

n=100 per-regime feature selector (mean $\pm$ sd, 10 optimizer seeds):
beats uniform random, not the majority-class baseline

