Causal diagrams for empirical research

Paloma Vaissman Uribe

October 20th, 2015

Judea Pearl: biographical sketch

Education

- B.S. degree in Electrical Engineering from the Technion, Haifa, Israel, in 1960;
- Master degree in Physics from Rutgers University, New Brunswick, New Jersey, in 1965;
- Ph.D. degree in Electrical Engineering from the Polytechnic Institute of Brooklyn, Brooklyn, NY in 1965.

Current Position

- UCLA Computer Science Department, Cognitive Systems Lab (joined in 1970).
- Before: worked at RCA Research Laboratories, Princeton, New Jersey, on superconductive parametric and storage devices and at Electronic Memories, Inc., Hawthorne, California, on advanced memory systems.

Judea Pearl: biographical sketch

A fair number of papers and authored 3 books:

- Heuristics, Addison-Wesley, 1984;
- Probabilistic Reasoning in Intelligent Systems, Morgan Kaufmann, 1988;
- Causality: Models, Reasoning, and Inference, Cambridge University Press, 2000; 2nd edition, 2009.

Pearl is the father of slain Wall Street Journal reporter Daniel Pearl (who was kidnapped and murdered by terrorists in Pakistan in early 2002, just four months after 9/11) and president of the Daniel Pearl Foundation.

Summary: Causal diagrams for empirical research

Pearl, 1995

- Show how graphical models can be used as a mathematical language for integrating statistical and subject-matter information.
- In particular, the paper develops a principled, nonparametric framework for causal inference, in which diagrams are queried to determine if the assumptions available are sufficient for identifying causal effects from nonexperimental data.

Classical example due to Cochran

The experiment:

- Experiment in which soil fumigants, X, are used to increase oat crop yields, Y, by controlling the eelworm population, Z.
- Note that X may also have direct effects, both beneficial and adverse, on yields beside the control of eelworms.

Classical case of counfounding bias:

- Controlled randomised experiments are infeasible: farmers insist on deciding for themselves which plots are to be fumigated;
- Farmers' choice of treatment depends on last year's eelworm population, Z_0, an unknown quantity strongly correlated with this year's population.

Classical example due to Cochran

- Fortunately, through laboratory analysis of soil samples, we can determine the eelworm populations (Z) **before and after the treatment**.
- Furthermore, because the fumigants are known to be active for a short period only, we can safely assume that they do not affect the growth of eelworms surviving the treatment.
- Instead, eelworm growth depends on the population of birds and other predators (B), which is correlated, in turn, with last year's eelworm population and hence with the treatment itself.

First step: construct a causal diagram

Z_1: the eelworm population, both size and type, **before** treatment;
Z_2: the eelworm population, both size and type, **after** treatment;
Z_3: the eelworm population, both size and type, **at the end of the season**;
Z_0: represents last year's eelworm population (because it is an unknown quantity, it is represented by a hollow circle);
B, the population of birds and other predators (also unknown).

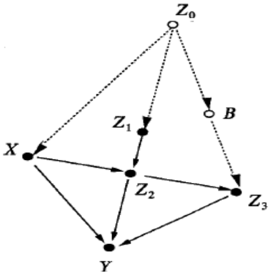

First step: construct a causal diagram

- Links in the diagram are of two kinds: those that connect unmeasured quantities are designated by dashed arrows, those connecting measured quantities by solid arrows.
- The missing arrow between Z_1 and Y signifies the investigator's understanding that pre-treatment eelworms cannot affect oat plants directly; their entire influence on oat yields is mediated by post-treatment conditions, namely Z_2 and Z_3.

Note

The purpose of the paper is not to validate or repudiate such domain-specific assumptions but, rather, to test whether a given set of assumptions is sufficient for quantifying causal effects from nonexperimental data, for example, estimating the total effect of fumigants on yields.

Conclusions based on the proposed method

(a) The total effect of X on Y can be estimated consistently from the observed distribution of X , Z_1, Z_2, Z_3 and Y.

(b) The total effect (assuming discrete variables throughout) is given by the formula:

$$pr(y|\check{x}) = \sum_{z_1}\sum_{z_2}\sum_{z_3} pr(y|z_2, z_3, x)p(z_2|z_1, x)\sum_{x'} pr(z_3|z_1, z_2, x')pr(z_1, x'), \quad (1)$$

where the symbol $\check{x}$, read 'x check', denotes that the treatment is set to level X = x by external intervention.

(c) Consistent estimation of the total effect would not be feasible if Y were confounded with Z_3; however, confounding Z_2 and Y will not invalidate the formula for $pr(y|\check{x})$.

Graphs and conditional independence

Conditional independence relationships implied by recursive product decompositions:

$$pr(x_1, ..., x_n) = \prod_i pr(x_i | pa_i), \quad (2)$$

where pa_i stands for the realisation of some subset of the variables that precede X, in the order $(X_1,..., X_n)$.

If we construct a directed acyclic graph (DAG) in which the variables corresponding to pa_i, are represented as the parents of X_i, then the independencies implied by the decomposition can be read off the graph using the **d-separation criterion**.

Definition: d-separation

Let X, Y and Z be three disjoint subsets of nodes in a DAG G, and let p be any path between a node in X and a node in Y, where by 'path' we mean any succession of arcs, regardless of their directions.

Then Z **is said to block** p if there is a node w on p satisfying one of the following two conditions:

Conditions for d-separation

- w has converging arrows along p, and neither w nor any of its descendants are in Z, or,
- w does not have converging arrows along p, and w is in Z.

Further, Z is said to d-separate X from Y, in G, if and only if, Z blocks every path from a node in X to a node in Y.

Definition: d-separation

'D' stands for dependence or directed. A path p is d-separated or blocked by a set of nodes Z if and only if (Pearl, 2000):

Conditions for d-separation

- p contains a **chain**, $u \leftarrow m \leftarrow v$, such that m is in Z, or
- p contains a **fork**, $u \leftarrow m \rightarrow v$, such that m is in Z, or
- p contains an **inverted fork (or collider)**, $u \rightarrow m \leftarrow v$, such that m is not in Z and no descendant of m is in Z.

Intuition: in **causal forks** and **causal chains** the two extreme variables are marginally dependent but become independent (blocked) once we condition on the middle variable.

Examples: d-separation

A causal chain: $X \to Z \to Y$

By (2):$P(X, Y, Z) = P(Y|Z)P(Z|X)P(X)$. But $P(X, Y, Z) = P(X, Y|Z)P(Z)$, so that $P(X, Y|Z) = \frac{P(Z|X)P(X)P(Y|Z)}{P(Z)} = P(X|Z)P(Y|Z)$.
(X and Y are independent, conditioning on Z, or they are d-separed)

A causal fork: $X \leftarrow Z \to Y$

By (2):$P(X, Y, Z) = P(Y|Z)P(X|Z)P(Z)$. But $P(X, Y, Z) = P(X, Y|Z)P(Z)$, so that $P(X, Y|Z) = P(X|Z)P(Y|Z)$.
(X and Y are independent, conditioning on Z, or they are d-separed)

Collider: $X \to Z \leftarrow Y$

By (2):$P(X, Y, Z) = P(X)P(Y)P(Z|X, Y)$. But $P(X, Y, Z) = P(X, Y|Z)P(Z)$, so that $P(X, Y|Z) = \frac{P(X)P(Y)P(Z|X,Y)}{P(Z)} \neq P(X|Z)P(Y|Z)$.
(X and Y are not independent, conditioning on Z, or they are not d-separed)

Graphs as models of interventions

Causal interpretation of directed graphs, based on nonparametric structural equations, which owes its roots to early works in econometrics.

In other words, each child-parent family in a directed graph G represents a deterministic function:

$$X_i = f_i(pa_i, \epsilon_i),\ i = (1, ..., n), \tag{3}$$

where pa_i, denote the parents of variable X_i, in G, and $\epsilon_i (1 \leq i \leq n)$ are mutually independent, arbitrarily distributed random disturbances.

Disturbances represent exogenous factors that the investigator chooses not to include in the analysis. If is judged to be influencing two or more variables, then must enter the analysis as an unmeasured, or **latent**, variable, to be represented in the graph by a **hollow node**.

Graphs as models of interventions

- This equational model is the nonparametric analogue of a structural equations model, with one exception: the functional form of the equations, as well as the distribution of the disturbance terms, **will remain unspecified**.

For the first example:

$Z_0 = f_0(\epsilon_0)$, $Z_2 = f_2(X, Z_1, \epsilon_2)$, $B = f_B(Z_0, \epsilon_B)$, $Z_3 = f_3(B, Z_2, \epsilon_3)$, $Z_1 = f_1(Z_0, \epsilon_1)$, $Y = f_Y(X, Z_2, Z_3, \epsilon_Y)$, $X = f_X(Z_0, \epsilon_X)$

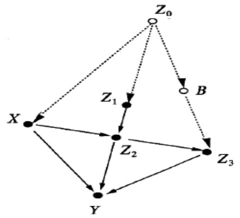

Fig. 1. A causal diagram representing the effect of fumigants, X, on yields, Y.

Graphs as models of interventions

- Characterising each child-parent relationship as a deterministic function, instead of by the usual conditional probability $pr(x|pa_i)$, imposes equivalent independence constraints on the resulting distributions, and leads to the same recursive decomposition that characterises DAG models (2).
- This occurs because each ϵ, is independent of all nondescendants of X.
- The functional characterisation $X_i = f_i(pa_i, \epsilon_i)$,also provides a convenient language for specifying how the resulting distribution would change in response to external interventions.

Definition: causal effect

- The simplest type of external intervention is one in which a single variable, say X_i, is forced to take on some fixed value x.
- Such an intervention, which we call atomic, amounts to lifting X_i from the influence of the old functional mechanism and placing it under the influence of a new mechanism that sets its value to x, while keeping all other mechanisms unperturbed.

Causal effect

- Given two disjoint sets of variables, X and Y, the causal effect of X on Y, denoted $pr(y|\check{x})$, is a function from X to the space of probability distributions on Y.

- For each realisation x of X , $pr(y|\check{x})$ gives the probability of $Y = y$ induced on deleting from the model (3) all equations corresponding to variables in X and substituting x for X in the remainder.

Graphical meaning of atomic intervention

In the case of an atomic intervention $set(X = x')$, this transformation can be expressed in a simple algebraic formula that follows immediately from (3) and Definition of Causal Effect:

$$pr(x_1, ..., x_n|\check{x}_i') = \begin{cases} pr(x_1, ..., x_n)/pr(x_i|pa_i) & if\ x_i = x_i' \\ 0 & if\ x_i \neq x_i' \end{cases} \tag{4}$$

This formula reflects the removal of the terms $pr(x_i|pa_i)$ from the product in (2), since pa_i no longer influence X.

Graphically, this is equivalent to **removing the links between** pa_i **and** X_i**, while keeping the rest of the network intact**.

Graphical meaning of atomic intervention

To better understanding go back to Pearl, 1993:

- The effect of the intervention $set(X_i = x_i')$ is encoded by adding to Γ a link $F_i \rightarrow X_i$, where F_i is a new variable taking values in $\{set(x_i), idle\}$, x_i' ranges over domain of X_i and *idle* represents no intervention.
- Thus the new parent set of X_i in the augmented network is $pa_i' = pa_i \cup \{F_i\}$ and is related to X_i by the conditional probability:

$$pr(x_i|pa_i') = \begin{cases} pr(x_i|pa_i), if\ F_i = idle, \\ 0, if\ F_i = set(x_i')\ and\ x_i \neq x_i', \\ 1, if\ F_i = set(x_i')\ and\ x_i = x_i'. \end{cases}$$

Graphical meaning of atomic intervention

The effect of the intervention $set(X_i = x_i')$ is to transform the original probability function $pr(x_1, ..., x_n)$ into a new function

$$pr(x_1, ..., x_n | \check{x}_i') = pr'(x_1, ..., x_n | F_i = set(x_i')),$$

where pr' is the directed markov field dictated by the augmented network $\Gamma' = \Gamma \cup \{F_i \to X_i\}$.

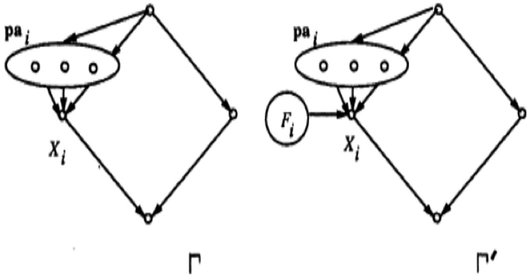

Graphical meaning of atomic intervention

The transformation exhibits the following properties:

- Intervention $set(X_i = x_i')$ can affect only the descendants of X_i in Γ.
- For any set S of variables, we have $pr(S|\check{x}_i', pa_i) = pr(S|x_i', pa_i)$

Intervention or passive?

In other words, given $X_i = x_i'$ and pa_i, it is superfluous to find out whether $X_i = x_i'$ was established by external intervention or not. This can be seen directly by the augmented network Γ', since $\{X_i\} \cup pa_i$ d-separates F_i from the rest of the network (thus S is independent of F_i, conditioning on (X_i, pa_i)).

- A necessary and sufficient condition for a external intervention $set(X_i = x_i')$ to have the same effect on X_j as the passive observation $X_i = x_i'$ is that X_i d-separates pa_i from X_j.

Estimation of effects of interventions from passive (nonexperimental) observations?

Not always possible, as this would require estimation of $pr(x_i|pa_i)$. The mere identification of pa_i (the direct causal factors of X_i) require substantive knowledge of the network, which is often unavailable and some members of pa_i may be unobservable or latent.

- The aim of the paper is to derive causal effects in situations such as Fig. 1, where some members of pa_i, may be unobservable.
- Assume we are given a causal diagram G together with nonexperimental data on a subset V_0 of observed variables in G, and we wish to estimate what effect the intervention $set(X_i = x_i)$ would have on some response variable X_j. In other words, we seek to estimate $p(x_j|\check{x}_i)$ from a sample estimate of $pr(V_0)$.

Concomitants

- The variables in $V_0 - \{X_i, X_j\}$, are commonly known as concomitants or cofounders or covariates (Cox, 1958, p. 48).
- In observational studies, concomitants are used to reduce confounding bias due to spurious correlations between treatment and response.

Simpson's paradox (Karl Pearson et al., 1899)

- Any statistical relationship between two variables may be reversed by including additional factors in the analysis.
- For example, we may find that students that smoke obtain higher grades than those who do not smoke, but, adjusting for age, smokers obtain lower grades in every age group and, further, adjusting for family income, smokers again obtain higher grades than non smokers in every income-age group.

Graphical test for concomitants

- The condition that renders a set Z of concomitants sufficient for identifying causal effect, also known as **ignorability**, has been given a variety of formulations, all requiring conditional independence judgments involving counterfactual variables (Rosenbaum & Rubin, 1983; Pratt & Schlaifer, 1988).
- Pearl (1993) shows that such judgments are equivalent to a simple graphical test, named the 'back-door criterion', which can be applied directly to the causal diagram.

Definition: the backdoor criterion

A set of variables Z satisfies the back-door criterion relative to an ordered pair of variables (X_i, X_j) in a DAG G if:

The criterion

(a) No node in Z is a descendant of X_i, **and**
(b) Z blocks every path between X_i, and X_j which contains an arrow into X_i. If X and Y are two disjoint sets of nodes in G, Z is said to satisfy the back-door criterion relative to (X, Y) if it satisfies it relative to any pair (X_i, X_j) such that $X_i \in X$ and $X_j \in Y$.

The name 'back-door' requires that only paths with arrows pointing at X_i be blocked; these paths can be viewed as entering X_i through the back door.

Definition: the back-door criterion

In Fig. 2, for example, the sets $Z_1 = \{X_3, X_4\}$ and $Z_2 = \{X_4, X_5\}$ meet the backdoor criterion, but $Z_3 = \{X_4\}$ does not, because X_4 does not block the path $(X_i, X_3, X_1, X_4, X_2, X_5, X_j)$.

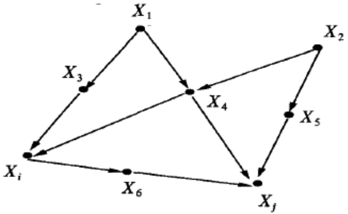

Fig. 2. A diagram representing the back-door criterion; adjusting for variables $\{X_3, X_4\}$ or $\{X_4, X_5\}$ yields a consistent estimate of $\text{pr}(x_j | \check{x}_i)$.

Identifiable causal effect

Theorem 1

If a set of variables Z satisfies the back-door criterion relative to (X, Y), then the causal effect of X on Y is identifiable and is given by the formula:

$$pr(y|\check{x}) = \sum_{z} pr(y|x, z)pr(z) \tag{5}$$

Identifiability means that $pr(y|\check{x})$ can be computed uniquely from any positive distribution of the observed variables that is compatible with DAG G.

A simple example of the back-door criterion

- A back-door path is a non-causal path. From the figure we can see that $A \leftarrow X \rightarrow Y$ is a path from A to Y that points into A.
- Back-door paths between A and Y generally indicate common causes (though not always). The simplest possible back-door path is the common confounding situation (see figure)
- When there are unblocked back-door paths, there are two sources of any association between A and Y: one causal (the effect of A on Y) and one non-causal (from the back-door path).Thus, with unblocked (meaning non d-separated) back-door paths, it's difficult to know if any association is a result of the causal effect or the back-door path.

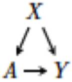

A simple example of the back-door criterion

How to tell if an effect is identifiable from the graph?

The back-door criterion states that an effect of A on Y is identifiable if either:
(a) No back door paths from A on Y (plausible only in a randomized experiment).
(b) Measured covariates are sufficient to block all back door paths from A on Y. (plausible in randomized and also in observational studies).

- The criterion tells is there confounding given a DAG, and
- if it is possible to removing the confounding, and
- what variables to condition on to eliminate the confounding.

The front door criterion

An alternative criterion, 'the front-door criterion', may be applied in cases where we cannot find observed covariates Z satisfying the back-door conditions.

Consider the diagram in Fig. 3. Although Z does not satisfy any of the back-door conditions, measurements of Z nevertheless enable consistent estimation of $pr(y|\check{x})$.

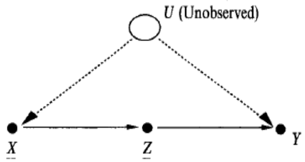

Fig. 3. A diagram representing the front-door criterion.

The front door criterion

The joint distribution associated with Fig. 3 can be decomposed into

$$p(x, y, z, u) = pr(u)pr(x|u)pr(z|x)pr(y|z, u), \tag{6}$$

and from (4), one can compute post intervention distributions (removing $pr(x|u)$)

$$p(y, z, u|\check{x}) = pr(y|z, u)pr(z|x)pr(u). \tag{7}$$

Summing over z and u then gives

$$p(y|\check{x}) = \sum_z pr(z|x) \sum_u pr(y|z, u)pr(u). \tag{8}$$

The front door criterion

Using the conditional independence assumptions implied by the graph: $pr(u|x,z) = pr(u|x)$ and $pr(y|x,z,u) = pr(y|z,u)$ yields

$$\begin{aligned}\sum_u pr(y|z,u)pr(u) &= \sum_x \sum_u pr(y|z,u)pr(u|x)p(x) \\ &= \sum_x \sum_u pr(y|x,z,u)pr(u|x,z)pr(x) \\ &= \sum_x pr(y|x,z)pr(x).\end{aligned} \tag{9}$$

Then, the causal effect of X on Y is given by

$$p(y|\check{x}) = \sum_z pr(z|x) \sum_{x'} pr(y|x',z)pr(x') \tag{10}$$

The front door criterion

Theorem 2

Suppose a set of variables Z satisfies the/allowing conditions relative to an ordered pair of variables (X, Y):
(a) Z intercepts all directed paths from X to Y,
(b) there is no back-door path between X and Z, and
(c) every back-door path between Z and Y is blocked by X. Then the causal effect of X on Y is identifiable and is given by (10).

Example: X = smoking, Y = lung cancer, Z = amount of tar deposited in a subject's lungs, and U = unobserved carcinogenic genotype that, also induces an inborn craving for nicotine. The theorem gives the means to quantify, from nonexperimental data, the causal effect of smoking on cancer, assuming that pr(x, y, z) is available and that smoking does not have any direct effect on lung cancer except that mediated by tar deposits.

The front door criterion as a two step application of the back-door criterion

First step: find the causal effect of X on Z, since there is no unblocked back-door path from X to Z in figure; so the effect is

$$pr(z|\check{x}) = pr(z|x).$$

Second step: compute causal effect from Z to Y, which we can no longer equate with the conditional probability $pr(y|z)$ because there is a back-door path $Z \leftarrow X \leftarrow U \rightarrow Y$ from Z to Y. However, since X d-separates this path, X can play the role of the concomitant, which allows us to compute the causal effect of Z to Y as

$$pr(y|\check{z}) = \sum_{x'} pr(y|x', z)pr(x').$$

Finally, we combine the two causal effects and get (10).

Preliminary notation

Let X, Y and Z be arbitrary disjoint sets of nodes in a DAG G. We denote by $G_{\overline{X}}$ the graph obtained by deleting from G all arrows pointing to nodes in X. Likewise, we denote by $G_{\underline{X}}$ the graph obtained by deleting from G all arrows emerging from nodes in X. To represent the deletion of both incoming and outgoing arrows, we denote $G_{\overline{X}\underline{Z}}$.

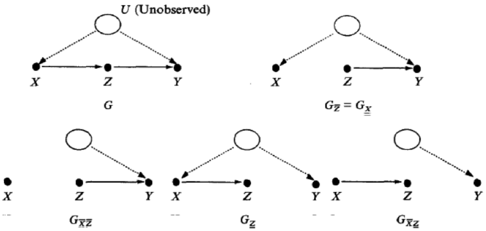

Fig. 4. Subgraphs of G used in the derivation of causal effects.

Inference rules

Theorem 3

Rule 1: (insertion/deletion of observations)

$$pr(y|\check{x}, z, w) = pr(y|\check{x}, w), if (Y \perp\!\!\!\perp Z|X, W)_{G_{\bar{X}}}$$

Rule 2:(action/observation exchange)

$$pr(y|\check{x}, \check{z}, w) = pr(y|\check{x}, z, w), if (Y \perp\!\!\!\perp Z|X, W)_{G_{\bar{X}\underline{Z}}}$$

Rule 3:(insertion/deletion of actions)

$$pr(y|\check{x}, \check{z}, w) = pr(y|\check{x}, w), if (Y \perp\!\!\!\perp Z|X, W)_{G_{\bar{X}\overline{Z(W)}}}$$

where $\overline{Z(W)}$ is the set of Z-nodes that are not ancestors of any W-node in $G_{\bar{X}}$.

Inference rules

- Rule 1 reaffirms d-separation as a valid test for conditional independence in the distribution resulting from the intervention $set(X = x)$, hence the graph $G_{\bar{X}}$. This rule follows from the fact that deleting equations from the system does not introduce any dependencies among the remaining disturbance terms.
- Rule 2 provides a condition for an external intervention $set(Z = z)$ to have the same effect on Y as the passive observation $Z = z$. The condition amounts to $X \cup W$ blocking all back-door paths from Z to Y in $G_{\bar{X}}$, since $G_{\bar{X}\underline{Z}}$ retains all, and only, such paths.
- Rule 3 provides conditions for introducing or deleting an external intervention $set(Z = z)$ without affecting the probability of $Y = y$. The validity of this rule stems, again, from simulating the intervention $set(Z = z)$ by the deletion of all equations corresponding to the variables in Z.

Inference rules

Corollary

A causal effect $q = pr(y_1, ..., y_k | \check{x}_1, ..., \check{x}_m)$ is identifiable in a model characterised by a graph G if there exists a finite sequence of transformations, each conforming to one of the inference rules in Theorem 3, which reduces q into a standard, i.e. check-free, probability expression involving observed quantities.

- Whether the three rules above are sufficient for deriving all identifiable causal effects remains an open question.

Graphs, structural equations and counterfactuals

- The primitive object of analysis in the potential-outcome framework is the unit-based response variable, denoted $Y_x(u)$, read: "the value that outcome Y would obtain in experimental unit u, had treatment X been x". This variable has a natural interpretation in structural equations model: consider a set T of equations

$$X_i = f_i(PA_i, U_i), (i = 1, ..., n),$$

which is similar to (3), except we no longer insist on the equations being recursive or on the U_i's being independent.

Graphs, structural equations and counterfactuals

- Let U stand for the vectors $U_1, \dots, U_n$, let X and Y be two disjoint subsets of observed variables, and let T_x be the submodel created by replacing the equations corresponding to variables in X with $X = x$. The structural interpretation of $Y_x(u)$ is given by

$$Y_x = Y_{T_x}(u).$$

- Namely, $Y_x(u)$ is the unique solution for Y under the realisation $U = u$ in the submodel T_x of T.
- While the term unit in the counterfactual literature normally stands for the identity of a specific individual in a population, a unit may also be thought of as the set of attributes that characterise that individual, the experimental conditions, which are represented as components of the vector u in structural modelling.

Graphs, structural equations and counterfactuals

- If U is treated as a random variable, then the value of the counterfactual $Y_x(u)$ becomes a random variable as well, denoted by $Y(x)$ or Y_x. The counterfactual analysis proceeds by imagining the observed distribution $pr(x_1, ..., x_n)$ as the marginal distribution of an augmented probability function pr* defined over both observed and counterfactual variables.
- Queries about causal effects, written $p(y|\check{x})$ in the structural analysis, are phrased as queries about the marginal distribution of the counterfactual variable of interest, written $pr^*(Y(x) = y)$.
- The new entities $Y(x)$ are treated as ordinary random variables that are connected to the observed variables via the logical constraints

$$X = x \rightarrow Y(x) = Y$$

Graphs, structural equations and counterfactuals

Main conceptual difference between the two approaches

Whereas the structural approach views the intervention $set(X = x)$ as an operation that changes a distribution but keeps the variables the same, the potential-outcome approach views the variable Y under intervention to be a different variable, Y_x (inferring probabilistic properties of Y_x, then becomes one of "missing data") .

Graphs, structural equations and counterfactuals

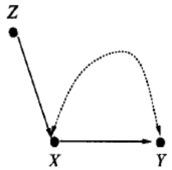

- For example, to communicate the understanding that in a randomised clinical trial, the way subjects react, Y, to treatments X is statistically independent of the treatment assignment Z, the analyst would write $Y(x) \perp\!\!\!\perp Z$.
- Likewise, to convey the understanding that the assignment process is randomised, hence independent of any variation in the treatment selection process, structurally written as $U_X \perp\!\!\!\perp U_Z$, the analyst would use the independence constraint $X(z) \perp\!\!\!\perp Z$.

Graphs, structural equations and counterfactuals

- To further formulate the understanding that Z does not affect Y directly, except through X, the analyst would write a, so called, "exclusion restriction": $Y_{xz} = Y_x$.
- A collection of constraints of this type might sometimes be sufficient to permit a unique solution to the query of interest. For example, if one can plausibly assume that, a set Z of covariates satisfies the conditional independence $Y_x \perp\!\!\!\perp X | Z$, then the causal effect $P(y|set(X = x)) = P^*(Y_x = y)$ can be evaluated to yield:

$$P^*(Y_x = y) = \sum_z P^*(Y_x = y|z)P(z) = \sum_z P^*(Y_x = y|x, z)P(z)$$
$$= \sum_z P^*(Y = y|x, z)P(z) = \sum_z P(y|x, z)P(z) \tag{11}$$

which is the same expression as (5): it mirrors therefore the "back-door" criterion.

References

- Cox, D. R. (1958). The Planning of Experiments. New York: John Wiley.
- Pearson, Karl; Lee, Alice; Bramley-Moore, Lesley (1899). Genetic (reproductive) selection: Inheritance of fertility in man, and of fecundity in thoroughbred racehorses. Philosophical Transactions of the Royal Society A 192: 257–330.
- Rosenbaum, P. & Rubin, D. (1983). The central role of propensity score in observational studies for causal effects. Biometrika 70, 41-55.
- Pratt, J. W. & Schlaieer, R. (1988). On the interpretation and observation of laws. J. Economet. 39, 23-52.
- Pearl, J. (1993): "Comment: Graphical models, causality, and intervention," Statistical Science, 8, 266–269.
- Pearl, J. (1995): "Causal diagrams for empirical research," Biometrika, 82, 669–710.