

Adapting to the Shifting Intent of Search Queries

Umar Syed, Aleksandrs Slivkins, Nina Mishra

- Our model:

Serving search engine results



Contextual bandit problem (Langford & Zhang NIPS 2007).

- New wrinkle: Context *only* reveals whether the best arm has changed, *not* identity of best new arm.
- Results: $O((k + d)(\log T))$ regret, where k is number of times best arm changes, and d is problem-specific constant.
- Previous algorithms: $\Omega(\sqrt{T})$ regret.