



Unbiased Learning to Rank from User Interactions

Harrie Oosterhuis

April 1, 2019

University of Amsterdam

oosterhuis@uva.nl

<https://staff.fnwi.uva.nl/h.r.oosterhuis>

Introduction

Learning to Rank is a **core task** in informational retrieval:

- Key component for **search** and **recommendation**.
- Directly impacts user experience.

Learning to Rank is a **core task** in informational retrieval:

- Key component for **search** and **recommendation**.
- Directly impacts user experience.

Traditionally learning to rank uses **annotated datasets**:

- **Relevance annotations** for query-document pairs provided by **human judges**.

Problems with Supervised Approach

Some of the most substantial limitations of **annotated datasets** are:

- **expensive** to make (Qin and Liu, 2013; Chapelle and Chang, 2011).

Problems with Supervised Approach

Some of the most substantial limitations of **annotated datasets** are:

- **expensive** to make (Qin and Liu, 2013; Chapelle and Chang, 2011).
- **unethical** to create in **privacy-sensitive settings** (Wang et al., 2016).

Problems with Supervised Approach

Some of the most substantial limitations of **annotated datasets** are:

- **expensive** to make (Qin and Liu, 2013; Chapelle and Chang, 2011).
- **unethical** to create in **privacy-sensitive settings** (Wang et al., 2016).
- **impossible** for small scale problems e.g. **personalization**.

Problems with Supervised Approach

Some of the most substantial limitations of **annotated datasets** are:

- **expensive** to make (Qin and Liu, 2013; Chapelle and Chang, 2011).
- **unethical** to create in **privacy-sensitive settings** (Wang et al., 2016).
- **impossible** for small scale problems e.g. **personalization**.
- **stationary**, cannot capture **future changes in relevancy** (Lefortier et al., 2014).

Problems with Supervised Approach

Some of the most substantial limitations of **annotated datasets** are:

- **expensive** to make (Qin and Liu, 2013; Chapelle and Chang, 2011).
- **unethical** to create in **privacy-sensitive settings** (Wang et al., 2016).
- **impossible** for small scale problems e.g. **personalization**.
- **stationary**, cannot capture **future changes in relevancy** (Lefortier et al., 2014).
- **not necessarily aligned with actual user preferences** (Sanderson, 2010),
i.e. annotators and users often disagree.

Learning from User Interactions

Learning from User Interactions: Advantages

Learning from user interactions solves the problems of annotations:

- Interactions are **virtually free** if you have users.
- User **behaviour** is indicative of their **preferences**.

Learning from User Interactions: Advantages

Learning from user interactions solves the problems of annotations:

- Interactions are **virtually free** if you have users.
- User **behaviour** is indicative of their **preferences**.
- Interactions give **implicit feedback**.

Learning from User Interactions: Difficulties

User interactions bring their **own difficulties**:

- **Noise:**
 - Users click for **unexpected reasons**.
 - Often clicks occur **not because** of relevancy.

Learning from User Interactions: Difficulties

User interactions bring their **own difficulties**:

- **Noise:**

- Users click for **unexpected reasons**.
- Often clicks occur **not because** of relevancy.
- Often clicks do not occur **despite** of relevancy.

Learning from User Interactions: Difficulties

User interactions bring their **own difficulties**:

- **Noise:**

- Users click for **unexpected reasons**.
- Often clicks occur **not because** of relevancy.
- Often clicks do not occur **despite** of relevancy.

- **Bias:** Interactions are affected by **factors other than relevancy**:

Learning from User Interactions: Difficulties

User interactions bring their **own difficulties**:

- **Noise:**

- Users click for **unexpected reasons**.
- Often clicks occur **not because** of relevancy.
- Often clicks do not occur **despite** of relevancy.

- **Bias:** Interactions are affected by **factors other than relevancy**:

- **Position bias:** **Higher ranked** documents get more attention.

Learning from User Interactions: Difficulties

User interactions bring their **own difficulties**:

- **Noise:**

- Users click for **unexpected reasons**.
- Often clicks occur **not because** of relevancy.
- Often clicks do not occur **despite** of relevancy.

- **Bias:** Interactions are affected by **factors other than relevancy**:

- **Position bias:** **Higher ranked** documents get more attention.
- **Item selection bias:** Interactions are **limited** to the **presented** documents.

Learning from User Interactions: Difficulties

User interactions bring their **own difficulties**:

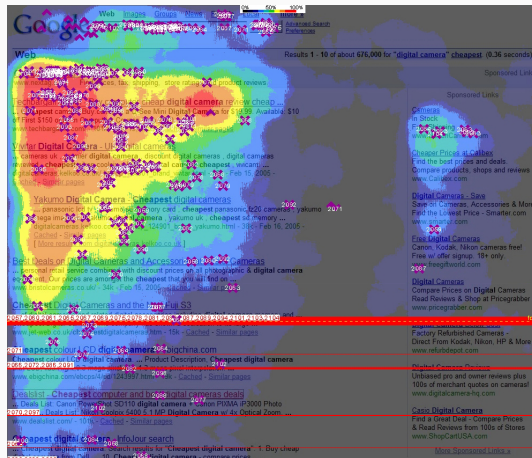
- **Noise:**

- Users click for **unexpected reasons**.
- Often clicks occur **not because** of relevancy.
- Often clicks do not occur **despite** of relevancy.

- **Bias:** Interactions are affected by **factors other than relevancy**:

- **Position bias:** **Higher ranked** documents get more attention.
- **Item selection bias:** Interactions are **limited** to the **presented** documents.
- **Presentation bias:** Results that are **presented differently** will be **treated different**.
- ...

The Golden Triangle



Source: <http://www.mediative.com/>

Goal of unbiased learning to rank:

- Optimize a ranker w.r.t. **relevance preferences** of users from their interactions.
- **Avoid** being **biased by other factors** that influence interactions.

Learning from Historical Interactions:

- Learn/estimate a **model of user behaviour** including their biases.
- Learn from historical data while **adjusting** for these **biases**.

Online Learning to Rank:

- Algorithms that can **intervene** during the learning process.
- Handle biases by having **control over displayed results**.

Learning from Historical Data

Learning from Historical Data: Introduction

For a search system that has **active users**:

- interaction logs are **easy to obtain** and with **no significant cost**.

Learning from Historical Data: Introduction

For a search system that has **active users**:

- interaction logs are **easy to obtain** and with **no significant cost**.

From interactions logs a **model of user behaviour** can be learned.

- User modelling is a long-established field (Chuklin et al., 2015).

Learning from Historical Data: Introduction

For a search system that has **active users**:

- interaction logs are **easy to obtain** and with **no significant cost**.

From interactions logs a **model of user behaviour** can be learned.

- User modelling is a long-established field (Chuklin et al., 2015).

By **modelling user biases** accurately, they can be **adjusted for effectively**.

History of this approach:

- **Click models:**

Learn models that predict user behaviour (Chuklin et al., 2015).

History of this approach:

- **Click models:**
Learn models that predict user behaviour (Chuklin et al., 2015).
- **Unbiased learning:**
Infer the effect of position bias on selection bias in click-data,
then correct for it in learning to rank for instant email-results (Wang et al., 2016).

History of this approach:

- **Click models:**
Learn models that predict user behaviour (Chuklin et al., 2015).
- **Unbiased learning:**
Infer the effect of position bias on selection bias in click-data,
then correct for it in learning to rank for instant email-results (Wang et al., 2016).
- **Counter-factual learning:**
Recast the approach with counter-factual learning theory,
and broaden applicability to general ranking problems (Joachims et al., 2017).



Under assumption of **binary relevance**:






$$r(y) \in \{0, 1\}$$

The **loss** for a single query \mathbf{q} , a ranking \mathbf{y} , and the relevance r :













$$\Delta(\mathbf{y}|\mathbf{q}, r) = \sum_{y \in \mathbf{y}} \text{rank}(y|\mathbf{y}) \cdot r(y). \quad (1)$$

Minimizing this loss is **trivial** if the relevances r are known (**full-information setting**).

Ranking	Relevance
document 1	
document 2	
document 3	
document 4	
document 5	

Ranking	Relevance	Position Bias
document 1	★	
document 2	✗	
document 3	✗	
document 4	✗	
document 5	★	

IPS: Position Bias as a Filter

Ranking	Relevance	Position Bias	Clicks
document 1			
document 2			
document 3			
document 4			
document 5			

IPS: Bias in Naive Approach

Assume that every **relevant** document that is **observed** is **clicked** (no noise).

Due to **biases** the **probability** of a document y being **observed** in ranking \bar{y} is:

$$Q(o(y) = 1 | \mathbf{q}, \bar{y}, r), \quad (2)$$

this is called the *propensity*.

IPS: Bias in Naive Approach

Assume that every **relevant** document that is **observed** is **clicked** (no noise).

Due to **biases** the **probability** of a document y being **observed** in ranking $\bar{\mathbf{y}}$ is:

$$Q(o(y) = 1 | \mathbf{q}, \bar{\mathbf{y}}, r), \quad (2)$$

this is called the *propensity*.

If we **naively estimate the loss** from clicks it is **biased**:

$$\begin{aligned} \mathbb{E}_o[\hat{\Delta}_{\text{naive}}(\mathbf{y} | \mathbf{q}, r)] &= \mathbb{E}_o\left[\sum_{y: o(y)=1 \wedge r(y)=1} \text{rank}(y | \mathbf{y}) \cdot r(y)\right] \\ &= \sum_{y \in \mathbf{y}} Q(o(y) = 1 | \mathbf{q}, \bar{\mathbf{y}}, r) \cdot \text{rank}(y | \mathbf{y}) \cdot r(y). \end{aligned} \quad (3)$$

IPS: Inverse Propensity Scoring

However, with **knowledge** of $Q(o(y) = 1|\mathbf{q}, \bar{\mathbf{y}}, r)$ it can be **corrected for** easily:

$$\hat{\Delta}_{IPS}(\mathbf{y}|\mathbf{q}, r) = \sum_{y:o(y)=1 \wedge r(y)=1} \frac{\text{rank}(y|\mathbf{y}) \cdot r(y)}{Q(o(y) = 1|\mathbf{q}, \bar{\mathbf{y}}, r)}. \quad (4)$$

IPS: Inverse Propensity Scoring

However, with **knowledge** of $Q(o(y) = 1|\mathbf{q}, \bar{\mathbf{y}}, r)$ it can be **corrected for** easily:

$$\hat{\Delta}_{IPS}(\mathbf{y}|\mathbf{q}, r) = \sum_{y:o(y)=1 \wedge r(y)=1} \frac{\text{rank}(y|\mathbf{y}) \cdot r(y)}{Q(o(y) = 1|\mathbf{q}, \bar{\mathbf{y}}, r)}. \quad (4)$$


The **expected loss** now gives an **unbiased estimate**:

$$\begin{aligned} \mathbb{E}_o[\hat{\Delta}_{IPS}(\mathbf{y}|\mathbf{q}, r)] &= \mathbb{E}_o\left[\sum_{y:o(y)=1 \wedge r(y)=1} \frac{\text{rank}(y|\mathbf{y}) \cdot r(y)}{Q(o(y) = 1|\mathbf{q}, \bar{\mathbf{y}}, r)}\right] \\ &= \sum_{y \in \mathbf{y}} \frac{Q(o(y) = 1|\mathbf{q}, \bar{\mathbf{y}}, r) \cdot \text{rank}(y|\mathbf{y}) \cdot r(y)}{Q(o(y) = 1|\mathbf{q}, \bar{\mathbf{y}}, r)} \\ &= \sum_{y \in \mathbf{y}} \text{rank}(y|\mathbf{y}) \cdot r(y) \\ &= \Delta(\mathbf{y}|\mathbf{q}, r). \end{aligned}$$

IPS: Propensity Inference

However, the propensities $Q(o(y) = 1 | \mathbf{q}, \bar{\mathbf{y}}, r)$ still have to be inferred.

One possible way is by **swapping document pairs**:



Assuming the difference in click probabilities is **only affected** by the **observance probability**.

Unbiased learning to rank from historical data:

- ① **Infer a user model** potentially from randomized rankings (e.g. swapped pairs).

Unbiased learning to rank from historical data:

- ① **Infer a user model** potentially from randomized rankings (e.g. swapped pairs).
- ② **Gather user-interaction logs** using the production ranker (**no randomization required**).

Unbiased learning to rank from historical data:

- ① **Infer a user model** potentially from randomized rankings (e.g. swapped pairs).
- ② **Gather user-interaction logs** using the production ranker (**no randomization required**).
- ③ Learn an **unbiased ranking model** by applying **counter-factual learning**, using the user model and interaction logs.

Unbiased learning to rank from historical data:

- ① **Infer a user model** potentially from randomized rankings (e.g. swapped pairs).
- ② **Gather user-interaction logs** using the production ranker (**no randomization required**).
- ③ Learn an **unbiased ranking model** by applying **counter-factual learning**, using the user model and interaction logs.

Note that:

- Once an accurate user model is obtained, **randomization is no longer needed**.
- Also proven to be unbiased when **click noise is present**.

Advantages:

- Can **learn unbiasedly** from **user interactions**.
- Learned ranking models **match user preferences** closer than annotations.
- Requires **no randomization** after an accurate user model is learned.

Advantages:

- Can **learn unbiasedly** from **user interactions**.
- Learned ranking models **match user preferences** closer than annotations.
- Requires **no randomization** after an accurate user model is learned.

Disadvantages:

- **Requires an accurate user model**, which may not always be feasible.
- **No comparison** with *online learning to rank* has been performed.

This is still a **very new and active** area of research.

Online Learning to Rank


Online Learning to Rank methods have **control over what to display** to the user Yue and Joachims (2009).

Simultaneously they:

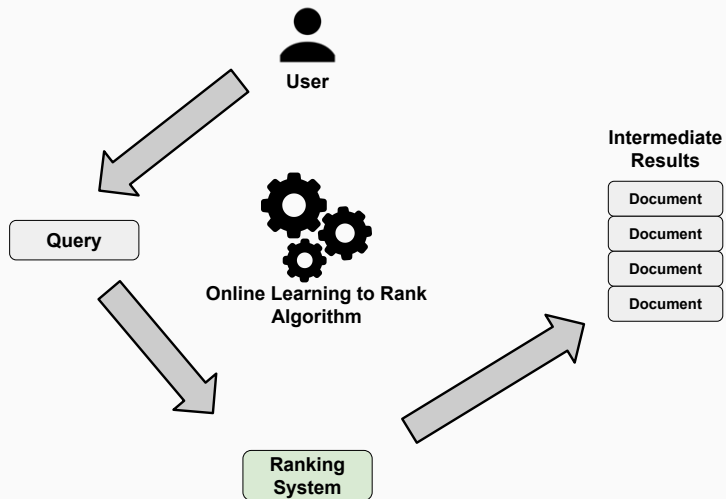
- **Decide** what **results to display** to the user.
- **Learn** from **user interactions** with chosen results.

These methods can be much **more efficient**,
because they have (more) **control over what data is gathered**.

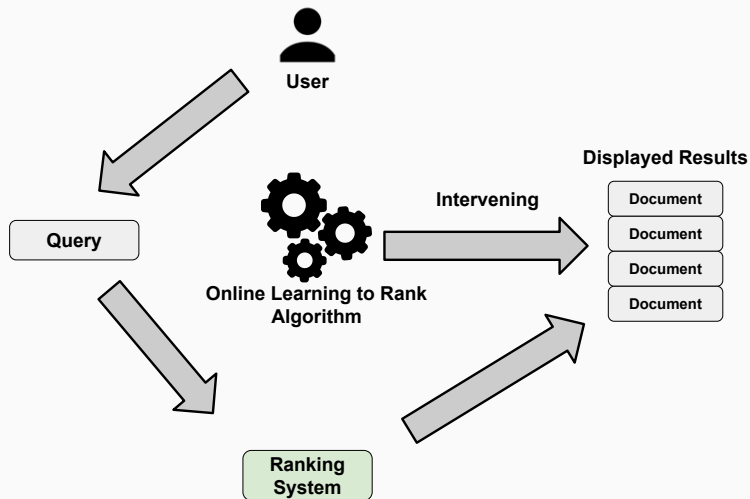
Online Learning to Rank: Visualization



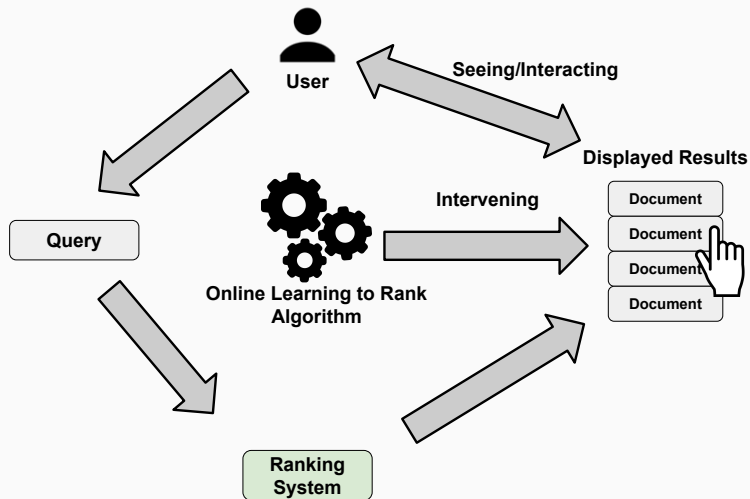
Online Learning to Rank: Visualization



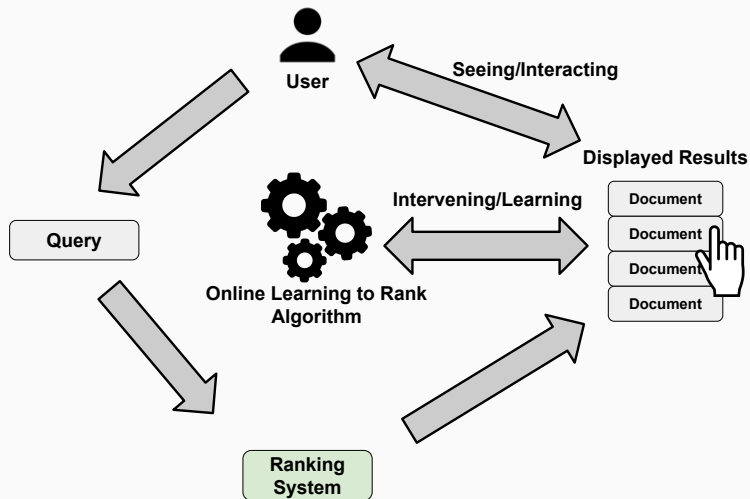
Online Learning to Rank: Visualization



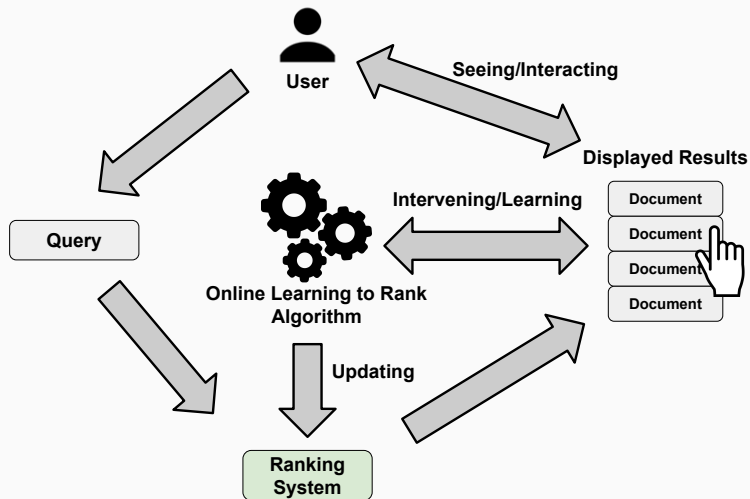
Online Learning to Rank: Visualization



Online Learning to Rank: Visualization



Online Learning to Rank: Visualization



Online learning to rank methods have the **potential advantages**:

- Learn the **true preferences of users** (unlike annotator approaches).

Online learning to rank methods have the **potential advantages**:

- Learn the **true preferences of users** (unlike annotator approaches).
- More **responsive** by **immediately adapting** to users.

Online learning to rank methods have the **potential advantages**:

- Learn the **true preferences of users** (unlike annotator approaches).
- More **responsive** by **immediately adapting** to users.

Online methods also bring a **large risk**:

Online learning to rank methods have the **potential advantages**:

- Learn the **true preferences of users** (unlike annotator approaches).
- More **responsive** by **immediately adapting** to users.

Online methods also bring a **large risk**:

- **Unreliable** methods could **severely worsen the user experience immediately**.

Advantages:

- Online methods learn **faster** by **intervening**, (than learning from historical data).
- By **immediately adapting** they can be more **responsive**.

Advantages:

- Online methods learn **faster** by **intervening**, (than learning from historical data).
- By **immediately adapting** they can be more **responsive**.
- More suitable for cases where **fewer user interactions** are available.

Advantages:

- Online methods learn **faster** by **intervening**, (than learning from historical data).
- By **immediately adapting** they can be more **responsive**.
- More suitable for cases where **fewer user interactions** are available.

Disadvantages:

- The algorithm requires **exploration**, and needs (some) **control** over search results.

Advantages:

- Online methods learn **faster** by **intervening**, (than learning from historical data).
- By **immediately adapting** they can be more **responsive**.
- More suitable for cases where **fewer user interactions** are available.

Disadvantages:

- The algorithm requires **exploration**, and needs (some) **control** over search results.
- **No comparisons** performed with *learning from historical data*.

Pairwise Differentiable Gradient Descent

Pairwise Differentiable Gradient Descent

We recently introduced **Pairwise Differentiable Gradient Descent** (Oosterhuis and de Rijke, 2018):

- Very different from previous Online Learning to Rank methods, that relied on sampling model variations similar to evolutionary approaches.

Intuition:

- A **pairwise** approach can be made **unbiased**, while being **differentiable**, without relying on online evaluation method or the sampling of models.

Pairwise Differentiable Gradient Descent optimizes a **Plackett Luce** ranking model, this models a **probabilistic distribution over documents**.

With the ranking scoring model $f(\mathbf{d}, \theta)$ the distribution is:

$$P(d|D, \theta) = \frac{\exp f(\mathbf{d}, \theta)}{\sum_{d' \in D} \exp f(\mathbf{d}', \theta)} \quad (5)$$

Pairwise Differentiable Gradient Descent optimizes a **Plackett Luce** ranking model, this models a **probabilistic distribution over documents**.


With the ranking scoring model $f(\mathbf{d}, \theta)$ the distribution is:

$$P(d|D, \theta) = \frac{\exp f(\mathbf{d}, \theta)}{\sum_{d' \in D} \exp f(\mathbf{d}', \theta)} \quad (5)$$

Confidence is explicitly modelled and **exploration** depends on the **available documents**, thus it **naturally varies per query** and even within the ranking.


Bias in Pairwise Inference

Similar to existing pairwise methods (Oosterhuis and de Rijke, 2017; Joachims, 2002), Pairwise Differentiable Gradient Descent infers **pairwise document preferences from user clicks**:



Bias in Pairwise Inference

Similar to existing pairwise methods (Oosterhuis and de Rijke, 2017; Joachims, 2002), Pairwise Differentiable Gradient Descent infers **pairwise document preferences from user clicks**:




This approach is **biased**:

- Some preferences are **more likely to be inferred** due to **position/selection bias**.


Reversed Pair Rankings

Let $R^*(d_i, d_j, R)$ be R but with the **positions** of d_i and d_j **swapped**:



Reversed Pair Rankings

Let $R^*(d_i, d_j, R)$ be R but with the **positions** of d_i and d_j **swapped**:




We assume:

- For a preference $d_i \succ d_j$ inferred from ranking R , if both are **equally relevant** the opposite preference $d_j \succ d_i$ is **equally likely** to be inferred from $R^*(d_i, d_j, R)$.

Reversed Pair Rankings

Let $R^*(d_i, d_j, R)$ be R but with the **positions** of d_i and d_j **swapped**:



We assume:

- For a preference $d_i \succ d_j$ inferred from ranking R , if both are **equally relevant** the opposite preference $d_j \succ d_i$ is **equally likely** to be inferred from $R^*(d_i, d_j, R)$.

Then scoring **as if** R and R^* are **equally likely to occur** makes the gradient **unbiased**.

Unbiasing the Pairwise Update

The **ratio** between the probability of the ranking and the reversed pair ranking indicates the **bias between the two directions**:

$$\rho(d_i, d_j, R) = \frac{P(R^*(d_i, d_j, R)|f, D)}{P(R|f, D) + P(R^*(d_i, d_j, R)|f, D)} \quad (6)$$

We use this ratio to **unbias the gradient estimation**:

$$\nabla f(\cdot, \theta) \approx \sum_{d_i \succ_{\mathbf{c}} d_j} \rho(d_i, d_j, R) \nabla P(d_i \succ d_j | D, \theta). \quad (7)$$

Unbiasedness of Pairwise Differentiable Gradient Descent

Under the reversed pair ranking assumption, we prove that **the expected estimated gradient** can be written as:

$$E[\nabla f(\cdot, \theta)] = \sum_{d_i, d_j} \alpha_{ij} (f'(\mathbf{d}_i, \theta) - f'(\mathbf{d}_j, \theta)). \quad (8)$$

Unbiasedness of Pairwise Differentiable Gradient Descent

Under the reversed pair ranking assumption, we prove that **the expected estimated gradient** can be written as:

$$E[\nabla f(\cdot, \theta)] = \sum_{d_i, d_j} \alpha_{ij} (f'(\mathbf{d}_i, \theta) - f'(\mathbf{d}_j, \theta)). \quad (8)$$

Where the weights α_{ij} will **match the user preferences** in expectation:

$$d_i =_{rel} d_j \Leftrightarrow \alpha_{ij} = 0, \quad (9)$$

$$d_i >_{rel} d_j \Leftrightarrow \alpha_{ij} > 0, \quad (10)$$

$$d_i <_{rel} d_j \Leftrightarrow \alpha_{ij} < 0. \quad (11)$$

Thus the estimated gradient is **unbiased w.r.t. document pair preferences**.

Pairwise Differentiable Gradient Descent: Method

Start with initial model θ_t .

Then indefinitely:

- 1 Wait for a user query.
- 2 **Sample** (without replacement) a **ranking** R from the document distribution:

$$P(d|D, \theta_t) = \frac{\exp^{f(\mathbf{d}, \theta_t)}}{\sum_{d' \in D} \exp^{f(\mathbf{d}', \theta_t)}}. \quad (12)$$

- 3 **Display** the ranking R to the user.
- 4 **Infer document preferences** from the **user clicks**: \mathbf{c} .
- 5 **Update** model according to the **estimated (unbiased) gradient**:

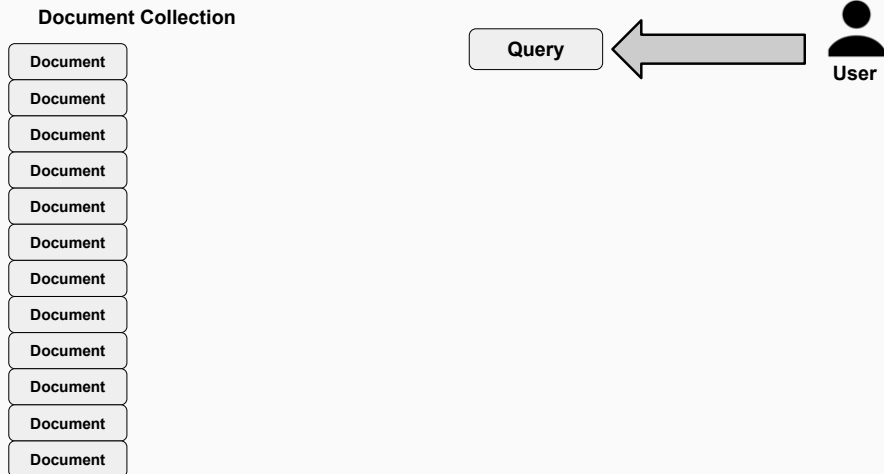
$$\nabla f(\cdot, \theta) \approx \sum_{d_i \succ_{\mathbf{c}} d_j} \rho(d_i, d_j, R) \nabla P(d_i \succ d_j | D, \theta). \quad (13)$$

Pairwise Differentiable Gradient Descent: Visualization

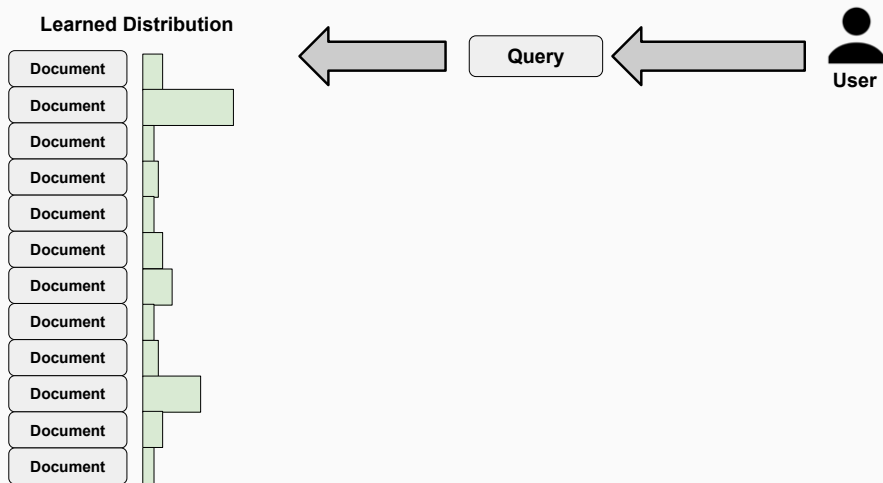
Document Collection



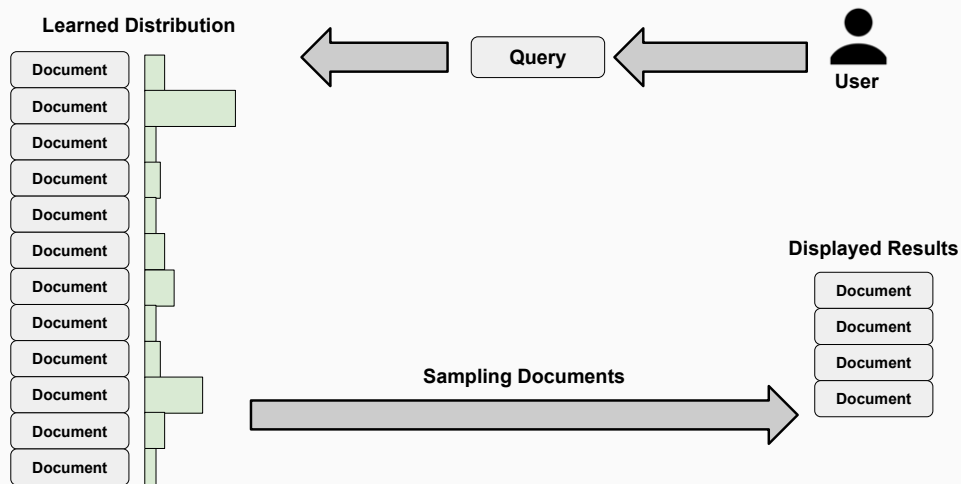
Pairwise Differentiable Gradient Descent: Visualization



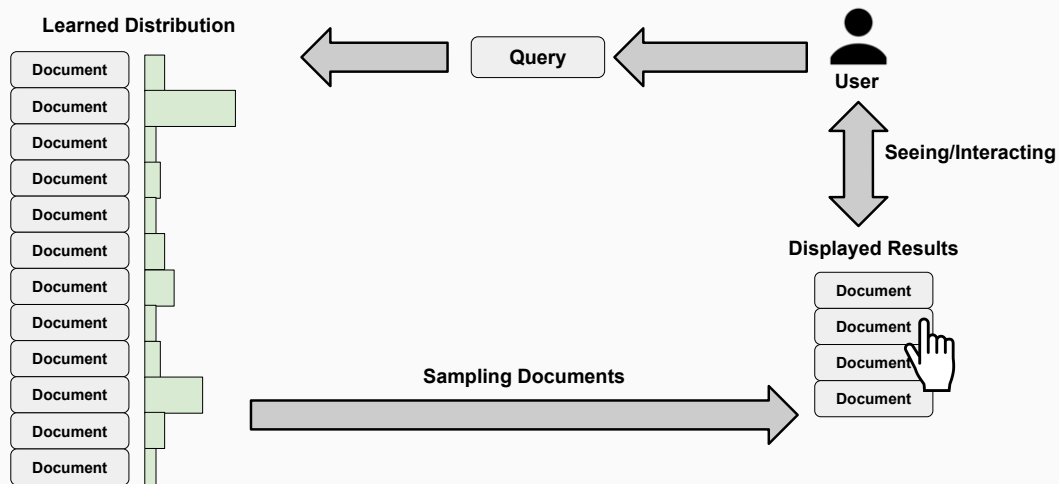
Pairwise Differentiable Gradient Descent: Visualization




Pairwise Differentiable Gradient Descent: Visualization



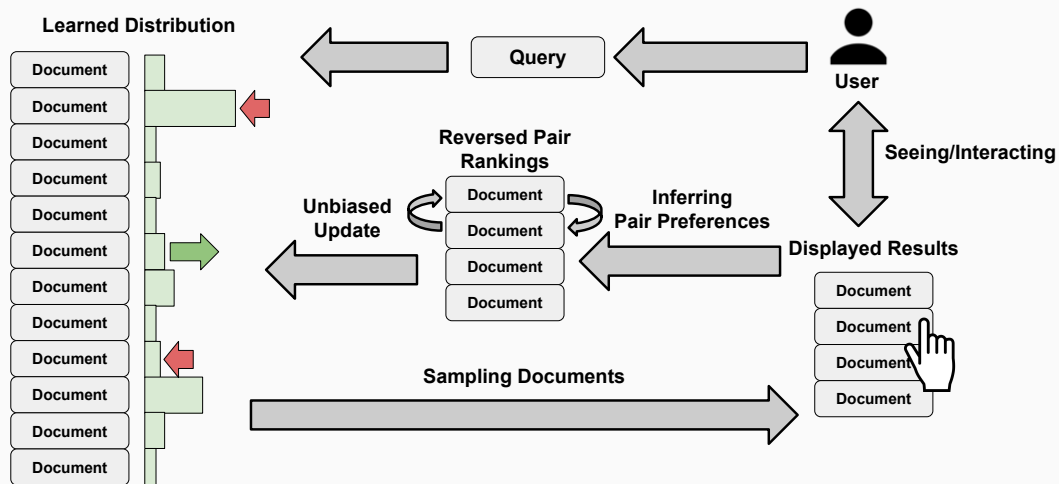
Pairwise Differentiable Gradient Descent: Visualization



Pairwise Differentiable Gradient Descent: Visualization



Pairwise Differentiable Gradient Descent: Visualization



Experimental Results

Experimental Setup

Comparison of Pairwise Differentiable Gradient Descent with **previous Online Learning to Rank methods**.

Simulations based on the annotated learning-to-rank datasets.

- **Largest available industry datasets**: MSLR-Web10k, Yahoo Webscope, Istella.

User behaviour simulated using **cascading click models**.



Experiments **repeated** under **varying levels** of noise and bias.

Pairwise Differentiable Gradient Descent: Results

Results **across all datasets** (MSLR-Web10k, Yahoo Webscope, Istella) we observe:



- **Large improvements** in performance of **convergence** under all levels of noise.
- Much **faster learning (better user experience)** under all levels of noise.

Pairwise Differentiable Gradient Descent: Results Short Term



**Results of simulations on the MSLR-WEB10k dataset,
a perfect user (left) and an informational user (right).**

Pairwise Differentiable Gradient Descent: Results Long Term



**Results of simulations on the MSLR-WEB10k dataset,
a perfect user (left) and an informational user (right).**

Conclusion

Conclusion

Take-away messages:

- **Supervised approaches** to learning to rank are **limited**.
 - **Annotations often disagree** with user preferences.

Conclusion

Take-away messages:

- **Supervised approaches** to learning to rank are **limited**.
 - **Annotations often disagree** with user preferences.
- User interactions solve this problem but bring **noise and biases**.

Take-away messages:

- **Supervised approaches** to learning to rank are **limited**.
 - **Annotations often disagree** with user preferences.
- User interactions solve this problem but bring **noise and biases**.
- **Counter-factual approaches** allow for **unbiased** learning to rank:
 - If an **accurate user model** can be learned, we can **adjust for biases**.
 - **Only** require randomization to **infer a user model**.

Take-away messages:

- **Supervised approaches** to learning to rank are **limited**.
 - **Annotations often disagree** with user preferences.
- User interactions solve this problem but bring **noise and biases**.
- **Counter-factual approaches** allow for **unbiased** learning to rank:
 - If an **accurate user model** can be learned, we can **adjust for biases**.
 - **Only** require randomization to **infer a user model**.
- **Online approaches** allow for **unbiased** and **responsive** learning to rank:
 - **Immediately adapt** to user behaviour.
 - Perform **randomization** at each step, though limited.

Conclusion

Take-away messages:

- **Supervised approaches** to learning to rank are **limited**.
 - **Annotations often disagree** with user preferences.
- User interactions solve this problem but bring **noise and biases**.
- **Counter-factual approaches** allow for **unbiased** learning to rank:
 - If an **accurate user model** can be learned, we can **adjust for biases**.
 - **Only** require randomization to **infer a user model**.
- **Online approaches** allow for **unbiased** and **responsive** learning to rank:
 - **Immediately adapt** to user behaviour.
 - Perform **randomization** at each step, though limited.
- Clear, that different situations suit different approaches.

Conclusion

Take-away messages:

- **Supervised approaches** to learning to rank are **limited**.
 - **Annotations often disagree** with user preferences.
- User interactions solve this problem but bring **noise and biases**.
- **Counter-factual approaches** allow for **unbiased** learning to rank:
 - If an **accurate user model** can be learned, we can **adjust for biases**.
 - **Only** require randomization to **infer a user model**.
- **Online approaches** allow for **unbiased** and **responsive** learning to rank:
 - **Immediately adapt** to user behaviour.
 - Perform **randomization** at each step, though limited.
- Clear, that different situations suit different approaches.
- Unclear, when which approach is preferred (still missing a direct comparison).

- O. Chapelle and Y. Chang. Yahoo! Learning to Rank Challenge Overview. *Journal of Machine Learning Research*, 14:1–24, 2011.
- A. Chuklin, I. Markov, and M. d. Rijke. Click models for web search. *Synthesis Lectures on Information Concepts, Retrieval, and Services*, 7(3):1–115, 2015.
- T. Joachims. Optimizing search engines using clickthrough data. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 133–142. ACM, 2002.
- T. Joachims, A. Swaminathan, and T. Schnabel. Unbiased learning-to-rank with biased feedback. In *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, pages 781–789. ACM, 2017.
- D. Lefortier, P. Serdyukov, and M. de Rijke. Online exploration for detecting shifts in fresh intent. In *CIKM 2014: 23rd ACM Conference on Information and Knowledge Management*. ACM, November 2014.

- H. Oosterhuis and M. de Rijke. Sensitive and scalable online evaluation with theoretical guarantees. In *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management*, pages 77–86. ACM, 2017.
- H. Oosterhuis and M. de Rijke. Differentiable unbiased online learning to rank. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 1293–1302. ACM, 2018.
- T. Qin and T.-Y. Liu. Introducing letor 4.0 datasets. *arXiv preprint arXiv:1306.2597*, 2013.
- M. Sanderson. Test collection based evaluation of information retrieval systems. *Foundations and Trends in Information Retrieval*, 4(4):247–375, 2010.
- X. Wang, M. Bendersky, D. Metzler, and M. Najork. Learning to rank with selection bias in personal search. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, pages 115–124. ACM, 2016.
- Y. Yue and T. Joachims. Interactively optimizing information retrieval systems as a dueling bandits problem. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 1201–1208. ACM, 2009.

Acknowledgments



All content represents the opinion of the author(s), which is not necessarily shared or endorsed by their employers and/or sponsors.