Association vs Causation

Arvid Sjölander

Department of Medical Epidemiology and Biostatistics
Karolinska Institutet

A short course on concepts and methods in Causal Inference

Outline

Outline

Preliminaries

- Suppose we are interested in the relation between an exposure, A, and an outcome, Y
- We assume for simplicity that both A and Y are binary
  - We use ‘0’ for ‘unexposed/no outcome’, and ‘1’ for ‘exposed/outcome’
- We assume that population data are available (infinite sample size)
  - No need for p-values, confidence intervals etc
- These conditions are often unrealistic, but are useful for pedagogical purposes
  - Will be relaxed later

Joint probability

- Suppose that the population proportions of A and Y are given by

		Y = 0	Y = 1
A	0	0.88	0.02
	1	0.09	0.01

- Among all subjects, 1% are both exposed and have the outcome
- We say that the **joint probability** of $(A = 1, Y = 1)$ is 0.01
- We denote this as $\Pr(A = 1, Y = 1) = 0.01$

Association Causation Subtle points

Marginal probability

		Y = 0	Y = 1
A	0	0.88	0.02
	1	0.09	0.01
	$\sum$	0.97	0.03

- Among all subjects, 3% have the outcome
- We say that the **marginal probability** of $Y = 1$ is 0.03
- We denote this as $\Pr(Y = 1) = 0.03$

Association Causation Subtle points

Conditional probability

		Y = 0	Y = 1
A	0	0.88	0.02
	1	0.09	0.01

- Among the exposed subjects, $\frac{0.01}{0.01+0.09} = 10\%$ have the outcome
- We say that the **conditional probability** of having the outcome, for exposed subjects, is 0.1
- We denote this as $\Pr(Y = 1|A = 1) = 0.1$

Association Causation Subtle points

Definition of association and independence

- We say that A and Y are **independent** if the risk of the outcome is the same for exposed and unexposed:

$$\Pr(Y = 1|A = 1) = \Pr(Y = 1|A = 0) = \Pr(Y = 1)$$

  - We sometimes write this as $Y \perp\!\!\!\perp A$.
- We say that A and Y are **associated** if the risk of the outcome is different for exposed and unexposed:

$$\Pr(Y = 1|A = 1) \neq \Pr(Y = 1|A = 0) \neq \Pr(Y = 1)$$

  - We sometimes write this as $Y \not\perp\!\!\!\perp A$

Example

		Y	
		0	1
A	0	0.88	0.02
	1	0.09	0.01

- *Are A and Y independent or associated in the table?*

Solution

Remark

- There may be several explanations to an association between A and Y
  - A causes Y
  - Y causes A ('reverse causation')
  - A and Y have common causes ('confounding')
- That A and Y are associated only means that certain values of A and Y tend to 'appear together'
  - Why this happens is a different question

Common association measures for binary variables

- The risk difference

$$RD = \Pr(Y = 1|A = 1) - \Pr(Y = 1|A = 0)$$

$Y \amalg A \Leftrightarrow RD = 0$

- The risk ratio

$$RR = \frac{\Pr(Y = 1|A = 1)}{\Pr(Y = 1|A = 0)}$$

$Y \amalg A \Leftrightarrow RR = 1$

- The odds ratio

$$OR = \frac{\Pr(Y = 1|A = 1)}{\Pr(Y = 0|A = 1)} \Big/ \frac{\Pr(Y = 1|A = 0)}{\Pr(Y = 0|A = 0)}$$

$Y \amalg A \Leftrightarrow OR = 1$

Example

		Y	
		0	1
A	0	0.88	0.02
	1	0.09	0.01

- *Compute RD , RR, and OR for the table*

Solution

Conditional association/independence

- Sometimes we wish to stratify data before analysis, e.g:
  - L = 'sex' (0=male, 1=female).
  - $\Pr(Y=1|A=a, L=1)$ is the conditional probability of having the outcome, for women with exposure level $A=a$
- **Definition**:
  - A and Y are conditionally independent, given L, if

$$\Pr(Y=1|A=1,L)=\Pr(Y=1|A=0,L)=\Pr(Y=1|L)$$

$$Y \amalg A \mid L$$

  - A and Y are conditionally associated, given L, if

$$\Pr(Y=1|A=1,L)\neq\Pr(Y=1|A=0,L)\neq\Pr(Y=1|L)$$

$$Y \not\amalg A \mid L$$

Technical note

- In principle, we could have that
  - $\Pr(Y=1|A=1,L)=\Pr(Y=1|A=0,L)$ for some values of L, and
  - $\Pr(Y=1|A=1,L)\neq\Pr(Y=1|A=0,L)$ for other values of L
- When we write $Y \amalg A \mid L$, we mean that $\Pr(Y=1|A=1,L)=\Pr(Y=1|A=0,L)$ for **all** values of L
- When we write $Y \not\amalg A \mid L$, we mean that $\Pr(Y=1|A=1,L)\neq\Pr(Y=1|A=0,L)$ for **at least one** value of L

Measures of conditional association

- Conditional risk difference

$$RD(L) = \Pr(Y = 1|A = 1, L) - \Pr(Y = 1|A = 0, L)$$

- Conditional risk ratio

$$RR(L) = \frac{\Pr(Y = 1|A = 1, L)}{\Pr(Y = 1|A = 0, L)}$$

- Conditional odds ratio

$$OR(L) = \frac{\Pr(Y = 1|A = 1, L)}{\Pr(Y = 0|A = 1, L)} \Big/ \frac{\Pr(Y = 1|A = 0, L)}{\Pr(Y = 0|A = 0, L)}$$

Outline

Association

Causation

Subtle points

Causal models

- The sufficient-component cause model (Rothman)
- Potential outcomes, counterfactuals (Rubin, Robins)
- Structural equations, causal diagrams (Pearl)

Relation between models

- All common causal models are essentially equivalent, from a mathematical perspective
  - different languages, same content
- To define ‘causation’, we will mostly rely on the potential outcome model, but borrow from the other models as well

Motivating example

- August has been smoking 5 cigs/day since he was 15 years old. At the age of 60 he develops liver cancer
- *Did the smoking cause the cancer?*

Human reasoning about cause and effects

- We mentally compare two scenarios:
  - the outcome when the exposure is present
  - the outcome when the exposure is absent

  everything else equal
- If the two outcomes differ, then we say that the exposure has a causal effect
  - causative or preventative

Ideal data

- Let Y_a be the outcome that we would observe, for a given subject, if the subject potentially received exposure level a
  - Y_1 is the outcome under exposure
  - Y_0 is the outcome under non-exposure
- Y_1 and Y_0 are referred to as **potential outcomes**
- Ideally - **and very unrealistically** - we could observe both potential outcomes for any given subject

subject	Y_1	Y_0
August	1	0
Selma	0	0
Fjodor	1	1

Subject-specific causal effects

subject	Y_1	Y_0
August	1	0
Selma	0	0
Fjodor	1	1

- A has a causal effect on Y, for a given subject, if the potential outcomes Y_1 and Y_0 differ for this subject
  - For August, the exposure has an effect: $Y_1 \neq Y_0$
  - For Selma and Fjodor, the exposure has not effect; $Y_1 = Y_0$

Observed data

- August is exposed ($A = 1$). Thus, for August
  - Y_1 is observed and equal to the factual outcome Y
  - Y_0 is unobserved, or **counterfactual**
- Selma and Fjodor are unexposed ($A = 0$). Thus, for Selma and Fjodor
  - Y_0 is observed and equal to the factual outcome Y
  - Y_1 is unobserved, or **counterfactual**

subject	A	Y	Y_1	Y_0
August	1	1	1	?
Selma	0	0	?	0
Fjodor	0	1	?	1

A fundamental problem of causation

- It is very difficult to say whether the exposure causes the outcome for a specific subject
  - because we cannot observe the same subject under two exposure levels simultaneously

Population effects

- Fortunately, it is much easier to make causal claims on population levels
  - e.g. ‘if everybody would quit smoking, then the incidence of liver cancer would decrease by 15%’
  - more later

Population causal effects

- $\mathrm{Pr}(Y_a = 1)$ is the proportion of subjects that would develop the outcome, if **everybody** would receive exposure level a
  - The probability of the outcome if everybody would receive a
- A has a population causal effect on Y if

$$\mathrm{Pr}(Y_1 = 1) \neq \mathrm{Pr}(Y_0 = 1)$$

- A has no population causal effect on Y if

$$\mathrm{Pr}(Y_1 = 1) = \mathrm{Pr}(Y_0 = 1)$$

Technical note

- In statistics, we use
  - upper case letters (e.g. A, Y) for random variables
  - lower case letters (e.g. a, y) for fixed numbers
- When writing Y_a, we consider the exposure to be fixed to a (0 or 1)
- When writing $\Pr(Y_a = 1)$, we consider a scenario where the exposure is fixed to a for everybody

Association vs Causation

- Association:

$A = 1$ $A = 0$

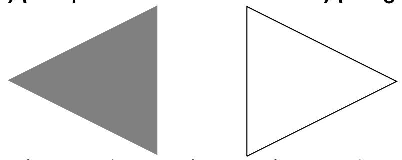

$$\Pr(Y = 1|A = 1) \text{ vs } \Pr(Y = 1|A = 0)$$

- Causation:

$a = 1$ $a = 0$

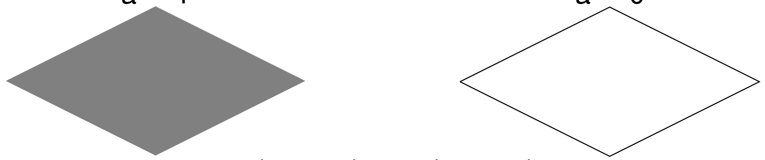

$$\Pr(Y_1 = 1) \text{ vs } \Pr(Y_0 = 1)$$

Measures of causal effects for binary variables

- The causal risk difference

$$CRD = \Pr(Y_1 = 1) - \Pr(Y_0 = 1)$$

Causal null hypothesis $\Leftrightarrow$ $CRD = 0$

- The causal risk ratio

$$CRR = \frac{\Pr(Y_1 = 1)}{\Pr(Y_0 = 1)}$$

Causal null hypothesis $\Leftrightarrow$ $CRR = 1$

- The causal odds ratio

$$COR = \frac{\Pr(Y_1 = 1)}{\Pr(Y_1 = 0)} \Big/ \frac{\Pr(Y_0 = 1)}{\Pr(Y_0 = 0)}$$

Causal null hypothesis $\Leftrightarrow$ $COR = 1$

Example

subject	Y_1	Y_0
1	0	0
2	1	0
3	0	0
4	1	1
5	0	0
6	1	1
7	1	1
8	1	1
9	0	0
10	1	0

- *Compute CRD, CRR, and COR*

Solution

Conditional causal effects

- Conditional causal risk difference

$$CRD(L) = \Pr(Y_1 = 1|L) - \Pr(Y_0 = 1|L)$$

- Conditional causal risk ratio

$$CRR(L) = \frac{\Pr(Y_1 = 1|L)}{\Pr(Y_0 = 1|L)}$$

- Conditional causal odds ratio

$$COR(L) = \frac{\Pr(Y_1 = 1|L)}{\Pr(Y_1 = 0|L)} \Big/ \frac{\Pr(Y_0 = 1|L)}{\Pr(Y_0 = 0|L)}$$

A brief remark

- We have seen that both association and causation can be quantified with risk differences, risk ratios, and odds ratios
- For convenience, we will mostly focus on risk ratios
- Everything that we say holds for risk differences and odds ratios as well

Outline

When is a counterfactual well defined?

- 'Well defined' = we have a clear understanding of what the counterfactual represents 'in real life'
- Are all counterfactuals well defined?
- If some counterfactuals are not well defined, then causal effects based on these are not well defined either

Example

- Define $A = 1$ if BMI>30, and $A = 0$ if BMI<30
- Certain diseases occur more frequently in obese than in non-obese, i.e.

$$\Pr(Y = 1|A = 1) > \Pr(Y = 1|A = 0)$$

- Does 'obesity' have a causal effect on the risk for disease?

$$\Pr(Y_1 = 1) \neq \Pr(Y_0 = 1)?$$

Quite a vague question

- Translated into plain English, the counterfactual comparison reads
  - 'what would the risk be if everybody had BMI>30 compared compared to if everybody had BMI<30?'
- But what does 'if everybody had BMI>30' really mean?
  - everybody short or heavy?
  - everybody fat or muscular?
  - everybody belly fat or hips fat?
- The outcome is probably very different under these alternative counterfactual scenarios
  - Unless we specify more precisely what scenario we refer to, the counterfactual outcome is not well defined

An important difference between association and causation

- In order for the causal effect of A on Y to be well defined we require that
  - we can tell whether an observed subject has $A = 1$ or $A = 0$
  - we agree on what it means that an observed subject with $A = 0$ **would have had** $A = 1$, and vice versa
- In order for the association between A and Y to be well defined, only the first condition is required
  - Because the concept of association is only based on factual observations, not on counterfactuals

'*Some counterfactuals are ill-defined, most are somewhat vague, but many are useful*'

Lewis, 1973

Summary

- Association is not equal to causation
- To define causation, we use **potential outcomes** and **counterfactuals**
- Not all counterfactuals (and causal effects) are well defined