EFFICIENT NONMYOPIC ACTIVE SEARCH

Shali Jiang, Gustavo Malkomes,
Geoff Converse, Alyssa Shofner,
Roman Garnett, Benjamin Moseley
Washington University in St. Louis

12.10.16

1. ACTIVE SEARCH

Finding interesting points

Active search[1]

- In *active search*, we consider active learning with an unusual goal: *locating as many members of a particular class as possible.*
- Numerous real-world examples:
  - drug discovery,
  - intelligence analysis,
  - product recommendation,
  - playing Battleship.

[1] Garnett, Krishnamurthy, Xiong, Schneider (CMU), Mann (Uppsala). ICML 2012.

Battleship!

Another definition

Active search is Bayesian optimization with *binary rewards* and *cumulative regret.*

Our approach

We approach this problem via *Bayesian decision theory.*

- We define a natural *utility function,* and
- The location of the next evaluation will be chosen by *maximizing the expected utility.*

The utility function (cumulative reward)

The natural utility function for this problem is *the number of interesting points found.*

The Bayesian optimal policy

The optimal policy may be derived by sequentially maximizing the expected utility of the *final dataset.* With a budget of B, at time t, we select

$$\begin{aligned} &\arg\max_{x_t} \mathbb{E}\big[u(\mathcal{D}_B) \mid x_t, \mathcal{D}_{t-1}\big] \\ =&\arg\max_{x_t} [\text{expected utility starting from point } x_t]. \end{aligned}$$

The Bayesian optimal policy

This may be written *recursively:*

$$
\begin{aligned}
&[\text{expected utility starting from point}] = \\
&\qquad [\text{current utility}] + \\
&\qquad \underbrace{[\text{expected utility of point}]}_{\text{exploitation, } < 1} + \\
&\qquad \underbrace{\mathbb{E}_{y_t}\Big[[\text{success of remaining search}]\Big]}_{\text{exploration, } < B - t}.
\end{aligned}
$$

Automatic *dynamic* tradeoff between exploration and exploitation!

Lookahead

- Unfortunately, the computational cost of computing the optimal policy is *expensive.* (Exponential in the number of points!)
- In practice, we use a *myopic approximation,* where we effectively pretend there is only a small number of observations remaining.

The Bayesian optimal policy

$$
\begin{aligned}
\text{[expected utility starting from point]} = \\
& \text{[current utility]} + \\
& \underbrace{\text{[expected utility of point]}}_{\text{exploitation, } < 1} + \\
& \underbrace{\mathbb{E}_{y_t}\Big[\text{[success of remaining search]}\Big]}_{\text{exploration, } < B - t}.
\end{aligned}
$$

ℓ-step myopic approximation

$$
\begin{aligned}
&[\text{expected utility of next few points}] = \\
&\qquad [\text{current utility}] + \\
&\qquad \underbrace{[\text{expected utility of point}]}_{\text{exploitation, } < 1} + \\
&\qquad \underbrace{\mathbb{E}_{y_t}\Big[[\text{success of next few points}]\Big]}_{\text{exploration, } < \ell}.
\end{aligned}
$$

(ℓ is normally 2–3).

Problems

- The dependence on the budget has been *lost!*
- Exploration is heavily *undervalued!*

Lookahead can always help

Theorem (Garnett, et al.)

Let $\ell, m \in \mathbb{N}^+, \ell < m$. *For any* $q > 0$, *there exists a search problem* $\mathcal{P}$ *such that*

$$\frac{\mathbb{E}_{\mathcal{D}}\big[u(\mathcal{D}) \mid m, \mathcal{P}\big]}{\mathbb{E}_{\mathcal{D}}\big[u(\mathcal{D}) \mid \ell, \mathcal{P}\big]} > q;$$

that is, the m*-step active-search policy can outperform the* ℓ*-step policy by any arbitrary degree.*

Our idea: Efficient nonmyopic active search

- Our idea is to approximate the remainder of the search differently. We assume that any remaining budget is selected *simultaneously in one big batch.*
- Similar idea to the GLASSES algorithm, in a different context (and in this case, *exact* and *efficient*).
- Exploration encouraged *correctly!* Automatic, *dynamic* tradeoff restored!

2. QUICK EXPERIMENT

CiteSeer data

- Includes papers from the 50 most popular venues present in the CiteSeer database.
- 42k nodes, 222k edges.
- We search for *NIPS* papers, 2.5k papers (6%).

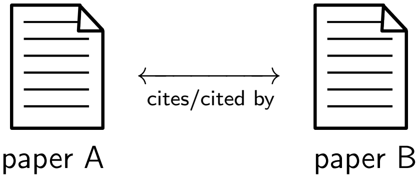

Experiment

- We select a single NIPS paper at random, and begin with that single positive observation.
- The one- and two-step myopic approximations were compared with our method (ENS).

Results

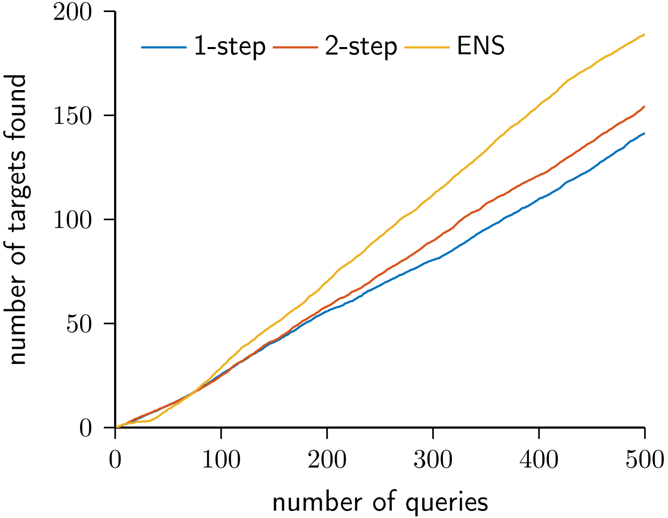

1-step
2-step
ENS
number of targets found
number of queries
0
50
100
150
200
0
100
200
300
400
500

Results: Zoom

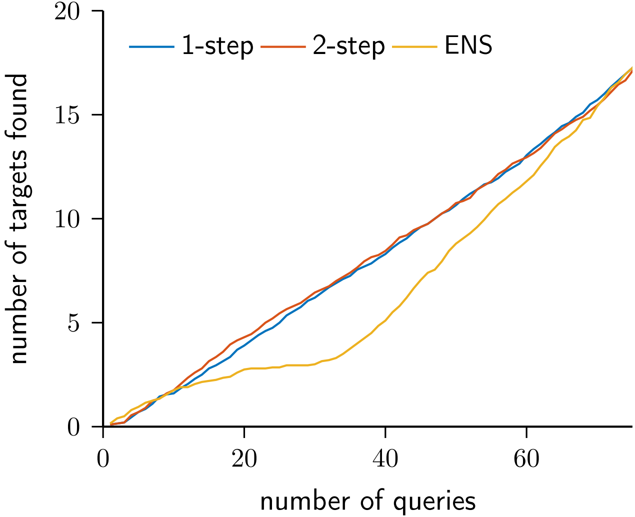

Results: Budget

	query number				
policy	100	300	500	700	900
one-step	25.5	80.5	141	209	273
two-step	24.9	89.8	155	220	287
ENS–900	25.9	94.3	163	239	**308**
ENS–700	28.0	105	188	**259**	
ENS–500	28.7	**112**	**189**		
ENS–300	26.4	105			
ENS–100	**30.7**				

2. THANK YOU!

Questions?