

Causal diagrams for encoding and evaluation of information bias

Eyal Shahr MD MPH

Professor, Division of Epidemiology and Biostatistics, Mel and Enid Zuckerman College of Public Health, The University of Arizona, Tucson, AZ, USA

Keywords

causal diagram, information bias, surrogate variables

Correspondence

Eyal Shahr
Division of Epidemiology and Biostatistics
Mel and Enid Zuckerman College of Public Health
The University of Arizona
1295 N. Martin Ave.
Tucson
AZ 85724
USA
E-mail: shahr@email.arizona.edu

Accepted for publication: 8 April 2008

doi:10.1111/j.1365-2753.2008.01031.x

Abstract

Background Epidemiologists and clinical researchers usually classify bias into three main categories: confounding, selection bias and information bias. Previous authors have described the first two categories in the logic and notation of causal diagrams, formally known as directed acyclic graphs (DAG).

Methods I examine common types of information bias – disease-related and exposure-related – from the perspective of causal diagrams.

Results Disease or exposure information bias always involves the use of an effect of the variable of interest – specifically, an effect of true disease status or an effect of true exposure status. The bias typically arises from a causal or an associational path of no interest to the researchers. In certain situations, it may be possible to prevent or remove some of the bias.

Conclusions Common types of information bias, just like confounding and selection bias, have a clear and helpful representation within the framework of causal diagrams.

Introduction

Epidemiologists and clinical researchers usually classify various kinds of bias into three main categories: confounding bias, selection bias and information (measurement) bias. Previous authors have described the first two categories in the logic and notation of causal diagrams, formally called ‘directed acyclic graphs’ (DAG) [1,2]. Briefly, confounding bias arises from a common cause of the exposure (putative cause) and the disease, whereas selection bias arises from conditioning on a common effect (a collider) by unnecessary restriction, stratification, or ‘statistical adjustment’. The DAG perspective of information bias was presented at a conference [3], and some aspects were mentioned in a few peer-reviewed publications [4–7]. To my knowledge, no article was devoted to this topic.

Surrogate for true disease status

The variables of interest to us are often beyond our reach. Even when we know precisely what we wish to measure, our measurement usually contains some error and we end up having to rely on a surrogate variable for the ‘real one’. In particular, measured (classified) disease status often differs from true disease status

because the measured value is influenced by variables other than true disease status. For example, in a study of risk factors for hospitalized stroke, factors such as consent to obtain a medical record, record quality, and the training level of the reviewer will affect classified stroke status. These ‘other variables’, as well as true stroke status, affect the chances of classifying an event as a stroke.

Figure 1 shows the general causal structure for this situation: E is the exposure (risk factor) of interest, D is true disease status, and D* is its surrogate – the measured version. The arrow from D to D* signifies that true disease status is, undoubtedly, an important cause of classified disease status, although not the only cause. The variables $M_1, M_2, \dots, M_n$ represent all other causes of D*, many of which are unknown or unmeasured. For pedagogical reasons, we assume that Fig. 1 fully captures causal reality. That is, no relevant causal connection is missing and E is measured with no meaningful error.

Disease information bias

Both the rationale behind using a surrogate variable, such as D*, and its drawbacks may be derived from the principles of causal diagrams. Although the effect of interest is $E \rightarrow D$, we can estimate




Figure 1 A directed acyclic graph showing the causes of classified disease status (D^*): E, exposure status; D, true disease status; M_1 , M_2 , ..., M_n , other causes of classified disease status (besides D).



Figure 2 A directed acyclic graph showing two causal paths from exposure status (E) to classified disease status (D^*): (1) through true disease status (D); (2) through a cause of classified disease status (M).

only the effect of E on D^* – even in the context of a randomized trial. As D^* is rarely correlated perfectly with D, if ever, the marginal association between E and D^* contains a component of bias; it is not estimating the effect of E on D (assuming the effect is not precisely zero.) The stronger the effects of those M variables (Fig. 1), the greater is the gap between the measured association (E with D^*) and the association of interest (E with D). Precisely for this reason researchers try to minimize the effect of measurement variables by means such as a uniform method of disease ascertainment and classification. Notice that the use of multiple classification methods within a single study might not always be desirable from that perspective. For instance, if incident stroke status is based either on neurological symptoms or on brain imaging, the chances of classifying a suspected event as a stroke depend on the method.

Our unavoidable need to use D^* , the measured outcome, instead of D, the true outcome, can lead to another classical type of disease information bias. If E affects D^* by pathways that do not involve D (Fig. 2), the marginal association of E with D^* reflects not only the path $E \rightarrow D \rightarrow D^*$, but also the path $E \rightarrow M \rightarrow D^*$, which is usually of no causal interest.

Example: replacement hormones and endometrial cancer

Figure 3 shows a possible set of causal connections that could create disease information bias when studying the effect of replacement hormones on endometrial cancer.

We are interested in the effect of replacement hormones (E) on endometrial cancer (D), but the statistical association between the two variables cannot be estimated directly because D is not available. We can only estimate the effect of replacement hormones on classified (diagnosed) endometrial cancer status (D^*). That variable, however, may be affected by other causes, which themselves may be affected by replacement hormones. For instance, it is possible to rationalize the existence of a causal path from the use




Figure 3 A directed acyclic graph showing two causal paths from the use of replacement hormones (E) to classified endometrial cancer status (D^*): (1) through true endometrial cancer status (D); (2) through the frequency of gynecological examinations (M).



Figure 4 A directed acyclic graph showing two paths from exposure status (E) to classified disease status (D^*): (1) a causal path through true disease status (D); (2) an associational path through a common cause (C) of exposure status (E) and a cause (M) of classified disease status.

of replacement hormones (E) to the diagnosis of endometrial cancer (D^*) via the frequency of gynecological examinations (M): gynecologists likely examine their patients more frequently after prescribing replacement hormones, and frequent gynecological examinations increase the chances of detecting (diagnosing) endometrial cancer. As a result, the association of replacement hormones with classified endometrial cancer contains a causal path of no interest to us – a classical type of information bias.

Notice that disease information bias could arise not only if E causes M, but also if E and M are associated via an unmeasured common cause (Fig. 4). For example, the age of the physician (perhaps a surrogate for training and experience) may affect hormone prescribing behavior as well as the frequency of gynecological examinations. On the assumption of that diagram, the marginal association of E and D^* still contains an associational path of no interest: $E \leftarrow C \rightarrow M \rightarrow D^*$. Interestingly, that structure may also be called ‘confounding’, because C is a common cause of E and D^* .

Depending on the nature of M, information bias goes by various names such as ‘unmasking bias’, ‘diagnosis suspicion bias’, and ‘ascertainment bias’. In most cases, however, the mechanism involves at least one path from E to D^* that does not pass through D.

Analytical approach to disease information bias

Researchers often ‘adjust’ for an intermediary variable on a causal pathway to block the pathway and thereby separate a direct effect from indirect effects. Although that practice has been debated, such conditioning is permissible under certain conditions [8–10]. On the assumptions encoded in Figs 2 and 3, conditioning on M will block the information path between E and D^* via M, as long as M is not an effect modifier of the effect of E on D^* [10]. For example, the bias may be reduced by adding the variable ‘frequency of gynecological examinations’ to a regression model of




Figure 5 A directed acyclic graph showing three causal paths from exposure status (E) to classified disease status (D*): (1) through true disease status (D); (2) through a cause of classified disease status (M); (3) through both D and M. Variable M is a collider on the path $E \rightarrow M \leftarrow D \rightarrow D^*$.




Figure 6 A directed acyclic graph showing three causal paths from aspirin use (E) to classified stomach cancer status (D*): (1) through true cancer status (D); (2) through melena status (M); (3) through both true cancer status and melena status. Melena status is a collider on the path Aspirin use → Melena status ← True stomach cancer status → Classified cancer status.

endometrial cancer (the dependent variable) and replacement hormones (the independent variable).

Nonetheless, we face another level of complexity if D happens to be a cause of M as well (Fig. 5). In that case, M becomes a collider on the path $E \rightarrow M \leftarrow D \rightarrow D^*$ and conditioning on M will induce an association between E and D, thereby opening a non-causal, associational path between E and D* ($E \rightarrow D \rightarrow D^*$). Figure 6 shows a hypothetical example. Melena may play the role of M in a study of the effect of aspirin (E) on stomach cancer (D), because taking aspirin and stomach cancer are both causes of melena. The presence of melena would trigger diagnostic studies, which in turn would increase the chances of detecting stomach cancer (D*).

The last situation is close to hopeless: we may be able to estimate the magnitude of the bias under various assumptions and simulations, but we cannot take any action to eliminate it. If we do not condition on M, the marginal association of E and D* contains a path of information bias ($E \rightarrow M \rightarrow D^*$). And if we do condition on M, the conditional association of E and D* contains selection bias because M is a collider on the path $E \rightarrow M \leftarrow D \rightarrow D^*$. Furthermore, to create an unsolvable situation, D need not be a cause of M: all that is needed is an unmeasured common cause of D and M, such as a common gene for both stomach cancer and stomach bleeding. Previous work has encoded the causal structure of that bias in studies of postmenopausal estrogen and endometrial cancer where M was vaginal bleeding [4].

Information bias depends on the causal question

Although we are usually interested in true disease status (D) as the outcome variable, it is possible to think of a case where classified, or detected disease status (D*) itself is also of causal interest. For




Figure 7 A directed acyclic graph showing the causes of classified exposure status (E*): E, true exposure status; D, disease status; M_1 , M_2 , ..., M_n , other causes of classified exposure status (besides E).




Figure 8 A directed acyclic graph showing two causes of reported family history of heart disease (E*): E, true family history status; M, subject memory.

example, we may be interested in the effect of a widely prescribed drug not only on the development of dementia, but also on the *detection* of dementia, perhaps because detected disease affects psychological well-being, medical care and health insurance. In that case, the association of E and D* provides all that we need, assuming it reflects all causal pathways from E to D* (and there are no other sources of bias).

Exposure information bias

Analogous to measured disease status (D*) which substitutes for true disease status (D), measured (classified) exposure status (E*) often substitutes for true exposure status (E). Assuming no meaningful measurement error of D, the causal structure is depicted in Fig. 7.

Again, the association of interest (E with D) cannot be estimated directly. We can only estimate the association of E* with D, which would differ from the association of interest (assuming E has some effect on D). Notice, however, that E is a common cause of E* and D, and that E* itself is not a cause of D. Therefore, the observed association between E* and D is generated by the path $E^* \leftarrow E \rightarrow D$, embedding two causal segments: $E \rightarrow D$, which is of interest, and $E \rightarrow E^*$, which is not. Interestingly, the idea of a common cause – the classical property of a confounder – serves us well here because the common cause happened to be the exposure itself.

As before, the surrogate E* should be strongly correlated with the true exposure, E. If E* is strongly affected by the M variables (Fig. 7), then its correlation with E would be weak, and its association with D will be a poor estimate of the association of interest. Figure 8 shows a simple example from a hypothetical study of the



Figure 9 A directed acyclic graph showing recall bias in a case–control study: E, true exposure status; E*, classified exposure status; D, true disease status.

relation between family history of heart disease and heart disease. Evidently, if the sample includes elderly people with memory problems, the association of heart disease with reported family history of heart disease would be very different from the association of interest.

A classical kind of exposure information bias occurs if $D \rightarrow E^*$ – if disease status supplies information on the measured exposure (Fig. 9). What is usually labelled ‘recall bias’ in case–control studies follows that causal structure: for various psychological reasons, cases may report their historical exposure differently from controls. Consequently, the marginal association between E^* and D would reflect not only the path $E^* \leftarrow E \rightarrow D$, but also the path $D \rightarrow E^*$.

Of course, the second path is of no interest to us; it produces information bias. No analytical remedy is available, however, and only preventive means are possible. For example, epidemiologists sometimes choose sick people as controls on the assumption that sick people – whether cases or controls – share the same state of mind and would therefore recall exposures similarly. In the coding of Fig. 9, they try to eliminate the arrow from D to E^* through a clever choice of the sample. As always, it is also possible to estimate the magnitude of the bias under hypothetical scenarios and various simulations.

Previous descriptions in light of causal diagrams

The types of information (or measurement) bias I discussed here were classically described as ‘differential measurement error’ or ‘differential misclassification’. For example:

Measurement error in one variable is *differential* with respect to a second variable if the magnitude of error in the first variable differs according to the true value of the second variable. [11, p. 348]

Differential misclassification occurs when the degree of misclassification differs between the groups being compared. [12, p. 145]

In differential misclassification, the rate of misclassification differs in different study groups. [13, p. 226]

Following organization of these descriptions in three columns, they all share the same format:

Column 1	Column 2	Column 3
‘error in the first variable	differs according to	the true value of the second’
‘degree of misclassification	differs between	the groups being compared’
‘rate of misclassification	differs in	different study groups’

If we code the first column as variable B^* and the third column as variable A, we may explain the mechanism behind the word ‘differs’ in the second column as any open (i.e. associational) path between the true variable (A) and the misclassified variable (B^*). Such an open path may be created by the structures $A \rightarrow B^*$ or $A \leftarrow C \rightarrow B^*$, or by conditioning on C in the following structure: $A \rightarrow C \leftarrow B^*$.

Limitations

I presented the DAG perspective of disease information bias assuming that exposure was measured with no meaningful error. Likewise, I presented the DAG perspective of exposure information bias assuming an error-free measurement of disease status. In practice, both exposure and disease are often measured with some error and therefore other mechanisms of information bias may be operating, including non-differential misclassification [3]. As the DAG-perspective is qualitative, it offers no inference on the magnitude and direction of the bias, which may be inferred by subject matter considerations, or quantitatively studied under various assumptions. Finally, DAG-based inference is only as good as the encoded assumptions, many of which reflect background causal knowledge [4]. If key assumptions are false, the inference is not helpful. For example, it is unlikely that any of the diagrams I presented here – for pedagogical reasons – indeed depicts all relevant causal connections.

Conclusion

The principles of DAG clearly explain how key categories of bias (confounding, selection and information) interfere with our attempts to estimate effects by marginal or conditional associations. These principles teach us that three kinds of paths may contribute to a measure of association between the exposure and the disease – besides the causal paths of interest: (1) naturally occurring confounding paths (by common causes of E and D); (2) induced, non-causal, associational paths because of conditioning on colliders (some of which may be called ‘sampling colliders’); and (3) causal or associational paths of no interest to us because of our using surrogates for E or D (information bias).

Common types of information bias always involve the use of an effect of true disease status or an effect of true exposure status. In certain situations, it may be possible to prevent or remove some of the bias.

Acknowledgements

I thank Miguel Hernán for sharing his notes, thoughts, and 2005 Society for Epidemiological Research presentation on this topic. I also thank several reviewers whose comments led to a minor revision of the manuscript.

References

- Greenland, S., Pearl, J. & Robins, J. M. (1999) Causal diagrams for epidemiologic research. *Epidemiology*, 10, 37–48.
- Hernán, M. A., Hernandez-Díaz, S. & Robins, J. M. (2004) A structural approach to selection bias. *Epidemiology*, 15, 615–625.

3. Hernán, M. A. (2005) A structural approach to observation bias. *American Journal of Epidemiology*, 161, S100 (abstract).
4. Robins, J. M. (2001) Data, design, and background knowledge in etiologic inference. *Epidemiology*, 12, 313–320.
5. Glymour, M. M., Weuve, J., Berkman, L. F., Kawachi, I. & Robins, J. M. (2005) When is baseline adjustment useful in analyses of change? An example with education and cognitive change. *American Journal of Epidemiology*, 162, 267–278.
6. Greenland, S. (2000) An introduction to instrumental variables for epidemiologists. *International Journal of Epidemiology*, 29, 722–729.
7. Greenland, S. (2003) Quantifying biases in causal models: classical confounding vs collider-stratification bias. *Epidemiology*, 14, 300–306.
8. Cole, S. R. & Hernán, M. A. (2002) Fallibility in estimating direct effects. *International Journal of Epidemiology*, 31, 163–165.
9. Kaufman, J. S., Maclehose, R. F. & Kaufman, S. (2004) A further critique of the analytic strategy of adjusting for covariates to identify biologic mediation. *Epidemiological Perspectives and Innovations*, 1, 4.
10. Petersen, M. L., Sinisi, S. E. & van der Laan, M. J. (2006) Estimation of direct causal effects. *Epidemiology*, 17, 276–284.
11. Kelsey, J., Whittemore, A., Evans, A. & Thompson, W. (1996) *Methods in Observational Epidemiology*. New York: Oxford University Press.
12. Szklo, M. & Nieto, F. J. (2000) *Epidemiology: Beyond the Basics*. Gaithersburg, MD: Aspen Publishers.
13. Gordis, L. (2004) *Epidemiology*. Philadelphia, PA: Elsevier Saunders.