

Personalizing Interventions in Behavioral Health

Problem Definition: Lack of patient adherence to treatment protocols is a main barrier to reducing the global disease burden of tuberculosis (TB). Using data from a completed RCT, we study the operational design of a treatment adherence support (TAS) platform that requires patients to verify their treatment adherence on a daily basis, with a focus on improving *personalization*.

Methods: We first focus on *personalized enrollment*. We trained a causal forest model to answer three research questions: (1) Was the effect of the intervention heterogeneous across individuals? (2) Was the intervention less effective for high-risk patients? (3) Can differentiated care improve program effectiveness and equity in treatment outcomes? We then focus on *personalized outreach*. Inspired by reinforcement learning, we provide a model-free approach to solving the problem of optimizing personalized interventions for patients to maximize some long-term outcome, in a setting where interventions are costly and capacity-constrained.

Results: For *personalized enrollment*, we find that individual intervention effects—the percentage point reduction in the likelihood of an unsuccessful treatment outcome—ranged from 4.2 to 12.4, with an average of 8.2. The intervention was beneficial for 76% of patients, and most beneficial for high-risk patients. Differentiated enrolment policies, targeted at high-risk patients, have the potential to (1) increase the average intervention effect of DAT services by up to 28.5% and (2) decrease the population average and standard deviation (across patients) of the probability of an unsuccessful treatment outcome by up to 8.5% and 31.5%, respectively. For *personalized outreach*, we show that under a natural set of structural assumptions on patient dynamics, our approach recovers at least 1/2 of the improvement possible between a naive baseline policy and the optimal policy. At the same time, our policy is both robust to estimation errors and interpretable. Numerically, we find that our policy can provide the same efficacy as the status quo with approximately half the capacity for interventions.