Supplementary Material

A. Proofs

Below are proofs for the regret bounds from Sections 5 and 6.

A.1. Proof of Theorem 2

First, we bound $\mathbf{E}[\|\mathbf{w}_{T+1}\|^2]$:

$$\begin{aligned}\mathbf{E}[\mathbf{w}_{T+1}^\top \mathbf{w}_{T+1}] &= \mathbf{E}[\mathbf{w}_T^\top \mathbf{w}_T + 2\mathbf{w}_T^\top \phi(\mathbf{x}_T, \bar{\mathbf{y}}_T) \\ &\quad - 2\mathbf{w}_T^\top \phi(\mathbf{x}_T, \mathbf{y}_T) + \|\phi(\mathbf{x}_T, \bar{\mathbf{y}}_T) - \phi(\mathbf{x}_T, \mathbf{y}_T)\|^2] \\ &\le \mathbf{w}_1^\top \mathbf{w}_1 + 2\sum_{t=1}^{T} \mathbf{E}[\mathbf{w}_t^\top \phi(\mathbf{x}_t, \bar{\mathbf{y}}_t) - \mathbf{w}_t^\top \phi(\mathbf{x}_t, \mathbf{y}_t)] + 4R^2 T \\ &\le (4R^2 + 2\Delta)T\end{aligned}$$

The first line utilizes the update rule from algorithm 2. The second line follows from $\|\phi(\mathbf{x}, \mathbf{y})\| \le R$ and repeating the inequality for $t = T-1, \cdots, 1$. The last inequality uses the premise on affirmativeness.

Using the update rule again, we get:

$$\begin{aligned}\mathbf{E}[\mathbf{w}_{T+1}^\top \mathbf{w}_*] &= \mathbf{E}[\mathbf{w}_T^\top \mathbf{w}_* + (\phi(\mathbf{x}_T, \bar{\mathbf{y}}_T) - \phi(\mathbf{x}_T, \mathbf{y}_T))^\top \mathbf{w}_*] \\ &= \sum_{t=1}^{T} \mathbf{E}[(U(\mathbf{x}_t, \bar{\mathbf{y}}_t) - U(\mathbf{x}_t, \mathbf{y}_t)\ )] \\ &\ge \alpha \sum_{t=1}^{T} (U(\mathbf{x}_t, \mathbf{y}_t^*) - \mathbf{E}[U(\mathbf{x}_t, \mathbf{y}_t)]) - \sum_{t=1}^{T} \xi_t\end{aligned}$$

where the last line uses Eq. (4). Using the Cauchy-Schwarz inequality and concavity of $\sqrt{x}$, we get $\mathbf{E}[\mathbf{w}_{T+1}^\top \mathbf{w}_*] \le \|\mathbf{w}_*\| \mathbf{E}[\|\mathbf{w}_{T+1}\|] \le \|\mathbf{w}_*\| \sqrt{\mathbf{E}[\|\mathbf{w}_{T+1}\|^2]}$ from which the claimed result follows.

A.2. Proof of Corollary 3

Note that:

$$\hat{\mathbf{y}}_t = \text{argmax}_{\mathbf{y}} \mathbf{w}_t^\top \phi(\mathbf{x}_t, \mathbf{y})$$

Therefore:

$$\forall t, \bar{\mathbf{y}}_t : \mathbf{w}_t^\top \phi(\mathbf{x}_t, \bar{\mathbf{y}}_t) \le \mathbf{w}_t^\top \phi(\mathbf{x}_t, \hat{\mathbf{y}}_t)$$

Hence:

$$\begin{aligned}&\forall t : \mathbf{E}\Big[\mathbf{w}_t^\top \phi(\mathbf{x}_t, \bar{\mathbf{y}}_t)\Big] - \mathbf{E}\Big[\mathbf{w}_t^\top \phi(\mathbf{x}_t, \mathbf{y}_t)\Big] \\ &\le \mathbf{w}_t^\top \phi(\mathbf{x}_t, \hat{\mathbf{y}}_t) - \mathbf{E}\Big[\mathbf{w}_t^\top \phi(\mathbf{x}_t, \mathbf{y}_t)\Big] \qquad (10)\end{aligned}$$

Given the condition of the corollary, and the above Equation 10, we get that:

$$\frac{1}{T} \sum_{t=1}^{T} \mathbf{E}\Big[\mathbf{w}_t^\top \phi(\mathbf{x}_t, \bar{\mathbf{y}}_t)\Big] - \mathbf{E}\Big[\mathbf{w}_t^\top \phi(\mathbf{x}_t, \mathbf{y}_t)\Big] \le \Omega$$

which using Theorem 2 gives us the corresponding regret bound.

A.3. Proof of Theorem 4

This proof is very similar to the one in (Raman et al., 2012), though it solves a different problem. In particular since:

$$\forall t : \mathbf{E}\Big[\mathbf{w}_t^\top \phi(\mathbf{x}_t, \mathbf{y}_t)\Big] \ge (1-\beta) \mathbf{w}_t^\top \phi(\mathbf{x}_t, \hat{\mathbf{y}}_t)$$

we have that:

$$\mathbf{E}[\mathbf{w}_t^\top (\phi(\mathbf{x}_t, \bar{\mathbf{y}}_t) - \phi(\mathbf{x}_t, \mathbf{y}_t))] \le \beta \mathbf{w}_t^\top \phi(\mathbf{x}_t, \hat{\mathbf{y}}_t)$$

From here on, the proof from (Raman et al., 2012) can be used, to prove the corresponding regret bound. Thus in other words, the perturbation can be thought of as a way to produce an $(1-\beta)$-*approximate* solution to the argmax problem.

A.4. Proof of Proposition 5

Consider the case when documents in positions i and $i+1$ (call them d_i and d_{i+1}) are swapped[2]:

$$\begin{aligned}&\mathbf{w}_t^\top (\gamma_i - \gamma_{i+1})(\phi(\mathbf{x}_t, d_i) - \phi(\mathbf{x}_t, d_{i+1})) \\ &\le \Big(1 - \frac{\gamma_{i+1}}{\gamma_i}\Big) \mathbf{w}_t^\top (\gamma_i \phi(\mathbf{x}_t, d_i) + \gamma_{i+1} \phi(\mathbf{x}_t, d_{i+1}))\end{aligned}$$

Note that this factor $1 - \frac{\gamma_{i+1}}{\gamma_i}$ is largest for $i = 1$. Thus we can state for every swapped pair:

$$\begin{aligned}&\mathbf{w}_t^\top (\gamma_i - \gamma_{i+1})(\phi(\mathbf{x}_t, d_i) - \phi(\mathbf{x}_t, d_{i+1})) \\ &\le \Big(1 - \frac{\gamma_2}{\gamma_1}\Big) \mathbf{w}_t^\top (\gamma_i \phi(\mathbf{x}_t, d_i) + \gamma_{i+1} \phi(\mathbf{x}_t, d_{i+1}))\end{aligned}$$

Summing this over all swapped pairs, and using the fact that each pair has some probability p to be swapped:

$$\begin{aligned}&\mathbf{w}_t^\top (\phi(\mathbf{x}_t, \hat{\mathbf{y}}_t) - \mathbf{E}[\phi(\mathbf{x}_t, \mathbf{y}_t)]) \\ &\le p\Big(1 - \frac{\gamma_2}{\gamma_1}\Big) \mathbf{w}_t^\top \phi(\mathbf{x}_t, \hat{\mathbf{y}}_t)\end{aligned}$$

A.5. Proof of Proposition 6

We prove a more general proposition here:

Proposition 7 *For $\Delta \ge 0$, dynamically setting the swap prob. of 3PR to be*

$$p_t \le max\Big(0, min\Big(1, c(\Delta \cdot t - R_t)\Big)\Big), \qquad (11)$$

[2]This holds assuming the inner products with documents are non-negative. Thus algorithmically this can be implemented by only ranking documents with non-negative scores.

for some positive constant c, has regret

$$\leq \frac{1}{\alpha T}\sum_{t=1}^{T}\xi_t + \frac{\|\mathbf{w}_*\|}{\alpha\sqrt{T}}\sqrt{4R^2 + 2\Delta + (\gamma_1 - \gamma_2)R\sqrt{\frac{4R^2+2\Delta}{T}}}.$$

Proof We prove this by using Theorem 2. In particular, we show:

$$\frac{1}{T}\sum_{t=1}^{T}\mathbf{w}_t^\top(\phi(\mathbf{x}_t, \bar{\mathbf{y}}_t) - \phi(\mathbf{x}_t, \mathbf{y}_t)) < \Delta + \Gamma\sqrt{\frac{4R^2+2\Delta}{T}} \quad (12)$$

where $\Gamma = (\gamma_1 - \gamma_2)R$. We will show this holds by induction on T. Note that this condition trivially holds for $T = 0$ (base case). Now assume it holds for $T = k - 1$. We will show it is true for $T = k$. Consider the cumulative affirmativeness $R_k = \sum_{i=1}^{k-1}\mathbf{w}_i^\top\phi(\mathbf{x}_i, \bar{\mathbf{y}}_i) - \mathbf{w}_i^\top\phi(\mathbf{x}_i, \mathbf{y}_i)$. There are 2 cases to consider:

- $R_k \geq k\Delta$: If this is the case $p_k = 0$ *i.e.,* no perturbation is performed for iteration k and hence $\mathbf{y}_k = \hat{\mathbf{y}}_k = \text{argmax}_{\mathbf{y}}\mathbf{w}_k^\top\phi(\mathbf{x}_k, \mathbf{y})$. Therefore $\mathbf{w}_k^\top(\phi(\mathbf{x}_k, \bar{\mathbf{y}}_k) - \phi(\mathbf{x}_k, \mathbf{y}_k)) \leq 0$; thus $R_{k+1} \leq R_k$ and hence the induction hypothesis is satisfied.
- $R_k < k\Delta$: We have $\|\mathbf{w}_k\| \leq \sqrt{k(4R^2+2\Delta)}$ as shown in the proof of Thm 2. As per the perturbation, for all $\mathbf{y}_k$ we have $\|\phi(\mathbf{x}_k, \hat{\mathbf{y}}_k) - \phi(\mathbf{x}_k, \mathbf{y}_k)\| \leq \Gamma$[3]. Next by Cauchy-Schwarz we get $\mathbf{w}_k^\top(\phi(\mathbf{x}_k, \hat{\mathbf{y}}_k) - \phi(\mathbf{x}_k, \mathbf{y}_k)) \leq \|\mathbf{w}_k\|\Gamma$. Thus $R_{k+1} \leq R_k + \Gamma\sqrt{k(4R^2+2\Delta)}$; hence satisfying the induction hypothesis.

Thus the induction holds for $T = k$. Since equation (12) holds for all $\mathbf{y}_t, \bar{\mathbf{y}}_t$, this condition is also satisfied under expectation (over $\mathbf{y}_t, \bar{\mathbf{y}}_t$). Hence the condition for Theorem 2 is satisfied, thus giving us the bound. Note that the second term on the RHS of Eq. (12) asymptotically disappears. ■

B. Additional Details of User Study

The ranking function in the ArXiv search engine used 1000 features which can be categorized into the following three groups.

- Features the corresponded to rank as per query similarity with different components of the document (authors, abstract, article *etc.*.). We used different similarity measures. For each of these document-components and similarity measures, we had multiple features of the form $\texttt{rank} \leq a$, where a was a value we varied to create multiple features (we used $2, 5, 10, 15, 25, 30, 50, 100, 200$).
- Second-order features the represented pairwise combinations of rank (for the default similarity measure) for 2 different document-components.
- Query-independent features representing the document age and the document category (e.g. AI, NLP, ML, Statistics *etc.*.).

Our baseline, was a hand-coded solution using 35 features considered the most important by us.

[3]This assumes that the document feature vectors are component-wise non-negative. If this is not true, then the bound still holds but with $\Gamma = 2R$