# **Efficient Online Learning for Optimizing Value of Information**

Yuxin Chen, Jean-Michel Renders, Morteza Haghir Chehreghani, Andreas Krause







- Greedy is near-optimal:  $cost(\pi) = (2 \ln(1/p_{min}) + 1) \cdot cost(\pi^*)$
- **Exponential run time** in terms of the number of available tests

### **Practical concerns**

• Existing approaches (e.g.,  $EC^2$ ) assume **a** known prior  $P(Y, X_T)$ 

# **Our Contribution**

A principled, yet efficient

- sampling-based online learning framework,
- built upon *submodular surrogates*

## **Iterative Filtering & Resampling**

merging hypotheses, resampling whenever necessary

**Thm**: With probability at least  $I - \eta$ , running EC<sup>2</sup> with dynamic hypothesis enumeration outputs an  $\epsilon$ -best decision with cost at most

$$2\ln\left(\frac{1-\eta}{\tilde{p}_{\min}}\right)+1\right)\cdot OPT$$

0

minimal probability of hypotheses in the sampled set

worst-case cost of the optimal algorithm with  $\epsilon$ -optimal decision

5 0.55 20 30 20 20 40 60 80 10 30 Number of iterations #Samples per root-cause #Samples per root-cause Results Full info • EC2+DHE consistently outperforms utility -5 the alternatives in terms of label OnlineVoI age -3 complexity (a.k.a., cost) static

10

 $\times 10^4$ 

Number of epochs

• Average utility (a.k.a reward) of OnlineVol approaches the optimal utility over time.