Statistics 263/363: Experimental Design

A. B. Owen

Autumn 2020/21

2:30–3:50 Tuesday & Thursday

`statweb.stanford.edu/~owen/courses/363`

and canvas

To gain understanding from data, we must work against: noise, bias, correlation, interaction and missingness. Often there is nothing we can do about some of those things, because we are just handed data with flaws embedded. Choosing or designing the data gives us much better possibilities. We all know that "correlation does not imply causation". Without a causal understanding, all we can do is predict outcomes, not confidently influence them. Injecting randomness into the independent variables provides the most convincing way to establish causality (though it is not perfect).

Experimental design is the science of choosing how to gather data. It has a long and continuing history spanning: agriculture, medicine and public health, education, engineering, computer experiments, A/B testing in e-commerce and more.

The design problem forces the statistical investigator to think carefully about the underlying domain topic, it's assumptions and costs, benefits, goals and prior history, even more than analyzing prior data does.

Pre-requisites

Completely deep and thorough understanding of basic probability (e.g. stat 116) and data analysis by regression. You should know those things well enough to explain them to other people, not just have a nodding aquantaince. You should also be able to program in R or python or Matlab or similar. Spreadsheets are not adequate.

Evaluation

Homework 50%, midterm 25%, project 25%. The project involves carrying out an experimental investigation and writing it up.

Course materials

They will all be online linked from the web page above or posted to canvas.

Level

This is a hybrid MS/PhD course. Lectures are at the PhD level. Evaluation is at the MS level. Some of the material touches on open problems and research areas.

Lecture topics

If we get one or two guest lectures then we will bump or move one of the more specialized lectures to work with the guest's schedule.

Sep 15	Introduction and Neyman-Rubin causal model
Sep 17	A/B testing in online settings
Sep 22	Bandits: upper confidence levels and Thompson sampling
Sep 24	Paired and blocked experiments. Randomization analysis
Sep 29	Factorial experiments: fixed/random/mixed effects. ANOVA tables.
Oct 01	2^k factorial experiments
Oct 06	Fractional factorials
Oct 08	Adjusting for pre-treatment variables and ANCOVA
Oct 13	Split plots and cluster randomized trials
Oct 15	Taguchi methods (robust design)
Oct 20	Catchup or summary or guest topic and review
Oct 22	Midterm on material up to class of Oct 15
Oct 27	Designs for response surfaces
Oct 29	Sequential clinical trials
Nov 03	Computer experiments 1
Nov 05	Computer experiments 2
Nov 10	Design for $p \gg n$ regression
Nov 12	Survey of other interesting designs (mixtures, crossovers, etc.)
Nov 17	Experimental / observational hybrids
Nov 19	Project presentations