TREATMENT PERSONALIZATION USING BAYESIAN DECISION THEORY.

DR. MAURITS C. KAPTEIN

Effective methods for the personalization of treatments are well sought for in the social sciences. For example, in the medical sciences physicians seek to select the most effective treatments for individual patients, and in on the Web firms seek content that is tailored to the recipient. Formally, such personalization attempts can be operationalized as a Bayesian decision problem: one first observes a number of features describing the recipient of the treatment. Next, given a probabilistic model of the expected outcomes given the features and the treatments (the posterior predictive distribution) and a loss function describing the feasibility of each expected outcome, we can select an action that minimizes the expected loss.

In this poster we explain this general approach, and show by numerical simulation that it is more effective than common (often frequentist) methods of personalized treatment selection. Furthermore, we highlight that loss functions can be specified both at the level of the receiving individual, but also at an aggregate level, representing (e.g.,) a policy maker. This latter extension highlights that personalizing with regard to aggregate goals does not always lead to the same actions as personalizing with regard to individual level goals. Furthermore, the aggregate perspective illustrates that an exploration-exploitation problem is introduced: at the aggregate a policy maker might prefer to experiment with novel treatments to improve model estimates and subsequently make better future decisions. This trade-off is—contrary to frequentist methods—quite naturally addressed within the Bayesian decision theoretic framework.