

Online Learning to Rank in Stochastic Click Models

Zoghi, M., Tunys, T., Ghavamzadeh, M., Kveton, B., Szepesvari, C., & Wen, Z.
ICML 2017

Kuhwan Jeong ¹

¹Department of Statistics, Seoul National University, South Korea

July, 2018

Introduction

- *Click model* is a probabilistic model of how users examine and click on a list of documents.
- Click models is designed to explain the so-called position bias in click data.
- Goal of learning to rank (LTR) is to provide a list of K documents out of L that maximizes the number of clicks made by users.
- In this paper, an online LTR algorithm, which can be used in a broad class of click models, is proposed.

Click Models

- All documents (items) are represented by $\mathcal{D} = \{1, \dots, L\}$.
- User is presented an ordered list of K documents out of L , denoted by $\mathcal{R} = (d_1, \dots, d_K)$.
- Click models are parametrized by *attraction function* α , where $\alpha(d)$ is the probability that item d is attractive, and *examination function* χ , where $\chi(\alpha, \mathcal{R}, k)$ is the probability that the k -th item in \mathcal{R} is examined.
- The probability of clicking k -th item in \mathcal{R} is $\chi(\alpha, \mathcal{R}, k)\alpha(d_k)$.
- The expected number of click on list \mathcal{R} is

$$r(\mathcal{R}) = \sum_{k=1}^K \chi(\alpha, \mathcal{R}, k)\alpha(d_k).$$

- Without loss of generality, let $\alpha(1) \geq \dots \geq \alpha(L)$.

Click Models - (1) Position-Based Model

- Position-based model (PBM) assumes the examination function depends only on the position

$$\chi(\alpha, \mathcal{R}, k) = \chi(k).$$

- The expected number of clicks on \mathcal{R} is

$$r(\mathcal{R}) = \sum_{k=1}^K \chi(k) \alpha(d_k).$$

- Under the assumption $\chi(1) \geq \dots \geq \chi(K)$, $r(\mathcal{R})$ is maximized by the list of K most attractive items

$$\mathcal{R}^* = (1, \dots, K)$$

where the k -th most attractive item is placed at position k .

Click Models - (2) Cascade Model

- In the cascade model (CM), the user scans a list of items $\mathcal{R} = (d_1, \dots, d_K)$ from d_1 to d_K .
- If item d_k is *attractive*, the user *clicks* on it and does not examine the remaining items.
- If item d_k is not attractive, the user *examines* item d_{k+1} .
- The examine function χ becomes

$$\chi(\alpha, \mathcal{R}, k) = \prod_{i=1}^{k-1} (1 - \alpha(d_i)).$$

- The expected number of clicks on \mathcal{R} is

$$r(\mathcal{R}) = \sum_{k=1}^K \chi(k) \alpha(d_k) = 1 - \prod_{k=1}^K (1 - \alpha(d_k)).$$

- Any permutation of $\{1, \dots, K\}$ is optimal in the CM.

Online Learning to Rank

Algorithm 1 BatchRank

```
1: // Initialization
2: for  $b = 1, \dots, 2K$  do
3:   for  $l = 0, \dots, T - 1$  do
4:     for all  $d \in \mathcal{D}$  do
5:        $c_{b,l}(d) \leftarrow 0, n_{b,l}(d) \leftarrow 0$ 
6:     end for
7:   end for
8: end for
9:  $\mathcal{B} \leftarrow \{1\}, b_{\max} \leftarrow 1,$ 
10:  $I_1 \leftarrow (1, K), B_{1,0} \leftarrow \mathcal{D}, l_1 \leftarrow 0$ 
11: // BatchRank
12: for  $t = 1, \dots, T$  do
13:   for all  $b \in \mathcal{B}$  do
14:     DisplayBatch( $b, t$ )
15:   end for
16:   for all  $b \in \mathcal{B}$  do
17:     CollectClicks( $b, t$ )
18:   end for
19:   for all  $b \in \mathcal{B}$  do
20:     UpdateBatch( $b, t$ )
21:   end for
22: end for
```

- Two key ideas are
 - ① randomizing the placement of items to avoid position biases and
 - ② recursively dividing the batches of items into more and less attractive items.
- \mathcal{B} : the set indices of batches
- $B_{b,l}$: the set of items in batch b at stage l
- I_b : the range of positions of items in batch b

Algorithm 2 DisplayBatch

1: $l \leftarrow l_b$
2: Let $d_1, \dots, d_{|B_{b,l}|}$ be a permutation of items in $B_{b,l}$ such that $n_{b,l}(d_1) \leq \dots \leq n_{b,l}(d_{|B_{b,l}|})$.
3: Let π be a random permutation of $\{I_b(1), \dots, I_b(2)\}$.
4: **for** $k = I_b(1), \dots, I_b(2)$ **do**
5: $d_k^t \leftarrow d_{\pi(k - I_b(1) + 1)}$
6: **end for**

Algorithm 3 CollectClicks

1: $l \leftarrow l_b, n_{\min} \leftarrow \min_{d \in B_{b,l}} n_{b,l}(d)$
2: **for** $k = I_b(1), \dots, I_b(2)$ **do**
3: **if** $n_{b,l}(d_k^t) = n_{\min}$ **then**
4: $c_{b,l}(d_k^t) \leftarrow c_{b,l}(d_k^t) + c_t(k)$
5: $n_{b,l}(d_k^t) \leftarrow n_{b,l}(d_k^t) + 1$
6: **end if**
7: **end for**

Algorithm 4 UpdateBatch

```

1:  $l \leftarrow l_b$ ,  $\text{len}(b) = I_b(2) - I_b(1) + 1$ 
2: if  $\min_{d \in B_{b,l}} n_{b,l}(d) = n_l$  then
3:   for all  $d \in B_{b,l}$  do
4:     Compute  $U_{b,l}(d)$  and  $L_{b,l}(d)$ .
5:   end for
6:   Let  $d_1, \dots, d_{|B_{b,l}|}$  be a permutation of items in  $B_{b,l}$  such that  $L_{b,l}(d_1) \geq \dots \geq L_{b,l}(d_{|B_{b,l}|})$ .
7:   // Find a split
8:    $s \leftarrow 0$ 
9:   for  $k = 1, \dots, \text{len}(b) - 1$  do
10:     $B_k^+ \leftarrow \{d_1, \dots, d_k\}$ ,  $B_k^- \leftarrow B_{b,l} \setminus B_k^+$ 
11:    if  $L_{b,l}(d_k) > \max_{d \in B_k^-} U_{b,l}(d)$  then
12:       $s \leftarrow k$ 
13:    end if
14:  end for
15:  if  $s = 0$ ,  $|B_{b,l}| > \text{len}(b)$  then
16:     $B_{b,l+1} \leftarrow \{d \in B_{b,l} : U_{b,l}(d) \geq L_{b,l}(d_{\text{len}(b)})\}$ 
17:     $l_b \leftarrow l_b + 1$ 
18:  else if  $s > 0$  then
19:    // Split
20:     $\mathcal{B} \leftarrow \{b_{\max} + 1, b_{\max} + 2\} \setminus \{b\}$ ,  $b_{\max} \leftarrow b_{\max} + 2$ 
21:     $I_{b_{\max}+1} \leftarrow (I_b(1), I_b(1) + s - 1)$ ,  $B_{b_{\max}+1,0} \leftarrow B_s^+$ ,  $l_{b_{\max}+1} \leftarrow 0$ 
22:     $I_{b_{\max}+2} \leftarrow (I_b(1) + s, I_b(2))$ ,  $B_{b_{\max}+2,0} \leftarrow B_s^-$ ,  $l_{b_{\max}+2} \leftarrow 0$ 
23:  end if
24: end if

```

Online Learning to Rank

- Any item $d \in B_{b,l}$ in stage l is explored n_l times where

$$n_l = \lceil 2^{2l+4} \log T \rceil.$$

- At the end of the stage, the probability of clicking on item d is estimated as

$$\hat{c}_{b,l}(d) = c_{b,l}(d)/n_l.$$

- KL-UCB upper and lower confidence bounds are

$$U_{b,l}(d) \leftarrow \operatorname{argmax}_{q \in [\hat{c}_{b,l}(d), 1]} \left\{ D_{\text{KL}}(\hat{c}_{b,l}(d) || q) \leq \frac{\log T + 2 \log \log T}{n_l} \right\},$$

$$L_{b,l}(d) \leftarrow \operatorname{argmin}_{q \in [0, \hat{c}_{b,l}(d)]} \left\{ D_{\text{KL}}(\hat{c}_{b,l}(d) || q) \leq \frac{\log T + 2 \log \log T}{n_l} \right\},$$

where $D_{\text{KL}}(p||q)$ denotes the *Kullback-Leibler divergence* between Bernoulli random variables with means p and q .

Experiments

- *Yandex* dataset
 - A dataset of 35M search sessions, each of which may contain multiple search queries
 - Each query is associated with displayed documents at positions 1 to 10 and their clicks.
- 60 frequent search queries are selected.
- For each query, CM and PBM are learned using PyClick.
- For each query, the goal is to rerank $L = 10$ most attractive items with the objective of maximizing the expected number of clicks at the first $K = 5$ positions.

Experiments

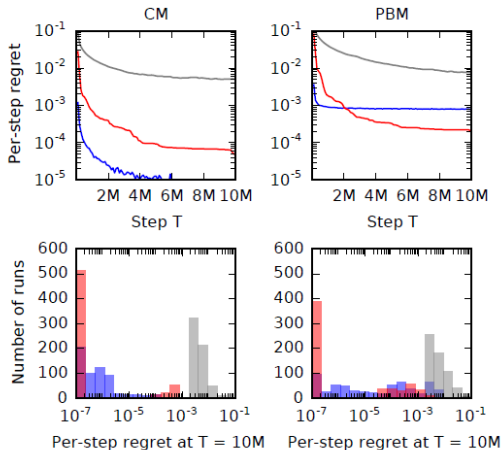


Figure 2. The comparison of BatchRank (red), CascadeKL-UCB (blue), and RankedExp3 (gray) in the CM and PBM. In the top plots, we report the per-step regret as a function of time T , averaged over 60 queries and 10 runs per query. In the bottom plots, we show the distribution of the regret at $T = 10M$.