf{O},\mathbf{S},\mathbf{L})$) if and only for any subset $\mathbf{W} \subseteq \mathbf{O}\backslash\{A,B\}$, A and B are d-connected given $\mathbf{W} \cup \mathbf{S}$ in $G(\mathbf{O},\mathbf{S},\mathbf{L})$.

(ii) There is an edge A $\rightarrow$ B or B $\leftarrow$ A in MAG($G(\mathbf{O},\mathbf{S},\mathbf{L})$) if and only if A is an ancestor of B but not $\mathbf{S}$ in $G(\mathbf{O},\mathbf{S},\mathbf{L})$;

(iii) There is an edge A $\leftarrow$* B or B *$\rightarrow$ A in MAG($G(\mathbf{O},\mathbf{S},\mathbf{L})$) if and only if A is **not** an ancestor of B or $\mathbf{S}$ in $G(\mathbf{O},\mathbf{S},\mathbf{L})$;

(iv) There is an edge A o—* B or B *—o A in MAG($G(\mathbf{O},\mathbf{S},\mathbf{L})$) if and only if A is an ancestor of $\mathbf{S}$ in $G(\mathbf{O},\mathbf{S},\mathbf{L})$.

The definition of "d-separation" given for DAGs can be applied directly to MAGs, as long as such concepts as "undirected path", "collider", etc., are given their obvious extensions to apply to MAGs. We include in the Appendix the definitions of terms such as "undirected path" etc. which apply both to directed graphs and MAGs.

The first step in forming a MAG for a graph is to form the ancestor matrix for the graph. Let n be the number of vertices in $\mathbf{O} \cup \mathbf{S} \cup \mathbf{L}$ and m the number of vertices in $\mathbf{O}$. Aho, Hopcroft, and Ullman (1974) describes a transitive closure algorithm for filling in such a matrix that is $O(n^3)$. Then each pair of vertices X and Y in $\mathbf{O}$ ($O(m^2)$) is adjacent in MAG($G_1(\mathbf{O},\mathbf{S},\mathbf{L})$) if and only if they are not d-separated ($O(n^2)$) given ($\mathbf{Ancestors}(\{X,Y\} \cup \mathbf{S}) \cap \mathbf{O}$) in $G_1(\mathbf{O},\mathbf{S},\mathbf{L})$ (where $\mathbf{Ancestors}(\mathbf{Z})$ is the set of vertices which have descendants in Z; see Lemma 5.) The orientation of each edge in the MAG ($O(m^2)$) then can be determined by examining the ancestor matrix. So forming a MAG is $O(n^3m)$.

If U is an acyclic undirected path containing X and B, and X is before B on U, then U(X,B) represents the unique subpath of U between X and B. If B is before X on U, by definition $U(X,B) = U(B,X)$. In MAG($G(\mathbf{O},\mathbf{S},\mathbf{L})$), U is a **discriminating path** for B if and only if U is an undirected path between X and Y with at least three edges, U contains B, B $\neq$ X, B adjacent to Y on U, X is not adjacent to Y, and for every vertex Q on U(X,B) except for the endpoints Q is a collider on U(X,B) and there is an edge Q $\rightarrow$ Y in

MAG(G(**O**,**S**,**L**)). If Y is adjacent to X and Z on a path U, and X and Z are not adjacent in thegraph, then Y is **unshielded**.

MAG(G_1(**O**,**S**,**L**)) and MAG(G_2(**O**,**S'**,**L'**)) have the **same basic colliders** if and only if they have (i) the same adjacencies; (ii) the same unshielded colliders (iii) if U is a discriminating path for X in MAG(G_1(**O**,**S**,**L**)), and the corresponding path U' in MAG(G_2(**O**,**S'**,**L'**)) is a discriminating path for X, then X is a collider on U in MAG(G_1(**O**,**S**,**L**)) if and only if X is a collider on U' in MAG(G_2(**O**,**S'**,**L'**)).

Theorem 1: DAGs G_1(**O**,**S**,**L**) and G_2(**O**,**S'**,**L'**) are d-separation equivalent if and only if MAG(G_1(**O**,**S**,**L**)) and MAG(G_2(**O**,**S'**,**L'**)) have the same basic colliders.

Theorem 1 is the basis of an $O(n^3m^2)$ algorithm for determining d-separation equivalence, where n is the maximum number of vertices in G_1(**O**,**S**,**L**) and G_2(**O**,**S'**,**L'**), and m is the number of vertices in **O**. The first step in determining d-separation equivalence is to form MAG(G_1(**O**,**S**,**L**)) and MAG(G_2(**O**,**S'**,**L'**)), which is $O(n^3m)$. Checking that the two MAGs have the same unshielded colliders is $O(m^3)$, and for each triple of vertices all of which are adjacent to each other, there is a simple algorithm that determines whether there is a discriminating path that examines each edge $O(m^2)$ in the MAG at most once. Hence, overall the algorithm is $O(n^3m^2)$.

Appendix

For our purposes we need to represent a variety of marks attached to the ends of edges. In general, we allow that the end of an edge can be marked out of by "—", or can be marked with ">", or can be marked with an "o". In order to specify completely the type of an edge, therefore, we need to specify the variables and **marks** at each end. For example, the left end of "A o→ B" can be represented as the ordered pair [A, o] and the right end can be represented as the ordered pair [B, >]. We will also call [A, o] the A end of the edge between A and B. The first member of the ordered pair is called an endpoint of an edge, e.g. in [A, o] the endpoint is A. The entire edge is a set of ordered pairs representing the endpoints, e.g. {[A, o], [B, >]}. Note that the edge {[B, >],[A, o]} is the same as {[A, o],[B, >]} since it doesn't matter which end of the edge is listed first. Note that a directed edge such as A → B has a mark "—" at the A end.

We say a **graph** is an ordered triple <**V**,**M**,**E**> where **V** is a non-empty set of vertices, **M** is a non-empty set of marks, and **E** is a set of sets of ordered pairs of the form {[V_1,M_1],[V_2,M_2]}, where V_1 and V_2 are in **V**, $V_1 \neq V_2$, and M_2 and M_2 are in **M**. If G = <**V**,**M**,**E**> we say that G is **over V**. (Directed graphs and MAGs are both special cases of graphs.)

In a graph, for a directed edge A → B, the edge is **out of** A and **into** B, and A is **parent** of B and B is a **child** of A. A sequence of edges <E_1,...,E_n> in G is an **undirected path** if and only if there exists a sequence of vertices <V_1,...,V_{n+1}> such that for $1 \leq i \leq n$ E_i has endpoints V_i and V_{i+1}, and $E_i \neq E_{i+1}$. A path U is **acyclic** if no vertex appears more than once in the corresponding sequence of vertices. We will assume that an undirected path is acyclic unless specifically mentioned otherwise. A sequence of edges <E_1,...,E_n> in G is a **directed path** D **from** V_1 to V_n if and only if there exists a sequence of vertices <V_1,..., V_{n+1}> such that for $1 \leq i \leq n$, there is a directed edge $V_i \rightarrow V_{i+1}$ on D. If there is an acyclic directed path from A to B or B = A then A is an **ancestor** of B, and B is a **descendant** of A. If **Z** is a set of variables, A is an **ancestor** of **Z** if and only if it is an ancestor of a member of **Z**, and similarly for **descendant**. If **X** is a set of vertices in G, let **Ancestors**(**X**) be the set of all ancestors of members of **X** in G(**O**,**S**,**L**). A vertex V is a **collider** on an undirected path U if and only if U contains a pair of distinct edges adjacent on the path and into V. The **orientation** of an acyclic undirected path between A and B is the set consisting of the A end of the edge on U that contains A, and the B end of the edge on U that contains B. Say that a vertex V on an undirected path U between X and Y is **active** on U given **Z** (**Z** not containing X and Y) if and only if either V is not a collider on U and not in **Z**, or V is a collider on U and is an ancestor of **Z**. For three disjoint sets of variables **A**, **B**, and **C**, **A** is **d-connected** to **B** given **C** in graph G, if and only if there is an undirected path from some member of **A** to a member of **B** such that every vertex on U is active given **C**; for three disjoint sets of variables **A**, **B**, and **C**, **A** is **d-separated** from **B** given **C** in graph G, if and only **A** is not d-connected to **B** given **C**.

In a **directed graph**, all of the edges are directed edges. A directed graph is **acyclic** if and only if it contains no directed cyclic paths. Lemma 1 is a simple generalization of Lemma 3.3.1 in Spirtes *et al.* (1993).

Lemma 1: In a directed acyclic graph G over a set of vertices **V**, if the following conditions hold:

(a) R is a sequence of vertices in **V** from A to B, $R \equiv \langle A \equiv X_0, \ldots X_{n+1} \equiv B \rangle$, such that $\forall i,\ 0 \leq i \leq n,\ X_i \neq X_{i+1}$ (the X_i are only *pairwise distinct* , i.e. not necessarily distinct),

(b) $\mathbf{Z} \subseteq \mathbf{V} \backslash \{A,B\}$,

(c) **T** is a set of undirected paths such that

(i) for each pair of consecutive vertices in R, X_i and X_{i+1}, there is a unique undirected path in **T** that d-connects X_i and X_{i+1} given $\mathbf{Z} \backslash \{X_i , X_{i+1}\}$,

(ii) if some vertex X_k in R is in **Z**, then the paths in **T** that contain X_k as an endpoint collide at X_k, (i.e. all such paths are directed into X_k)

(iii) if for three vertices X_{k-1}, X_k, X_{k+1} occurring in R, the d-connecting paths in **T** between X_{k-1} and X_k, and X_k and X_{k+1}, collide at X_k then X_k has a descendant in **Z**,

then there is a path U in G that d-connects $A \equiv X_0$ and $B \equiv X_{n+1}$ given **Z** that contains only edges occurring in **T**.

U is an **inducing path** between X and Y in $G(\mathbf{O},\mathbf{S},\mathbf{L})$, if and only U is an acyclic undirected path such that every member of $\mathbf{O} \cup \mathbf{S}$ on U is a collider on U, and every collider on U is an ancestor of $\{X,Y\} \cup \mathbf{S}$. (This is a generalization of the concept of inducing path that was introduced in Verma and Pearl 1990). The following sequence of lemmas state that for every subset **W** of **O**, X and Y are d-connected given $\mathbf{W} \cup \mathbf{S}$ in $G(\mathbf{O},\mathbf{S},\mathbf{L})$ if and only there is an inducing path between X and Y in $G(\mathbf{O},\mathbf{S},\mathbf{L})$. For space reasons we do not present the proofs here, but they are simple modifications of the proofs that appear in Spirtes et al. (1993), in which the case of latent variables without selection bias is considered. (There is no analog of Lemma 4 in Spirtes *et al.* 1993, but the proof is very similar to that of Lemma 2 and Lemma 3.)

Lemma 2: In directed graph $G(\mathbf{O},\mathbf{S},\mathbf{L})$, if there is an inducing path between A and B that is out of A and into B, then for any subset **Z** of $\mathbf{O} \backslash \{A,B\}$ there is an undirected path C that d-connects A and B given $\mathbf{Z} \cup \mathbf{S}$ that is out of A and into B.

Lemma 3: If $G(\mathbf{O},\mathbf{S},\mathbf{L})$ is a directed acyclic graph, and there is an inducing path U between A and B that is into A and into B then for every subset **Z** of $\mathbf{O} \backslash \{A,B\}$ there is an undirected path C that d-connects A and B given $\mathbf{Z} \cup \mathbf{S}$ that is into A and into B.

Lemma 4: If $G(\mathbf{O},\mathbf{S},\mathbf{L})$ is a directed acyclic graph, and there is an inducing path U between A and B that is out of A and out of B then for every subset **Z** of $\mathbf{O} \backslash \{A,B\}$ there is an undirected path C that d-connects A and B given $\mathbf{Z} \cup \mathbf{S}$.

Lemma 5: If $G(\mathbf{O},\mathbf{S},\mathbf{L})$ is a directed acyclic graph and an undirected path U in $G(\mathbf{O},\mathbf{S},\mathbf{L})$ d-connects A and B given $((\mathbf{Ancestors}(\{A,B\} \cup \mathbf{S}) \cap \mathbf{O}) \cup \mathbf{S}) \backslash \{A,B\}$ then U is an inducing path between A and B.

The following lemma follows from a simple application of d-separation to discriminating paths.

Lemma 6: In $\mathrm{MAG}(G(\mathbf{O},\mathbf{S},\mathbf{L}))$, if U is a discriminating path for B between X and Y, and B is a collider on U then B is no set that d-separates X and Y, and if B is not a collider on U, B is in every set that d-separates X and Y.

Note that it follows directly from the definition of $\mathrm{MAG}(G(\mathbf{O},\mathbf{S},\mathbf{L}))$ that there are no edges A o—B in $\mathrm{MAG}(G(\mathbf{O},\mathbf{S},\mathbf{L}))$, and if there is an edge A o—* B in $\mathrm{MAG}((G(\mathbf{O},\mathbf{S},\mathbf{L}))$, then the A endpoint of every edge in $\mathrm{MAG}(G(\mathbf{O},\mathbf{S},\mathbf{L}))$ is "o". Hence if A is a collider on any path in $\mathrm{MAG}(G(\mathbf{O},\mathbf{S},\mathbf{L}))$, the A endpoint of no edge in $\mathrm{MAG}(G(\mathbf{O},\mathbf{S},\mathbf{L}))$ is a "o".

Let $l(U,C_i,\mathbf{Z})$ be the length of a shortest directed path from collider C_i on U to a member of **Z**. Let U be a **minimal d-connecting path** between X and Y given **Z** if and only if U is a d-connecting path between X and Y given **Z**, and there is no other path V d-connecting X and Y given **Z** such that either V has fewer edges than U, or V has the same number of edges as U and the sum over j of $l(V,D_j,\mathbf{Z})$ is less than the sum over i of $l(U,C_i,\mathbf{Z})$. Say that two undirected paths U and U' which contain a vertex C **disagree** at C if C is a collider on U but not on U', or vice-versa.

Lemma 7: If U is a minimal d-connecting path between A and B given **R** in $\mathrm{MAG}(G(\mathbf{O},\mathbf{S},\mathbf{L}))$, and E is an edge between C and D in U, but C and D are not adjacent on U, and U' is the result of substituting E in for $U(C,D)$ in U, then either $U(C,D)$ is into C and E is out of C, or $U(C,D)$ is into D, and E is out of D.

Proof. If C and D are both active on *U'*, then *U'* d-connects A and B given **R**, and is shorter than *U*, contradicting the assumption. Hence either C is not active on *U'*, or D is not active on *U'*. If *U'* agrees with *U* at C and D, then C and D are both active on *U'*. Hence *U'* disagrees with *U* at C or D.

If C is a collider on *U'*, but not on *U*, it follows that *E* is an edge D *→ C, there is an edge M *→ C on *U*, and *U*(C,D) is out of C. If there is no collider on *U*(C,D) then C is either an ancestor of D, or an ancestor of a vertex with a "o" endpoint. If C is an ancestor of a vertex with a "o" endpoint, then C is an ancestor of **S** in *G*(**O**,**S**,**L**), and hence there cannot be a D *→ C edge in MAG(*G*(**O**,**S**,**L**)). If C is an ancestor of D, this contradicts the D *→ C edge. It follows that there is a collider on *U*(C,D) and hence C is an ancestor of the first collider on *U*(C,D). It follows that C is an ancestor of **R**. Hence C is active on *U'*.

Similarly, if D is a collider on *U'* but not on *U*, D is active on *U'*. It follows that either C is a collider on *U* but not on *U'*, or *D* is a collider on *U* but not on *U'*. Hence either *U*(C,D) is into C and *E* is out of C, or *U*(C,D) is into D, and *E* is out of D. ∴

Lemma 8: If *U* is a minimal d-connecting path between A and B given **R** in MAG(*G*(**O**,**S**,**L**)), *U* contains C *—* F *—* D, and C and D are adjacent in MAG(*G*(**O**,**S**,**L**)), then MAG(*G*(**O**,**S**,**L**)) contains one of the following subgraphs:

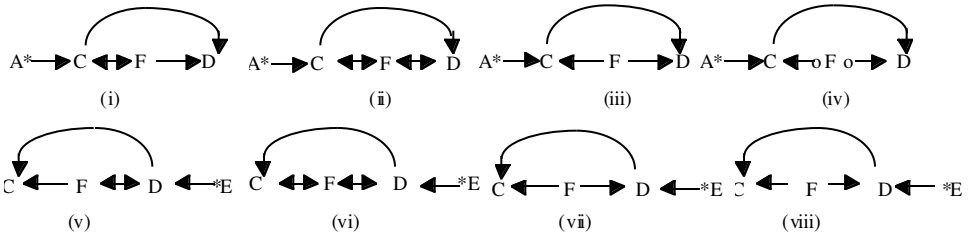

Figure 1

Proof. Let *E* be the edge between C and D, and *U'* be the result of substituting *E* in for *U*(C,D). By Lemma 7, either *U*(C,D) is into C and *E* is out of C, or *U*(C,D) is into D, and *E* is out of D. Suppose first that E is out of C and *U*(C,D) is into C. Then MAG(*G*(**O**,**S**,**L**)) contains either (i), (ii), (iii) or (iv) of Figure 1 or one of the following subgraphs (ix), (x), or (xi) of Figure 2:

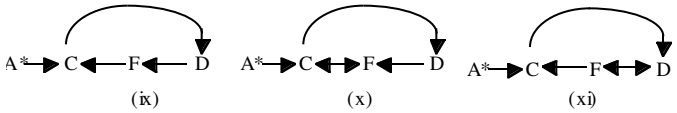

Figure 2

However (ix) contains a cycle. (x) is impossible because C ↔ F implies C is not an ancestor of F in *G*(**O**,**S**,**L**), but there is a path C → D → F in MAG(*G*(**O**,**S**,**L**)), and hence a directed path from C to F in *G*(**O**,**S**,**L**). (xi) is impossible because F ↔ D implies F is not an ancestor of D in *G*(**O**,**S**,**L**), but there is a path F → C → D in MAG(*G*(**O**,**S**,**L**)), and hence a directed path from F to D in *G*(**O**,**S**,**L**).

Similarly, it can be shown that if there is an edge D → C, the only possible subgraphs are (v), (vi), (vii), and (viii). ∴

Lemma 9: If *U* is a minimal d-connecting path between X and Y given **R** in MAG(*G*(**O**,**S**,**L**)), *U* contains the subpath A *→ B ←* D *→ C, and MAG(*G*(**O**,**S**,**L**)) contains the edge B → C, then *U* contains a unique subpath *U*(F,C) that is a discriminating path for D.

Proof. We will show that for each n ≥ 1, if *U* contains a vertex M such that *U*(M,D) is of length n, for every vertex Q on *U*(M,D) except for the endpoints Q is a collider on *U*(M,D), and for each vertex Q on *U*(M,D) except possibly for D there is an edge Q → C in MAG(*G*(**O**,**S**,**L**)), then *U* contains a vertex F such that *U*(F,D) is of length n + 1, and either *U*(F,D) is a discriminating path for D, or *U* contains an edge F ↔ M, and MAG(*G*(**O**,**S**,**L**)) contains an edge F → C.

By hypothesis, there is a path *U*(B,D) of length 1 such that every vertex Q between B and D is a collider on *U* and there is an edge B → C in MAG(*G*(**O**,**S**,**L**)).

Suppose that U contains a vertex M such that U(M,D) is of length n, and for every vertex Q on U(M,D) except for the endpoints Q is a collider on U(M,D) and for each vertex Q on U(M,D) except possibly for D there is an edge Q $\rightarrow$ C in MAG(G(**O**,**S**,**L**)). Because U is a minimal d-connecting path, by Lemma 7 there is an edge F *$\rightarrow$ M on U. The edge between F and M is either F $\rightarrow$ M, F o$\rightarrow$ M, or F $\leftrightarrow$ M. If U(F,C) is not a discriminating path for D then there is an edge between F and C in MAG(G(**O**,**S**,**L**)). By Lemma 7, there are two cases: (i) the edge between F and C is out of C (and hence C $\rightarrow$ F), or (ii) the edge between F and C is F $\rightarrow$ C, and the edge between F and M is F $\leftrightarrow$ M.

If (i), there is no edge F o$\rightarrow$ M because C $\rightarrow$ F is into F. If (i) and the edge between F and M isF $\rightarrow$ M, then there is a cycle F $\rightarrow$ M $\rightarrow$ C and C $\rightarrow$ F. If the edge between F and M is F $\leftrightarrow$ M then there is a contradiction because F is not an ancestor of M but there is a path M $\rightarrow$ C $\rightarrow$ F. It follows that case (ii) holds. If U does not contains a discriminating subpath, then for every subpath of U there is a longer subpath of U. Because U is of finite length, it follows that it contains a discriminating subpath U(M,C) for D.

We will now show that U is unique. Because the edge between B and C is oriented as B $\rightarrow$ C, M lies between X and C. No subpath of U(M,C) is a discriminating path for D because all of the vertices on U(M,C) except for M are adjacent to C. No path containing U(M,C) is a discriminating path for D because M is not adjacent to C. ∴

In MAG(G(**O**,**S**,**L**)), a vertex V is a **hidden vertex** on a discriminating path U if and only if there are vertices X and Y on U such that V is adjacent to X and Y on U, and X and Y are adjacent in MAG(G(**O**,**S**,**L**)).

Lemma 10: In MAG(G(**O**,**S**,**L**)), if U is a minimal d-connecting path between X and Y given **Z**, then there is no pair of distinct vertices B, J such that B is a hidden vertex on the discriminating path U(I,K) for J, and J is a hidden vertex on the discriminating path U(A,C) for B.

Proof. Suppose on the contrary that B is a hidden vertex on U(I,K), and J is a hidden vertex on U(A,C). Because B is hidden on U(I,K), C lies on U(I,K). C ≠ K because B lies on U(I,K), the only vertex adjacent to K on U(I,K) is J, and B ≠ J. C ≠ J because otherwise J is not a hidden vertex on U(A,C). C ≠ I because otherwise J, which is on U(A,C) but not equal to B or A is an ancestor of C = I, and hence by repeated applications of Lemma 1 through Lemma 5 there is an inducing path between I and K in G(**O**,**S**,**L**). But by definition of discriminating path I and K are not adjacent in MAG(G(**O**,**S**,**L**)). Hence C ≠ I. Because K is on U(A,C) but not equal to A, B, or C, there is a directed path from K to C. Similarly, there is a directed path from C to K. Hence, there is a directed cycle in MAG(G(**O**,**S**,**L**)), which is a contradiction. ∴

Lemma 11: In a MAG(G(**O**,**S**,**L**)), if U is a minimal d-connecting path between X and Y given **Z**, then there is no triple of distinct hidden vertices X, Y, Z on U such that X is a hidden vertex on the discriminating path for Y on U, Y is a hidden vertex on the discriminating path for Z on U, and Z is between Y and X on U.

Proof. Let U_Y be the discriminating path for Y on U, and similarly for U_Z. Because X is a hidden vertex on the discriminating path for Y on U, every vertex between X and Y is on the discriminating path for Y on U. Hence Z is on the discriminating path for Y on U. Because Z is between Y and X, and neither Y nor X is an endpoint of U_Y, Z is not an endpoint of U_Y; it follows that each of the vertices adjacent to Z on U are also on U_Y. Because Z is a hidden vertex on U, and both of the vertices adjacent to Z are also on the discriminating path for Y on U, Z is a hidden vertex on U_Y. By hypothesis, Y is a hidden vertex on U_Z. But this contradicts Lemma 10. ∴

Lemma 12: In a MAG(G(**O**,**S**,**L**)), if U is a minimal d-connecting path between X and Y given **Z**, then there is no quadruple of distinct hidden vertices A_i, A_{r+1}, A_{i+1}, A_r in that order on U such that A_i is a hidden vertex on the discriminating path for A_{i+1} on U, and A_r is a hidden vertex on the discriminating path for A_{r+1} on U.

Proof. Let U_{i+1} be the discriminating path for A_{i+1} on U, and similarly for U_{r+1}. Suppose contrary to the hypothesis there is a quadruple of distinct hidden vertices A_i, A_{r+1}, A_{i+1}, A_r in that order on U such that A_i is a hidden vertex on the discriminating path for A_{i+1} on U, and A_r is a hidden vertex on the discriminating path for A_{r+1} on U. A_{r+1} is on U_{i+1} because it is between A_i and A_{i+1}, and A_i is on U_{i+1}. A_{r+1} is a hidden vertex on U, and both of the vertices adjacent to A_{r+1} on U are also on U_{i+1}, because A_{r+1} is not an endpoint of U_{i+1}.

Hence A_{r+1} is a hidden vertex on U_{i+1}. Similarly, A_{i+1} is a hidden vertex on U_{r+1}. But this contradicts Lemma 10.

Lemma 13: In a MAG(G(**O**,**S**,**L**)), if U is a minimal d-connecting path between X and Y given **Z**, then there is no sequence of length greater than 1 of distinct vertices $<A_1,A_2,...,A_n>$ such that for each pair of vertices A_i, A_{i+1} that are adjacent in the sequence, A_i is a hidden vertex on the discriminating path of A_{i+1} on U, and A_n is a hidden vertex on the discriminating path of A_1 on U. (Note that the subscripts of the vertices do not necessarily reflect the order in which they occur on U.)

Proof. Suppose without loss of generality that n is greater than 1, and A_1 is to the right of A_n on U. Let r be the highest index such that A_r is to the right of A_1 if such a vertex exists; otherwise let r = 1. We will now show that A_{r+1} is to the left of A_n. If r = 1, every vertex except A_1 is to the left of A_1, so A_2 is to the left of A_1, and $A_2 \neq A_n$ by Lemma 10. By Lemma 11 then A_2 is not between A_1 and A_n, so it is to the left of A_n. If r ≠ 1, then $A_{r+1} \neq A_n$ by Lemma 11, and A_{r+1} is not between A_1 and A_n by Lemma 12. Hence A_{r+1} is to the left of A_n.

We will now show that some vertex A_{s+1} whose index is greater than r+1 is to the right of A_r. A_{r+2} is not between A_{r+1} and A_r by Lemma 11, and hence not equal to A_n. There are two cases. If A_{r+2} is to the right of A_r, then we are done. Suppose then that A_{r+2} is to the left of A_{r+1}. It follows that there is some vertex with index greater than r+2 (e.g. A_n) on the other side of A_{r+1}. It follows that for some s > r such that A_s and A_{s+1} are on opposite sides of A_{r+1} (where A_s is to the left of A_{r+1}.) A_{s+1} is not between A_{r+1} and A_r by Lemma 12, and hence $A_{s+1} \neq A_n$. So A_{s+1} is to the right of A_r. But this is a contradiction, because r is the highest index such that A_r is to the right of A_1 ∴

We will now recursively define the order of a discriminating path for a hidden variable on a minimal d-connecting path. If U is a minimal d-connecting path between X and Y given **Z**, and W is a hidden variable on U such that the discriminating path for W on U contains no hidden variables other than W, then W is a **0-order hidden variable** on U. If U is a minimal d-connecting path between X and Y given **Z**, and W is a hidden variable on U such that the maximum order of any other hidden variable on the discriminating path for W on U is n−1, then W is an **n^{th}-order hidden variable** on U.

Lemma 13 guarantees that this recursive definition is sound, because it guarantees that if U is a minimal d-connecting path between X and Y given **Z** that contains hidden variables, then there is a 0 order hidden variable on U and also that the definition of the order of any hidden variable W on U is not defined in terms of the order of W.

Lemma 14: If there is an edge A *→ B in MAG(G(**O**,**S**,**L**)), then in G(**O**,**S**,**L**) there is an inducing path between A and B that is into B.

Proof. By the definition of a MAG and Lemma 5 there is an inducing path U between A and B in G(**O**,**S**,**L**), and B is not an ancestor of A or **S** in G(**O**,**S**,**L**). Suppose that U is out of B. If there are no colliders on U, then U is a directed path from B to A, and B is an ancestor of A, which is a contradiction. If there is a collider C on U closest to B, and it is an ancestor of B, A, or **S**. If C is an ancestor of B then there is a cycle in G(**O**,**S**,**L**) which is a contradiction. If C is an ancestor of A or **S**, then B is an ancestor of A or **S** which is a contradiction. Hence U is into B. ∴

Lemma 15: If MAG(G_1(**O**,**S**,**L**)) and MAG(G_2(**O**,**S'**,**L'**)) have the same basic colliders, U is a discriminating path between X_1 and Y for F in MAG(G_1(**O**,**S**,**L**)), U' is the path corresponding to U in MAG(G_2(**O**,**S'**,**L'**)), and every vertex (except for the endpoints and possibly F) is a collider on U', then U' is a discriminating path for F in MAG(G_2(**O**,**S'**,**L'**)).

Proof. Suppose that the vertices on U preceding F are X_1, ..., X_n. By definition, in MAG(G_1(**O**,**S**,**L**)) X_1 is not adjacent to Y, X_2 is a collider on U, and X_2 is an unshielded non-collider on the concatenation of $U(X_1,X_2)$ and the edge $X_2 \rightarrow$ Y. Hence in MAG(G_2(**O**,**S'**,**L'**)), X_1 is not adjacent to Y, X_2 is a collider on U', and X_2 is an unshielded non-collider on the concatenation of $U'(X_1,X_2)$ and the edge between X_2 and Y. It follows that the edge between X_2 and Y is oriented as $X_2 \rightarrow$ Y in MAG(G_2(**O**,**S'**,**L'**)).

Suppose for each X_i, $2 \leq i \leq m-1$, in MAG(G_2(**O**,**S'**,**L'**)) X_i is a collider on U', and the edge between X_i and Y is oriented as $X_i \rightarrow$ Y. In MAG(G_2(**O**,**S'**,**L'**)), let V' be the concatenation of $U'(X_1,X_{m-1})$ and the edge between X_{m-1} and Y. Every vertex on V' between X_1 and X_m is a collider by hypothesis, and for each X_i between X_1 and X_m, there is an edge $X_i \rightarrow$ Y by hypothesis. Hence V' is a discriminating path for X_m. If V

is the corresponding path in MAG(G_1(**O**,**S**,**L**)), V is a discriminating path, and X_m is a non-collider on V. Hence X_m is a non-collider on V'. It follows that the edge between X_m and Y is oriented as $X_m \rightarrow Y$. By induction, U' is a discriminating path for F in MAG(G_2(**O**,**S'**,**L'**)). ∴

Lemma 16: If MAG(G_1(**O**,**S**,**L**)) and MAG(G_2(**O**,**S'**,**L'**)) have the same basic colliders, U is a minimal d-connecting path between X and Y given **Z** in MAG(G_1(**O**,**S**,**L**)), U' is the corresponding path in MAG(G_2(**O**,**S'**,**L'**)), then F is a collider on U if and only if F is a collider on U'.

Proof. If F is not a hidden vertex on U, then because MAG(G_1(**O**,**S**,**L**)) and MAG(G_2(**O**,**S'**,**L'**)) have the same basic colliders, F is not a hidden vertex on U', and by definition F is a collider on U if and only if F is a collider on U'.

Suppose F is a hidden vertex on U. By Lemma 8, MAG(G(**O**,**S**,**L**)) contains one of the subgraphs of type (i) through (viii) in Figure 1. By Lemma 9, U contains a discriminating path U(M,N) for F.

Suppose first that F is a zero order hidden vertex on U. Then all of the vertices on U(M,F) except for the endpoints are unshielded colliders in MAG(G_1(**O**,**S**,**L**)). Because MAG(G_1(**O**,**S**,**L**)) and MAG(G_2(**O**,**S'**,**L'**)) have the same basic colliders, all of the vertices on U(M,F) are unshielded colliders. Hence by Lemma 15, U'(A,B) is a discriminating path in MAG(G_2(**O**,**S'**,**L'**)). It follows that F is a collider on U if and only if F is a collider on U'.

Suppose that for all $0 \leq i < n$, the i^{th} order hidden vertices on U are oriented the same way on U'. Now consider an n^{th} order hidden vertex on U. There is a subpath U(M,N) that is a discriminating path for F. By the induction hypothesis, all of the colliders on U(M,N) are colliders on U'(M,N). Hence by Lemma 15, U'(A,B) is a discriminating path in MAG(G_2(**O**,**S'**,**L'**)). It follows that F is a collider on U if and only if F is a collider on U'. ∴

Lemma 17: If A and B are d-connected given **R** in MAG(G(**O**,**S**,**L**)), then A and B are d-connected given **R** ∪ **S** in G(**O**,**S**,**L**).

Proof. Suppose that in MAG(G(**O**,**S**,**L**)) A and B are d-connected given **R** by a minimal d-connecting path U. Then each vertex on U is active given **R**. For each edge X *—* Y in MAG(G(**O**,**S**,**L**)), there is an inducing path between X and Y in G(**O**,**S**,**L**). By Lemma 2, Lemma 3, and Lemma 4, there is a path that d-connects X and Y given **R** ∪ **S**\{X,Y} in G(**O**,**S**,**L**). Choose all such d-connecting path for each pair of vertices X and Y adjacent on U; call this collection of d-connecting paths T . If a vertex X is on U, say that X is **active in** T **given R** ∪ **S** whenever either (i) there are vertices C and D on U adjacent to X, there is a path in T between X and C that is into X, there is a path in T between X and D that is into X, and X is an ancestor of **R** ∪ **S** in G(**O**,**S**,**L**), or (ii) there is a d-connecting path in T containing X as an endpoint that is out of X, and X is not in **R**. Consider the following three cases.

U contains a subpath C *—o F o—* D, and F is active on U given **R**. Hence there is a path X_1 in T that d-connects C and F given (**R** ∪ **S**)\{C,F}, and a path X_2 in T that d-connects F and D given (**R** ∪ **S**)\{F,D}. F is not in **R** because F is active on U given **R**, and F is not a collider on U. F is active in T given **R** ∪ **S** if X_1 and X_2 collide at F because it is an ancestor of **S** in G(**O**,**S**,**L**), and is active in T if X_1 and X_2 do not collide at F because it is not in **R** ∪ **S**, so F is active in T given **R** ∪ **S**.

U contains a subpath C *→ F ←* D, and F is active on U given **R**. It follows that in MAG(G(**O**,**S**,**L**)) F has a descendant in **R**. Hence F has a descendant in **R** in G(**O**,**S**,**L**). By Lemma 14 there is an inducing path between C and F that is into F, and an inducing path between D and F that is into F. It follows from Lemma 2 and Lemma 3 that there is a path X_1 in T that d-connects C and F given (**R** ∪ **S**)\{C,F} that is into F, and a path X_2 in T that d-connects F and D given (**R** ∪ **S**)\{F,D} that is into F. F is active in T given **R** ∪ **S** because X_1 and X_2 collide at F in G(**O**,**S**,**L**), and F is an ancestor of **R** in G(**O**,**S**,**L**).

U contains a subpath C *—* F → D, and F is active on U given **R**. (The case where U contains a subpath C ← F *—* D is analogous.) Because F is active on U given **R**, F is not in **R**. There is a directed path from F to D in G(**O**,**S**,**L**) that does not contain any vertices in **S**. There are two cases. If the directed path contains a member of **R**, then F is an ancestor of **R**, and hence F is active in T given **R** ∪ **S** regardless of whether or not the d-connecting paths collide at F. If the directed path does not contain a member of **R**, the directed path d-connects F and D given **R** ∪ **S** and is out of F. It follows that F is active in T given **R** ∪ **S**. It follows from Lemma 1 that there is a path in G(**O**,**S**,**L**) that d-connects A and B given **R** ∪ **S**. ∴

Lemma 18: If X and Y are d-connected given $\mathbf{Z} \cup \mathbf{S}$ in $G(\mathbf{O},\mathbf{S},\mathbf{L})$, then X and Y are d-connected given $\mathbf{Z}$ in MAG($G(\mathbf{O},\mathbf{S},\mathbf{L})$).

Proof. Suppose that U is a minimal d-connecting path between X and Y given $\mathbf{Z} \cup \mathbf{S}$ in $G(\mathbf{O},\mathbf{S},\mathbf{L})$. We will perform a series of operation which show how to construct a path U' in MAG($G(\mathbf{O},\mathbf{S},\mathbf{L})$) which d-connects A and B given $\mathbf{Z}$ in MAG($G(\mathbf{O},\mathbf{S},\mathbf{L})$). The operations are illustrated with Figure 3 ($\mathbf{Z} = O_2,O_3$).

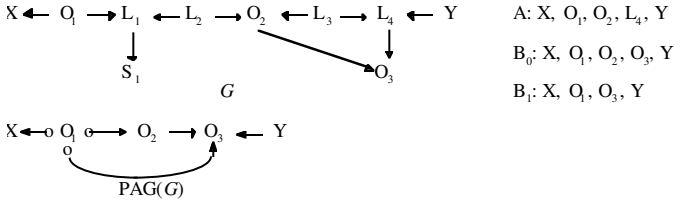

Figure 3

First form the following sequence of vertices. A(0) = X. If A(n−1) ≠ Y, then A(n) is the first vertex on U after A(n−1) such that either it is a non-collider on U and is in $\mathbf{O}$, or it is a collider on U and not an ancestor of $\mathbf{S}$. The last vertex in the sequence is Y because Y is in $\mathbf{O}$ and not a collider on U. Suppose the length of the sequence is n, i.e. A(n) = Y.

Note that for $1 \le i \le n$, if A(i) is a collider on U then it is an ancestor of $\mathbf{Z}$, because U d-connects X and Y given $\mathbf{Z} \cup \mathbf{S}$.

Suppose A(i) is in $\mathbf{O}$, but not a collider on U. Then for $1 < i \le n-1$, either U(A(i−1),A(i)) or U(A(i),A(i+1)) is out A(i). Suppose without loss of generality that U(A(i−1),A(i)) is out of A(i). Then A(i) is an ancestor of $\mathbf{S}$ or A(i−1) because either U(A(i−1),A(i)) contains no colliders in which case A(i) is an ancestor of A(i−1), or it does contain a collider, in which case the first collider is an ancestor of a member of $\mathbf{S}$, and A(i−1) is an ancestor of the first collider. Similarly, if U(A(i),A(i+1)) is out A(i) then A(i) is an ancestor of A(i+1) or $\mathbf{S}$. So if A(i) is in $\mathbf{O}$, but not a collider on U, then A(i) is an ancestor of A(i−1), A(i+1) or $\mathbf{S}$.

Now form the sequence of vertices where for $1 \le i \le n$, $B_0(i) = A(i)$ if A(i) in $\mathbf{O}$, and otherwise $B_0(i) = O_i$, where O_i is the first vertex in $\mathbf{O}$ on a shortest path from A(i) to $\mathbf{Z}$. (Such a path exists because U d-connects X and Y given $\mathbf{Z} \cup \mathbf{S}$, and no A(i) that is a collider on U is an ancestor of $\mathbf{S}$.) We will now show that for $1 \le i \le n$, there is an inducing path between $B_0(i)$ and $B_0(i+1)$.The path between A(i) and A(i+1) d-connects A(i) and A(i+1) given (**Ancestors**($\{B_0(i),B_0(i+1)\} \cup \mathbf{S}$) ∩ $\mathbf{O}$) ∪ $\mathbf{S}$)\{A(i),A(i+1)}, because every collider on U(A(i),A(i+1)) is an ancestor of $\mathbf{S}$, and no non-collider on U(A(i),A(i+1)) is in $\mathbf{O} \cup \mathbf{S}$. The path from A(i) to O_i (if there is one) d-connects A(i) and O_i given (**Ancestors**($\{B_0(i),B_0(i+1)\} \cup \mathbf{S}$) ∩ $\mathbf{O}$) ∪ $\mathbf{S}$)\{A(i),O_i} because by construction it is a directed path that contains no member of $\mathbf{O}$ except for O_i, and no member of $\mathbf{S}$. Similarly, the path from A(i+1) to O_{i+1} (if there is one) d-connects A(i+1) and O_{i+1} given (**Ancestors**($\{B_0(i),B_0(i+1)\} \cup \mathbf{S}$) ∩ $\mathbf{O}$) ∪ $\mathbf{S}$)\{A(i+1),O_{i+1}}. By Lemma 1 $B_0(i)$ and $B_0(i+1)$ are d-connected given **Ancestors**($\{B_0(i),B_0(i+1)\} \cup \mathbf{S}$) ∩ $\mathbf{O}$) ∪ $\mathbf{S}$)\{A(i), A(i+1}, and by Lemma 5 there is an inducing path between $B_0(i)$ and $B_0(i+1)$. Because for $1 \le i \le n-1$ there is an inducing path between $B_0(i)$ and $B_0(i+1)$, there is a path B_0 in MAG($G(\mathbf{O},\mathbf{S},\mathbf{L})$) on which $B_0(i)$ is the i^{th} vertex.

If $B_0(i)$ is not a non−collider on U, any edge that contains $B_0(i)$ on any of the paths used to construct the inducing path between $B_0(i)$ and $B_0(i+1)$ or $B_0(i-1)$, is into $B_0(i)$. Hence the inducing path between $B_0(i)$ and $B_0(i+1)$ and the inducing path between $B_0(i)$ and $B_0(i-1)$ are into $B_0(i)$.

If $B_0(i)$ is on U but not a collider on U, then $B_0(i)$ is not in $\mathbf{Z} \cup \mathbf{S}$ because U d-connects X and Y given $\mathbf{Z} \cup \mathbf{S}$. In addition, either $B_0(i)$ is an ancestor of $B_0(i-1)$, $B_0(i+1)$ or $\mathbf{S}$, because $B_0(i) = A(i)$, A(i) is an ancestor of A(i−1), A(i+1) or $\mathbf{S}$, and A(i−1) is an ancestor of $B_0(i-1)$, and A(i+1) is an ancestor of $B_0(i+1)$. It follows that $B_0(i)$ is not a collider on B_0, and is not in $\mathbf{Z} \cup \mathbf{S}$.

By construction, if $B_0(i)$ is not a non-collider on U, $B_0(i)$ is in $\mathbf{O}$, an ancestor of $\mathbf{Z}$, and not an ancestor of $\mathbf{S}$. However, $B_0(i)$ may be in $\mathbf{Z}$ but not be a collider on B_0 in MAG($G(\mathbf{O},\mathbf{S},\mathbf{L})$) (if in $G(\mathbf{O},\mathbf{S},\mathbf{L})$ it is an ancestor of either its predecessor or successor on U). The following algorithm removes all non-colliders on B_0 which are in $\mathbf{Z}$.

k = 0;

Repeat

If there is a triple of vertices $B_k(i-1)$, $B_k(i)$, $B_k(i+1)$ such that the inducing paths between $B_k(i-1)$ and $B_k(i)$, and $B_k(i)$ and $B_k(i+1)$ collide at $B_k(i)$, but $B_k(i)$ is in **Z** and an ancestor of $B_k(i-1)$ or $B_k(i+1)$, form sequence B_{k+1} by removing $B_k(i)$ from the sequence (i.e. for $1 \leq j < i$, set $B_{k+1}(j) = B_k(j)$, and for $i \leq j \leq n-1$, set $B_{k+1}(j) = B_k(j+1)$);

k := k + 1;

until there is no such triple of vertices in the sequence B_k.

At each stage of the algorithm, if there is a triple of vertices $B_k(i-1)$, $B_k(i)$, $B_k(i+1)$ such that the inducing paths between $B_k(i-1)$ and $B_k(i)$, and $B_k(i)$ and $B_k(i+1)$ collide at $B_k(i)$, but $B_k(i)$ is an ancestor of $B_k(i-1)$ or $B_k(i+1)$, then by Lemma 1 through Lemma 5 there is an inducing path between $B_k(i-1)$ and $B_k(i+1)$. Hence for every k and each i, $1 \leq i \leq$ length of sequence B_k, there is an inducing path between $B_k(i)$ and $B_k(i+1)$. It follows that there is path B_k in MAG(G(**O**,**S**,**L**)) such that the i^{th} vertex on the path is $B_k(i)$.

Suppose first that $B_k(i)$ is a non-collider on U. We will show $B_k(i)$ is a non-collider on B_k in MAG(G(**O**,**S**,**L**)), and not in $\mathbf{Z} \cup \mathbf{S}$. We have already shown this for B_0. In addition, we have shown that if $B_0(i)$ is a non-collider on U either $B_0(i)$ is an ancestor of $B_0(i-1)$, $B_0(i+1)$ or **S.** Suppose for $1 \leq i \leq$ length(B_{k-1}), if $B_{k-1}(i)$ is a non-collider on U either $B_{k-1}(i)$ is an ancestor of $B_{k-1}(i-1)$, $B_{k-1}(i+1)$ or **S.** It is an ancestor of $B_k(i-1)$, $B_k(i+1)$ or **S**, unless $B_{k-1}(i)$ is an ancestor of $B_{k-1}(i-1)$, and $B_{k-1}(i-1)$ was removed at the k^{th} step of the algorithm, or $B_{k-1}(i)$ is an ancestor of $B_{k-1}(i+1)$, and $B_{k-1}(i+1)$ was removed at the k^{th} step of the algorithm. Suppose without loss of generality that $B_{k-1}(i)$ is an ancestor of $B_{k-1}(i+1)$, and $B_{k-1}(i+1)$ was removed at the k^{th} step of the algorithm. It follows that $B_{k-1}(i+1)$ is an ancestor of $B_{k-1}(i+2)$, $B_{k-1}(i)$, or **S**. $B_{k-1}(i+1)$ is not an ancestor of $B_{k-1}(i)$ because G(**O**,**S**,**L**) is acyclic. Hence, it follows that it is an ancestor of either $B_{k-1}(i+2)$ or **S**. It follows that $B_k(i) \equiv B_{k-1}(i)$ is an ancestor of **S** or $B_{k-1}(i+2) \equiv B_k(i+1)$. Hence $B_k(i)$ is not a collider on B_k, and not in $\mathbf{Z} \cup$ **S.**

Suppose that $B_k(i)$ is not a non-collider on U. If k = 0, then $B_k(i)$ is in **O** and an ancestor of **Z**, and for every other value of k, B_k is a subsequence of B_0, so $B_k(i)$ is in in **O** and an ancestor of **Z** if it is not a non-collider on U. We will now show that the inducing paths between $B_k(i)$ and $B_k(i-1)$, and $B_k(i)$ and $B_k(i+1)$ are both into $B_k(i)$. We have already shown that the inducing path between $B_0(i)$ and $B_0(i+1)$ and the inducing path between $B_0(i)$ and $B_0(i-1)$ are into $B_0(i)$. Suppose the same is true for each $B_{k-1}(i)$ that is not a non-collider on U. This will also be true for each $B_k(i)$ unless $B_{k-1}(i)$ is an ancestor of $B_{k-1}(i-1)$, and the vertex $B_{k-1}(i-1)$ was removed at the k^{th} step of the algorithm, or $B_{k-1}(i)$ is an ancestor of $B_{k-1}(i+1)$, and the vertex $B_{k-1}(i+1)$ was removed at the k^{th} step of the algorithm. Suppose without loss of generality that $B_{k-1}(i)$ is an ancestor of $B_{k-1}(i+1)$, and $B_{k-1}(i)$ was removed at the k^{th} step of the algorithm. The edge on the inducing path between $B_{k-1}(i)$ and $B_{k-1}(i+1)$ that contains $B_{k-1}(i)$ is into $B_{k-1}(i)$ and so is every edge on the inducing path between $B_{k-1}(i+1)$ and $B_{k-1}(i+2)$ that contains $B_{k-1}(i)$ (because it is in **O**.) $B_{k-1}(i)$ is not on any directed path from $B_{k-1}(i+1)$ to $B_{k-1}(i+2)$ because by hypothesis, $B_{k-1}(i)$ is an ancestor of $B_{k-1}(i+1)$. Since all of the edges that contain $B_{k-1}(i)$ on the paths used to construct the inducing path between $B_{k-1}(i)$ and the vertex $B_{k-1}(i+2)$ are into $B_{k-1}(i)$, the inducing path between $B_{k-1}(i) \equiv B_k(i)$ and $B_{k-1}(i+2) \equiv B_k(i+1)$ is into $B_k(i)$.

If $B_k(i)$ is not a non-collider on U, then the inducing paths between $B_k(i)$ and $B_k(i-1)$, and $B_k(i)$ and $B_k(i+1)$ are both into $B_k(i)$. It follows that if the algorithm exits at stage k, each $B_k(i)$ that is not a non-collider on U is not an ancestor of either $B_k(i-1)$ or $B_k(i+1)$. Hence it is a collider on B_k, and is a member of **O** that is an ancestor of **Z**.

Hence each vertex on B_k is active, and hence B_k d-connects X and Y given **Z** in MAG(G(**O**,**S**,**L**)). ∴

Lemma 19: If MAG(G(**O**,**S**,**L**)) contains A *→ B → H, and an edge A *—* H, then (i) the edge between A and H is into H, and (ii) if A *—* H has a different orientation at A than A *→ B, then the edges are oriented as A → H and A ↔ B.

Proof. Because there is an edge into H, A *—* H is not oriented as A *—o H. Because there is an edge A *→ B, B is not an ancestor of A. If the edge between A and H is oriented as A ← H, then B is an ancestor of A, which is a contradiction. Hence the edge between A and H is into H. If the edge A *→ H has a different orientation at A than A *→ B, then either A ↔ B and A → H, or A → B and A ↔ H. If A → B and A ↔ H, then A is an ancest or of H (A → B → H) which contradicts A ↔ H. ∴

Lemma 20: If MAG(G_1(**O**,**S**,**L**)) and MAG(G_2(**O**,**S'**,**L'**)) have the same basic colliders, U is a minimal d-connecting path between X and Y given **Z** in MAG(G_1(**O**,**S**,**L**)), F is a collider on U, H is an ancestor of **Z** and there is an F → H edge in MAG(G_1(**O**,**S**,**L**)), then there is an F → H edge in MAG(G_2(**O**,**S'**,**L'**)).

Proof. If F is a collider on U, by Lemma 16 both MAG(G_1(**O**,**S**,**L**)) and MAG(G_2(**O**,**S'**,**L'**)) contain X_0 *→ F ←* Y_0. If there is no edge between X_0 and H in MAG(G_1(**O**,**S**,**L**)), then F is an unshielded non-collider on X_0 *→ F *—* H in MAG(G_1(**O**,**S**,**L**)), and hence F is an unshielded non-collider on X_0 *→ F *—* H in MAG(G_2(**O**,**S'**,**L'**)). It follows that F *—* H is oriented as F o→ H or F → H in MAG(G_2(**O**,**S'**,**L'**)). It is not oriented as F o→ H because F is a collider on U and hence not an ancestor of **S**. It follows that the edge is oriented as F → H. Similarly, if there is no edge between Y_0 and H, F *—* H is oriented as F → H in MAG(G_2(**O**,**S'**,**L'**)). Suppose then that X_0 and Y_0 are both adjacent to H.

There is a vertex N on U between X and F that is either (i) not adjacent to H, or (ii) the edge between N and H is not into H, or (iii) U agrees with the concatenation of U(X,N) and the edge between N and H at N (since X itself trivially satisifes condition (iii) if it is adjacent to H.) Similarly, there is a vertex M on U between Y and F that is either (i) not adjacent to H, or (ii) the edge between M and H is not into H, or (iii) U agrees with the concatenation of U(Y,M) and the edge between M and H at M. Let X_{n+1} be the closest such vertex on U to F, and Y_{o+1} be the closest such vertex on U to F. Let X_n through X_0 be the be the vertices on U between X_{n+1} and F.

We will now show by induction that for $0 \le i \le n$, the edge between X_i and its successor on U is into X_i, and there is an edge X_i → H in MAG(G_1(**O**,**S**,**L**)). If $X_0 = X_{n+1}$, then it is trivially true (because there are no vertices between X_{n+1} and F). Suppose X_0 is between X_{n+1} and F. We have already shown that there is an edge X_0 *→ F, and an edge F → H. By definition of X_{n+1}, the edge between X_0 and H disagrees with the edge between X_0 and F at X_0. By Lemma 19 it follows that the edge between X_0 and F is X_0 → F, and the edge between X_0 and F is into X_0. Suppose for $0 \le i \le m-1$ that the edge between X_i and its successor on U is into X_i, and the edge between X_i and H is oriented as X_i → H. If X_m is between X_{n+1} and H, by definition of X_{n+1}, X_m is adjacent to H and U disagrees with the concatenation of U(X,X_{m-1}) and the edge between X_{m-1} and H at X_{m-1}; hence the edge between X_m and X_{m-1} is X_m *→ X_{m-1}. By lemma 19, the edge between X_m and X_{m-1} is into X_m, and there is an edge X_m → H. Hence every vertex X_i between X_{n+1} and F is a collider on U, and there is an edge X_i → H. Similarly, every vertex Y_i between Y_{o+1} and F is a collider on U, and there is an edge Y_i → H.

If X_{n+1} is adjacent to H, then by Lemma 19 X_{n+1} *—* is into H, and by definition of X_{n+1}, U agrees with the concatenation of U(X,X_{n+1}) and the edge between X_{n+1} and H at X_{n+1}. Similarly, if If Y_{o+1} is adjacent to H, then by Lemma 19 Y_{o+1} *—* is into H, and by definition of Y_{o+1}, U agrees with the concatenation of U(Y_{o+1},Y) and the edge between Y_{o+1} and H at Y_{o+1}. In that case, the concatenation of U(X,X_{n+1}), the edge between X_{n+1} and H, the edge between H and Y_{o+1}, and U(Y_{o+1},Y) d-connects X and Y given **Z**, and is shorter than U. This is a contradiction. It follows that either X_{n+1} or Y_{o+1} is not adjacent to H. Suppose without loss of generality that it is the former.

If X_{n+1} is not adjacent to H there is a path V between X_{n+1} and H consisting of the concatenation of U(X_{n+1},F) with the edge between F and H. Let V' be the corresponding path in MAG(G_2(**O**,**S'**,**L'**)). By definition, V is a discriminating path for F in MAG(G_1(**O**,**S**,**L**)), and F is a non-collider on the path. Because all of the colliders on V are also colliders on U which is minimal, by Lemma 16 they are all colliders on V'. By Lemma 15, V' is a discriminating path for F. Hence F is a non-collider on V' in MAG(G_2(**O**,**S'**,**L'**)), and the edge between F and H is oriented as F → H in MAG(G_2(**O**,**S'**,**L'**)). ∴

Lemma 21: If MAG(G_1(**O**,**S**,**L**)) and MAG(G_2(**O**,**S'**,**L'**)) have the same basic colliders, U is a minimal d-connecting path between X and Y given **Z** in MAG(G_1(**O**,**S**,**L**)), A is a collider on U and B is a member of **Z** that is the endpoint of a shortest path D from A to **Z** in MAG(G_1(**O**,**S**,**L**)), then B is a descendant of A in MAG(G_2(**O**,**S'**,**L'**)).

Proof. Let D' be the path corresponding to D in MAG(G_2(**O**,**S'**,**L'**)), and U' be the path corresponding to U in MAG(G_2(**O**,**S'**,**L'**)). By Lemma 16, A is a collider on U'. By Lemma 20, the first edge on D' is out of A. If D' is not a directed path in MAG(G_2(**O**,**S'**,**L'**)) then it contains a collider F. Because D does not contain a collider, and MAG(G_1(**O**,**S**,**L**)) and MAG(G_2(**O**,**S'**,**L'**)) have the same basic colliders, F is a shielded collider in MAG(G_2(**O**,**S'**,**L'**)). It follows then that in MAG(G_1(**O**,**S**,**L**)) there is a vertex E and D contains a subpath E → F → H, and an edge between E and H. The edge between E and H is not oriented as

H $\rightarrow$ E, else MAG(G_1(**O**,**S**,**L**)) contains a cycle; it is nor oriented as E $\leftrightarrow$ H because E is an ancestor of H in MAG(G_1(**O**,**S**,**L**)); it is neither E *—o H nor E o—* H, because A then would be an ancestor of **S** in G_1(**O**,**S**,**L**) and hence not a collider on *U*. Hence it is oriented as E $\rightarrow$ H. But then *D* is not a shortest directed path from A to a member of **Z**, contrary to our assumption. $\therefore$

Theorem 1: DAGs G_1(**O**,**S**,**L**) and G_2(**O**,**S'**,**L'**) are d-separation equivalent if and only if MAG(G_1(**O**,**S**,**L**)) and MAG(G_2(**O**,**S'**,**L'**)) have the same basic colliders.

Proof. If X and Y are adjacent in MAG(G_1(**O**,**S**,**L**)) but not in MAG(G_2(**O**,**S'**,**L'**)), then for some subset **V** of **O**, X and Y are d-separated given **V** $\cup$ **S** in G_2(**O**,**S'**,**L'**), but not d-separated given **V** $\cup$ **S'** in G_2(**O**,**S'**,**L'**). If C *$\rightarrow$ F $\leftarrow$* D is an unshielded collider in MAG(G_1(**O**,**S**,**L**)) but not in MAG(G_2(**O**,**S'**,**L'**)), then every set that d-separates C and D in MAG(G_1(**O**,**S**,**L**)) does not contain F, but every set that d-separates C and D in MAG(G_2(**O**,**S'**,**L'**)) does contain F. If *U* is a discriminating path between X and Y for F in MAG(G_1(**O**,**S**,**L**)) and the corresponding path *U'* is a discriminating path for F in MAG(G_2(**O**,**S'**,**L'**)), and F is a collider on *U* but not on *U'*, then by Lemma 6 there is a set **Z** that contains F that d-separates X and Y in MAG(G_2(**O**,**S'**,**L'**)) but not in MAG(G_1(**O**,**S**,**L**)).

If MAG(G_1(**O**,**S**,**L**)) and MAG(G_2(**O**,**S'**,**L'**)) have the same basic colliders, then by Lemma 16 and Lemma 21, MAG(G_1(**O**,**S**,**L**)) and MAG(G_2(**O**,**S'**,**L'**)) have the same d-separation relations. By Lemma 17 and Lemma 18, X and Y are d-separated given **R** in MAG(G_1(**O**,**S**,**L**)) if and only if X and Y are d-separated given **R** $\cup$ **S** in G_1(**O**,**S**,**L**). Similarly, X and Y are d-separated given **R** in MAG(G_2(**O**,**S'**,**L'**)) if and only if X and Y are d-separated given **R** $\cup$ **S'** in G_2(**O**,**S'**,**L'**). It follows that G_1(**O**,**S**,**L**) and G_2(**O**,**S'**,**L'**) are d-separation equivalent. $\therefore$

Bibliography

Aho, A., Hopcroft, J., Ullman D., (1974) The Design and Analysis of Computer Algorithms, Addison-Wesley: Reading, MA.

Frydenberg, M., (1990), The Chain Graph Markov Property, Scandinavian Journal of Statistics, 17, 333-353.

Lauritzen, S., Dawid, A., Larsen, B., Leimer, H., (1990) Independence properties of directed Markov fields, Networks, 20, 491-505.

Pearl, J., (1988). Probabilistic Reasoning in Intelligent Systems, Morgan Kaufman: San Mateo, CA.

Richardson, T. (1994b). Properties of Cyclic Graphical Models. MS Thesis, Carnegie Mellon University.

Richardson, T. (1996). A discovery algorithm for directed cyclic graphs. Uncertainty in Artificial Intelligence, Proceedings, 12th Conference, Morgan Kaufman, CA.

Spirtes, P., Verma, T. (1992) Equivalence of Causal Models with Latent Variables. Technical Report CMU-PHIL-33, Department of Philosophy, Carnegie Mellon University, October, 1992.

Spirtes, P., Glymour, C., and Scheines, R., (1993) Causation, Prediction, and Search, (Springer-Verlag Lecture Notes in Statistics 81, New York).

Verma, T. and Pearl, J. (1990), Equivalence and synthesis of causal models, in Proc. Sixth Conference on Uncertainty in AI. Association for Uncertainty in AI, Inc., Mountain View, CA.