UBC Seminar October 9, 2018

Adjusting for Bias Induced by Informative Dose Selection Procedures

Nancy Flournoy
University of Missouri, Columbia MO, USA

Assaf Oron
Institute for Disease Modeling, Bellevue WA, USA

Many fields such acute toxicity studies, Phase I cancer trials, sensory studies and psychometric testing use informative dose allocation procedures. In this talk, we explain how such adaptive designs induce bias, and in the context of dose-finding designs we show how to modify frequency data to adjust for this bias.

To provide context, we start the talk with a general discussion of issues in inference following adaptive designs. Then, we assume a binary response Y has a monotone positive response probability to a stimulus or treatment X, and we consider designs that sequentially select X values for new subjects in a way that concentrates treatments in a certain region of interest under the dose-response curve. We discuss how data analysis at the end of a study is affected by choosing the stimulus value for each subject sequentially according to some informative sampling rule.

Without loss of generality, we call a positive response a toxicity and the stimulus a dose. For simplicity, we restrict this talk to the case of a univariate treatment X and binary Y, and further assume that treatments are limited to a finite set $\{d_1, d_2, \ldots, d_M\}$ of M values we call doses. Now suppose n subjects receive treatments that were sequentially selected (according so some rule using data from prior subjects) from the restricted set of M doses. Let N_m and T_m denote the number of subjects receiving treatment d_m and the number of toxicities observed on treatment d_m, respectively. Define $F_m \equiv P\{Y = 1|X = d_m\} = \mathrm{E}[Y|X = d_m]$.

Then it is often said that the distribution of T_m given N_m is Binomial with parameters (F_m, N_m). But taking N_m as fixed is not the same as conditioning on this random variable, and conditioning on informative dose assignments is not the same as conditioning on summary dose frequencies. Indeed, it is easy to show that the observed dose-specific toxicity rate, T_m/N_m, is biased for F_m. From first principals, we obtain

$$\mathrm{E}\left[\frac{T_m}{N_m}\right] = F_m - \frac{\mathrm{Cov}[T_m/N_m, N_m]}{\mathrm{E}[N_m]}.$$

The observed toxicity rate is biased for F_m because adaptive allocations, by design, induce a correlation between toxicity rates and allocation frequencies.

This bias impacts inference procedures: Isotonic regression methods use dose-specific toxicity rates directly. Standard likelihood-based methods mask the bias by providing first-order linear approximations. We illustrate these biases using isotonic and likelihood-based regression methods in some well known (small sample size) adaptive methods including selected up-and-down designs, interval designs, and the continual reassessment method. Then we propose a bias adjustment inspired by Firth (1993).

[Nancy Flournoy; University of Missouri – http://web.missouri.edu/ flournoyn/]
[flournoyn@missouri.edu – https://en.wikipedia.org/wiki/Nancy_Flournoy]