Introduction to experimental design

Oriol Camps, Victor Peña

Experiments and observational studies

Experimenting involves intervening in the mechanism or system from which the data is being collected. When we simply observe what happens and collect data, we can discover associations (or correlations) between variables, but we cannot infer cause and effect relationships from these associations. The reason is that there are often confounding variables.

Example 1. Let's imagine that hospital A and hospital B receive a similar annual volume of patients with a serious illness. The death rate for this disease in hospital A is (for example) 5 times higher than in hospital B. Can we conclude that hospital A is worse than hospital B? Note that not only can we not conclude this, but it could very well be the other way around: it could be that the most severe cases were sent to hospital A, precisely because it has more resources or more specialized doctors. Here there is a confounding variable: the severity of the patients received in each hospital. If we ignore it, it can lead us to erroneous conclusions. If we assign a randomly chosen group of patients to each hospital, we may see that the association between hospital A and the number of deaths disappears or even reverses.

Example 2. A fire alarm will usually go off when there is fire (although there may be false positives and false negatives). There is a clear association between the presence of fire and the activation of the alarm, but it is the fire that causes the activation, not the other way around. If we experiment, we can more easily distinguish causal relationships from non-causal ones:

- If you light a fire, you will notice that the alarm goes off.
- If you activate the alarm electronically, you will naturally check that no fire appears. Correlation is bidirectional; causality is unidirectional.

It is important to see that much of applied statistics uses observational data, and that for example in econometrics and social sciences it is difficult to experiment. In this course, we will study different experimental designs and how they affect data analysis and statistical inference.

Some terminology

- Response (Y): The outcome we are interested in studying. In fact, the response may be multivariate, although we won't cover that in this course.
- Factors/Treatments (Xs): The variables we set during the experiment. We will use the terms "factor" and "treatment" interchangeably. In this half of the course, these are always categorical variables. The goal of an experiment is finding out how the factors (the Xs) affect the response (the Ys).
- Levels: The values that the factors can take on. For example, if X is a variable recording whether a patient takes a placebo or a new treatment, the levels are "placebo" and "treatment".
- Covariates: Additional variables we can observe during the experiment, but we do not set ourselves (that is, we cannot set the values of these variables to any values we want, in contrast to the factors). Sometimes there are no covariates and we only record the response Y and the factors X.

Example 3. A critical characteristic of steering wheels is that they must be hard enough not to break, but soft enough so that, in the event of an accident, the steering wheel breaks before the driver's ribs do. A manufacturer of automotive steering wheels has problems with the hardness of their products. The manufacturing process involves injecting polyurethane into a mold. The company decided to design an experiment where they set the injection pressure, the ratio of the two polyurethane components, and the injection temperature to different values and and see how they affect the hardness of the steering wheels. Here, the response Y is the hardness, and the factors are the pressure, the ratio, and the temperature. The levels are the different values that the pressure, ratio, and temperature can take on in the experiment. There are no additional covariates.

Blocking and randomizing

Randomization and blocking are tools for avoiding biases introduced by confounding variables, like the severity of the cases in Example 1.

The basic principles are

- Block: Group experimental units into known blocks to eliminate known causes of variability. The effects of blocks on the response are not of interest.
- Randomize: Assign the experimental units to treatments randomly to reduce potential biases.

In Example 1, we saw an example of randomization: the idea was assigning the patients to hospitals randomly to avoid the fact that, if we don't do that, hospital A might be getting more severe cases than hospital B.

For an example of blocking, let's reconsider Example 3. Suppose we are told that some of the experiments are going to be run by a worker in the day shift, whereas some other ones are going to be run by a worker in the night shift. To avoid "worker" or "time of day" effects, we can consider the worker who runs the experiment as a block. Then, we can decide which experiments are run in the day shift and which ones are run in the night shift randomly to avoid further biases.

George Box has a nice quote that nicely captures how we should think about randomizing and blocking when designing experiments:

> Block what you can; randomize what you cannot

Random and fixed effects

Loosely, there are two types of factors: random and fixed effects. There are no agreed upon definitions in the literature, but these are enough for our purposes:

- Random effects: The levels of the factors we observe are a random sample from a population of possible levels. We want to make inferences about all levels of the population, even though we only experiment with those we observe in the sample.
- Fixed effects: The levels we observe are all the levels we care about. We do not want to make inferences about other possible levels.

In the first half of the course, we will only consider fixed effects.