



Vukosi Marivate

Department of Computer Science

Rutgers University

Estimating Uncertainty in Personalized Sequential Decision Making Using Observational Clinical Data

vukosi@cs.rutgers.edu

As the amount of data collected from individuals increases, there are more opportunities to use it to offer personalized experiences (e.g., using electronic health records to offer personalized treatments). We advocate applying techniques from batch reinforcement learning to predict the range of effectiveness that policies might have for individuals.

We identify three sources of uncertainty and present a method that addresses all of them. It handles the uncertainty caused by population mismatch by modeling the data as a latent mixture of different subpopulations of individuals. Also, it explicitly quantifies data sparsity by accounting for the limited data available about the underlying models, and incorporates intrinsic stochasticity to yield estimated percentile ranges of the effectiveness of a policy for a particular new individual.

Using this approach, we highlight some interesting variability in policy effectiveness amongst HIV patients given a prior patient treatment dataset. Our approach highlights the potential benefit of taking into account individual variability and data limitations when performing batch policy evaluation for new individuals.

Work done jointly with Jessica Chemali, Emma Brunskill and Michael Littman.