

Optimizing Search and Recommender Systems based on Position-Biased User Interactions

Harrie Oosterhuis

April 30, 2021

Radboud University, Nijmegen

harrie.oosterhuis@ru.nl

Based on the WWW'20 tutorial: *Unbiased Learning to Rank: Counterfactual and Online Approaches*
(Harrie Oosterhuis, Rolf Jagerman, and Maarten de Rijke).

Who are we?



Harrie Oosterhuis
Assistant Professor at
Radboud University



Rolf Jagerman
Software Engineer at
Google Research



Maarten de Rijke
Professor at
University of Amsterdam

Introduction

Example of Ranking for Search

[All](#) [Images](#) [Videos](#) [News](#) [Maps](#) | [Meanings](#) [Settings](#) ▾

Netherlands ▾ [Safe search: moderate](#) ▾ [Any time](#) ▾

Florence Nightingale - Wikipedia

https://en.wikipedia.org/wiki/Florence_Nightingale

Florence Nightingale, OM, RRC, DStJ (/ˈnɑːtɪnɡɛɪl/; 12 May 1820 - 13 August 1910) was an English social reformer, statistician and the founder of modern nursing. **Nightingale** came to prominence while serving as a manager and trainer of nurses during the Crimean War, in which she organised care for wounded soldiers at Constantinople.

Florence Nightingale - Biography, Facts & Nursing - HISTORY

<https://www.history.com/topics/womens-history/florence-nightingale-1>

Florence Nightingale (1820-1910), known as "The Lady With the Lamp," was a British nurse, social reformer and statistician best known as the founder of modern nursing. Her experiences as a nurse...

Florence Nightingale | Biography & Facts | Britannica

<https://www.britannica.com/biography/Florence-Nightingale>

Florence Nightingale, byname Lady with the Lamp, (born May 12, 1820, Florence [Italy]—died August 13, 1910, London, England), British nurse, statistician, and social reformer who was the foundational philosopher of modern nursing. **Nightingale** was put in charge of nursing British and allied soldiers in Turkey during the Crimean War.

Florence Nightingale | National Women's History Museum

<https://www.womenshistory.org/education-resources/biographies/florence-nightingale>


Often called "the Lady with the Lamp," **Florence Nightingale** was a caring nurse and a leader. In addition to writing over 150 books, pamphlets and reports on health-related issues, she is also credited with creating one of the first versions of the pie chart. However, she is mostly known for making hospitals a cleaner and safer place to be.

Florence Nightingale - Quotes, Education & Facts - Biography

<https://www.biography.com/scientist/florence-nightingale>

Florence Nightingale was a trailblazing figure in nursing who greatly affected 19th- and 20th-century policies around proper medical care. She was known for her night rounds to aid the wounded,...

Florence Nightingale



Florence Nightingale, was an English social reformer, statistician and the founder of modern nursing. Nightingale came to prominence while serving as a manager and trainer of nurses during the Crimean War, in which she organised care for wounded soldiers at Constantinople. [Wikipedia](#)

Born: May 12, 1820, Florence, Grand Duchy of Tuscany

Died: August 13, 1910, Mayfair, London, England, UK

Nationality: British

[Share Feedback](#)

Example of Ranking for Recommendation



Unbiased Learning to Rank: Counterfactual and Online Approaches

Harrie Oosterhuis* Rolf Jagerman* Maarten de Rijke***

April 21, 2020

*University of Amsterdam
**Ahold Delhaize
oosterhuis@uva.nl, rolf.jagerman@uva.nl, derijke@uva.nl

The Web Conference 2020 Tutorial

UNIVERSITY OF AMSTERDAM
Unbiased Learning to Rank: Counterfactual and Online Approaches - The Web Conference 2020 Tutorial
1,569 views · Apr 22, 2020

69 1 SHARE SAVE ...

- Taking the Counterfactual Online: Efficient and Unbiased...**
Harrie Oosterhuis
88 views · 7 months ago
20:01
- Michael Manapat: Counterfactual evaluation of...**
PyData
9.1K views · 5 years ago
44:32
- Neural Learning To Rank by Bhaskar Mitra (Guest lecture a...**
Bhaskar Mitra
633 views · 1 year ago
1:37:01
- TF-Ranking: Learning-to-Rank in Tensorflow - Michael Bendersky**
BCS Member Groups
63 views · 1 month ago
27:39
- DirectStorage for Windows**
Microsoft Game Stack
3.8K views · 1 week ago
29:54
- Learning to Rank from Clicks - Zach Alexander & Tracy Backe...**
Lucidworks
2.6K views · 3 years ago
30:11
- Counterfactual Evaluation and Learning from Logged User...**
Microsoft Research
2.6K views · 4 years ago
1:16:06
- Deep Learning from Logged Interventions - Thorsten...**
Criteo Eng
1.7K views · 3 years ago
45:31
- Leon Bottou Microsoft - Counterfactual Baseline**

Learning to Rank (LTR) is:

“... the task to automatically construct a ranking model using training data, such that the model can sort new objects according to their degrees of relevance, preference, or importance.”

— Liu et al. (2009)

Learning to Rank (LTR) is:

“... the task to automatically construct a ranking model using training data, such that the model can sort new objects according to their degrees of relevance, preference, or importance.”

— Liu et al. (2009)

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

- Key component for **search** and **recommendation**.

Learning to Rank: Problem Definition

The **ranking** R of **ranker** f_θ over a document set D is:

$$R = (R_1, R_2, R_3, \dots),$$

where documents are **ordered** by their (descending) **scores**:

$$f_\theta(R_1) \geq f_\theta(R_2) \geq f_\theta(R_3) \geq \dots,$$

and **every document** is in the ranking:

$$d \in D \iff d \in R.$$

For this talk, we will cast **the goal of LTR** as:

- Find the **parameters** θ for the **model** f_θ ,
where sorting documents d according to their scores $f_\theta(d)$
results in the **most optimal rankings**.

We will later define what is *optimal* according to well-known ranking metrics.

Limitations of Annotated Datasets

Traditionally, learning to rank is **supervised** through **annotated datasets**:

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

However, over time **several limitations** of this approach have become apparent.

Limitations of the Annotated Datasets

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

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

Limitations of the Annotated Datasets

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).

Limitations of the Annotated Datasets

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**.

Limitations of the Annotated Datasets

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).

Limitations of the Annotated Datasets

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.

Limitations of the Supervised Approach

Annotated datasets are **valuable** and have an **important place in research and development**.

However, the supervised approach is:

- **Unavailable** for practitioners without a **considerable budget**.
- **Impossible** for certain ranking problems.
- Often **misaligned** with *true* user preferences.

Therefore, there is a **need** for an **alternative** learning to rank approach.

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 **behavior** 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 **behavior** is indicative of their **preferences**.

User interactions also bring their **own difficulties**:

- 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 differently**.
 - ...

The Golden Triangle

The image shows a screenshot of a Google search results page for the query "digital camera cheapest". The search results are displayed in a grid format, with a prominent "Golden Triangle" of results highlighted in yellow. The search results include:

- Technical Camera**: Cheap digital camera review cheap ...
- Vivitar Digital Camera**: UK digital camera ...
- Yakumo Digital Camera**: Cheapest digital cameras ...
- Best Deals on Digital Cameras and Accessories**: Personal retail service company with discount prices on all photographic & digital camera equipment. Our prices are among the cheapest that you will find on ...
- Cheapest Digital Cameras and the Canon Full S3**
- Cheapest colour LCD digital camera**: ebchina.com
- Cheapest digital camera**: Product Description, Cheapest digital camera
- ebchina.com**: vbcps4.mtd1243997.htm
- ebchina.com**: Cheap computer and home digital cameras deals
- Deals List**: Canon PowerShot SD110 digital camera - Canon PIXMA iP3000 Photo
- Deals List**: Canon PowerShot 5400 5.1 MP Digital Camera w/ 4x Optical Zoom
- Cheapest digital camera**: InfoLour search

The search results are displayed in a grid format, with a prominent "Golden Triangle" of search results highlighted in yellow. The search results include:

- Technical Camera**: Cheap digital camera review cheap ...
- Vivitar Digital Camera**: UK digital camera ...
- Yakumo Digital Camera**: Cheapest digital cameras ...
- Best Deals on Digital Cameras and Accessories**: Personal retail service company with discount prices on all photographic & digital camera equipment. Our prices are among the cheapest that you will find on ...
- Cheapest Digital Cameras and the Canon Full S3**
- Cheapest colour LCD digital camera**: ebchina.com
- Cheapest digital camera**: Product Description, Cheapest digital camera
- ebchina.com**: vbcps4.mtd1243997.htm
- ebchina.com**: Cheap computer and home digital cameras deals
- Deals List**: Canon PowerShot SD110 digital camera - Canon PIXMA iP3000 Photo
- Deals List**: Canon PowerShot 5400 5.1 MP Digital Camera w/ 4x Optical Zoom
- Cheapest digital camera**: InfoLour search

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.

Counterfactual Learning to Rank

The remainder of this talk will cover the following topics:

- **Counterfactual Evaluation**
 - Evaluating unbiasedly from historical interactions.
- **Propensity-weighted LTR**
 - Learning unbiasedly from historical interactions.
- **Estimating Position Bias**
- **Practical Considerations**

Counterfactual Evaluation

Counterfactual Evaluation: Introduction

Evaluation is incredibly **important before deploying** a ranking system.

However, with the **limitations of annotated datasets**,
can we **evaluate** a ranker **without deploying** it or **annotated data**?

Counterfactual Evaluation: Introduction

Evaluation is incredibly **important before deploying** a ranking system.

However, with the **limitations of annotated datasets**,
can we **evaluate** a ranker **without deploying** it or **annotated data**?

Counterfactual Evaluation:

Evaluate a new ranking function f_θ using **historical interaction data** (e.g., clicks) collected from a previously deployed ranking function f_{deploy} .

Counterfactual Evaluation: Full Information

If we **know** the **true relevance labels** ($y(d_i)$ for all i), we can compute any additive linearly decomposable IR metric as:

$$\Delta(f_\theta, D, y) = \sum_{d_i \in D} \lambda(\text{rank}(d_i \mid f_\theta, D)) \cdot y(d_i),$$

where λ is a rank weighting function, e.g.,

Average Relevant Position	$ARP : \lambda(r) = r,$
Discounted Cumulative Gain	$DCG : \lambda(r) = \frac{1}{\log_2(1 + r)},$
Precision at k	$Prec@k : \lambda(r) = \frac{\mathbf{1}[r \leq k]}{k}.$

Counterfactual Evaluation: Full Information

$$y(d_1) = 1$$

Document d_1

$$y(d_2) = 0$$

Document d_2

$$y(d_3) = 0$$

Document d_3

$$y(d_4) = 1$$

Document d_4

$$y(d_5) = 0$$

Document d_5

We often do not know the true relevance labels $y(d_i)$, but can only observe implicit feedback in the form of, e.g., clicks:

- A click c_i on document d_i is a **biased and noisy indicator** that d_i is relevant
- A missing click does **not** necessarily indicate non-relevance.

Counterfactual Evaluation: Clicks

$$y(d_1) = 1$$

Document d_1

$$y(d_2) = 0$$

Document d_2

$$y(d_3) = 0$$

Document d_3

$$y(d_4) = 1$$

Document d_4

$$y(d_5) = 0$$

Document d_5

Counterfactual Evaluation: Clicks

$$y(d_1) = 1$$

Document d_1



$$y(d_2) = 0$$

Document d_2

$$y(d_3) = 0$$

Document d_3

$$y(d_4) = 1$$

Document d_4

$$y(d_5) = 0$$

Document d_5

Counterfactual Evaluation: Clicks

$$y(d_1) = 1$$

Document d_1



$$c_1 = 1$$

$$y(d_2) = 0$$

Document d_2

$$y(d_3) = 0$$

Document d_3

$$y(d_4) = 1$$


Document d_4

$$y(d_5) = 0$$

Document d_5

Counterfactual Evaluation: Clicks

$$y(d_1) = 1$$

Document d_1 



$$c_1 = 1$$

$$y(d_2) = 0$$

Document d_2



$$y(d_3) = 0$$

Document d_3

$$y(d_4) = 1$$

Document d_4

$$y(d_5) = 0$$

Document d_5

Counterfactual Evaluation: Clicks

$$y(d_1) = 1$$

Document d_1



$$c_1 = 1$$

$$y(d_2) = 0$$

Document d_2



$$c_2 = 0$$

$$y(d_3) = 0$$

Document d_3

$$y(d_4) = 1$$

Document d_4

$$y(d_5) = 0$$

Document d_5

Counterfactual Evaluation: Clicks

$$y(d_1) = 1$$

Document d_1



$$c_1 = 1$$

$$y(d_2) = 0$$

Document d_2



$$c_2 = 0$$

$$y(d_3) = 0$$

Document d_3



$$y(d_4) = 1$$


Document d_4

$$y(d_5) = 0$$

Document d_5

Counterfactual Evaluation: Clicks

$$y(d_1) = 1$$

Document d_1 



$$c_1 = 1$$


$$y(d_2) = 0$$

Document d_2



$$c_2 = 0$$

$$y(d_3) = 0$$

Document d_3 



$$c_3 = 1$$

$$y(d_4) = 1$$


Document d_4

$$y(d_5) = 0$$

Document d_5

Counterfactual Evaluation: Clicks

$$y(d_1) = 1$$

Document d_1 



$$c_1 = 1$$


$$y(d_2) = 0$$

Document d_2



$$c_2 = 0$$

$$y(d_3) = 0$$

Document d_3 



$$c_3 = 1$$

$$y(d_4) = 1$$

Document d_4




$$y(d_5) = 0$$

Document d_5

Counterfactual Evaluation: Clicks

$$y(d_1) = 1$$

Document d_1 



$$c_1 = 1$$


$$y(d_2) = 0$$

Document d_2



$$c_2 = 0$$

$$y(d_3) = 0$$

Document d_3 



$$c_3 = 1$$

$$y(d_4) = 1$$

Document d_4




$$c_4 = 0$$

$$y(d_5) = 0$$

Document d_5

Counterfactual Evaluation: Clicks

$$y(d_1) = 1$$

Document d_1 



$$c_1 = 1$$


$$y(d_2) = 0$$

Document d_2



$$c_2 = 0$$

$$y(d_3) = 0$$

Document d_3 



$$c_3 = 1$$

$$y(d_4) = 1$$

Document d_4



$$c_4 = 0$$


$$y(d_5) = 0$$

Document d_5



Counterfactual Evaluation: Clicks

$$y(d_1) = 1$$

Document d_1 



$$c_1 = 1$$


$$y(d_2) = 0$$

Document d_2



$$c_2 = 0$$

$$y(d_3) = 0$$

Document d_3 



$$c_3 = 1$$

$$y(d_4) = 1$$

Document d_4



$$c_4 = 0$$

$$y(d_5) = 0$$

Document d_5



$$c_5 = 0$$

Remember that there are many reasons why a click on a document may **not** occur:

- **Relevance**: the document may not be relevant.
- **Observance**: the user may not have examined the document.
- **Miscellaneous**: various random reasons why a user may not click.

Remember that there are many reasons why a click on a document may **not** occur:

- **Relevance**: the document may not be relevant.
- **Observance**: the user may not have examined the document.
- **Miscellaneous**: various random reasons why a user may not click.

Some of these reasons are considered to be:

- **Noise**: averaging over many clicks will remove their effect.
- **Bias**: averaging will **not** remove their effect.

Counterfactual Evaluation: Examination User Model

If we **only** consider **examination** and **relevance**, a user click can be modelled by:

- The probability of document d_i **being examined** ($o_i = 1$) in a ranking R :

$$P(o_i = 1 \mid R, d_i).$$

Counterfactual Evaluation: Examination User Model

If we **only** consider **examination** and **relevance**, a user click can be modelled by:

- The probability of document d_i **being examined** ($o_i = 1$) in a ranking R :

$$P(o_i = 1 \mid R, d_i).$$

- The probability of a **click** $c_i = 1$ on d_i given its **relevance** $y(d_i)$ and whether it was **examined** o_i :

$$P(c_i = 1 \mid o_i, y(d_i)).$$

Counterfactual Evaluation: Examination User Model

If we **only** consider **examination** and **relevance**, a user click can be modelled by:

- The probability of document d_i **being examined** ($o_i = 1$) in a ranking R :

$$P(o_i = 1 \mid R, d_i).$$

- The probability of a **click** $c_i = 1$ on d_i given its **relevance** $y(d_i)$ and whether it was **examined** o_i :

$$P(c_i = 1 \mid o_i, y(d_i)).$$

- **Clicks only occur on examined documents**, thus the probability of a click in ranking R is:

$$P(c_i = 1 \wedge o_i = 1 \mid y(d_i), R) = P(c_i = 1 \mid o_i = 1, y(d_i)) \cdot P(o_i = 1 \mid R, d_i).$$

Counterfactual Evaluation: Naive Estimator

A **naive way** to estimate is to assume clicks are a unbiased relevance signal:

$$\hat{\Delta}_{NAIVE}(f_{\theta}, D, c) = \sum_{d_i \in D} \lambda(\text{rank}(d_i | f_{\theta}, D)) \cdot c_i.$$

Counterfactual Evaluation: Naive Estimator

A **naive way** to estimate is to assume clicks are a unbiased relevance signal:

$$\hat{\Delta}_{NAIVE}(f_{\theta}, D, c) = \sum_{d_i \in D} \lambda(\text{rank}(d_i | f_{\theta}, D)) \cdot c_i.$$

Even if **no click noise** is present: $P(c_i = 1 | o_i = 1, y(d_i)) = y(d_i)$, this estimator is **biased** by the observation probabilities:

$$\begin{aligned} \mathbb{E}_o[\hat{\Delta}_{NAIVE}(f_{\theta}, D, c)] &= \mathbb{E}_o \left[\sum_{d_i: o_i=1 \wedge y(d_i)=1} \lambda(\text{rank}(d_i | f_{\theta}, D)) \right] \\ &= \sum_{d_i: y(d_i)=1} P(o_i = 1 | R, d_i) \cdot \lambda(\text{rank}(d_i | f_{\theta}, D)). \end{aligned}$$

Counterfactual Evaluation: Naive Estimator Bias

The biased estimator **weights documents** according to their **observation probabilities** in the ranking R displayed during **logging**:

$$\mathbb{E}_o[\hat{\Delta}_{NAIVE}(f_\theta, D, c)] = \sum_{d_i: y(d_i)=1} P(o_i = 1 \mid R, d_i) \cdot \lambda(\text{rank}(d_i \mid f_\theta, D)).$$

In rankings, **documents at higher ranks** are more likely to be examined: **position bias**.

Counterfactual Evaluation: Naive Estimator Bias

The biased estimator **weights documents** according to their **observation probabilities** in the ranking R displayed during **logging**:

$$\mathbb{E}_o[\hat{\Delta}_{NAIVE}(f_\theta, D, c)] = \sum_{d_i: y(d_i)=1} P(o_i = 1 \mid R, d_i) \cdot \lambda(\text{rank}(d_i \mid f_\theta, D)).$$

In rankings, **documents at higher ranks** are more likely to be examined: **position bias**.

Position bias causes **logging-policy-confirming** behavior:

- Documents displayed at **higher ranks during logging** are incorrectly considered as **more relevant**.

Inverse Propensity Scoring

Inverse Propensity Scoring (IPS) estimators can remove bias:

- First introduced by Wang et al. (2016) and Joachims et al. (2017).
- **Main idea:** weight clicks depending on their *observation probability*

Counterfactual Evaluation: Inverse Propensity Scoring

Inverse Propensity Scoring (IPS) estimators can remove bias:

- First introduced by Wang et al. (2016) and Joachims et al. (2017).
- **Main idea**: weight clicks depending on their *observation probability*
- Clicks near the **top** of the ranked list:
 - Have **high** observation probability \Leftrightarrow Get assigned **small** weight
- Clicks near the **bottom** of the ranked list:
 - Have **low** observation probability \Leftrightarrow Get assigned **large** weight

Counterfactual Evaluation: Inverse Propensity Scoring

Counterfactual evaluation accounts for bias using **Inverse Propensity Scoring (IPS)**:

$$\hat{\Delta}_{IPS}(f_{\theta}, D, c) = \sum_{d_i \in D} \frac{\lambda(\text{rank}(d_i | f_{\theta}, D))}{P(o_i = 1 | R, d_i)} \cdot c_i,$$

where

- $\lambda(\text{rank}(d_i | f_{\theta}, D))$: (weighted) rank of document d_i by ranker f_{θ} ,
- c_i : observed click on the document in the log,
- $P(o_i = 1 | R, d_i)$: observation probability of d_i in ranking R displayed during logging.

This is an **unbiased estimate** of any additive linearly decomposable IR metric.

Counterfactual Evaluation: Proof of Unbiasedness

If no click noise is present, this provides an **unbiased estimate**:

$$\begin{aligned}\mathbb{E}_o[\hat{\Delta}_{IPS}(f_\theta, D, c)] &= \mathbb{E}_o \left[\sum_{d_i \in D} \frac{\lambda(\text{rank}(d_i | f_\theta, D))}{P(o_i = 1 | R, d_i)} \cdot c_i \right] \\ &= \mathbb{E}_o \left[\sum_{d_i: o_i=1 \wedge y(d_i)=1} \frac{\lambda(\text{rank}(d_i | f_\theta, D))}{P(o_i = 1 | R, d_i)} \right] \\ &= \sum_{d_i: y(d_i)=1} \frac{P(o_i = 1 | R, d_i) \cdot \lambda(\text{rank}(d_i | f_\theta, D))}{P(o_i = 1 | R, d_i)} \\ &= \sum_{d_i \in D} \lambda(\text{rank}(d_i | f_\theta, D)) \cdot y(d_i) \\ &= \Delta(f_\theta, D, y).\end{aligned}$$

Counterfactual Evaluation: Robustness of Noise

So far we have **no click noise**: $P(c_i = 1 \mid o_i = 1, y(d_i)) = y(d_i)$.

However, the IPS approach still works without these assumptions, as long as:

$$y(d_i) > y(d_j) \Leftrightarrow P(c_i = 1 \mid o_i = 1, y(d_i)) > P(c_j = 1 \mid o_j = 1, y(d_j)).$$

Since we can prove **relative differences** are inferred unbiasedly:

$$\mathbb{E}_{o,c}[\hat{\Delta}_{IPS}(f_\theta, D, c)] > \mathbb{E}_{o,c}[\hat{\Delta}_{IPS}(f_{\theta'}, D, c)] \Leftrightarrow \Delta(f_\theta, D) > \Delta(f_{\theta'}, D).$$

Propensity-weighted Learning to Rank

Propensity-weighted Learning to Rank (LTR)

The inverse-propensity-scored estimator can unbiasedly estimate performance:

$$\hat{\Delta}_{IPS}(f_{\theta}, D, c) = \sum_{d_i \in D} \frac{\lambda(\text{rank}(d_i | f_{\theta}, D))}{P(o_i = 1 | R, d_i)} \cdot c_i.$$

How do we **optimize** for this **unbiased performance estimate**?

- It is **not differentiable**.
- **Common problem for all ranking metrics**.

The inverse-propensity-scored estimator can unbiasedly estimate performance:

$$\hat{\Delta}_{IPS}(f_{\theta}, D, c) = \sum_{d_i \in D} \frac{\lambda(\text{rank}(d_i | f_{\theta}, D))}{P(o_i = 1 | R, d_i)} \cdot c_i.$$

How do we **optimize** for this **unbiased performance estimate**?

- It is **not differentiable**.
- **Common problem for all ranking metrics**.

Solution:

- Optimize a **bound** on the metric instead (out of the scope of this talk).

Propensity-weighted LTR: Walkthrough

Overview of the approach:

- Obtain a **model of position bias**.
- Acquire a **large click-log**.
- Then for every click in the log:
 - Compute the **propensity of the click**:

$$P(o_i = 1 \mid R, d_i).$$

- Calculate the **gradient** of the **bound** on the **unbiased estimator**:

$$\nabla_{\theta} \left[\frac{\lambda(\overline{\text{rank}}(d_i \mid f_{\theta}, D))}{P(o_i = 1 \mid R, d_i)} \right].$$

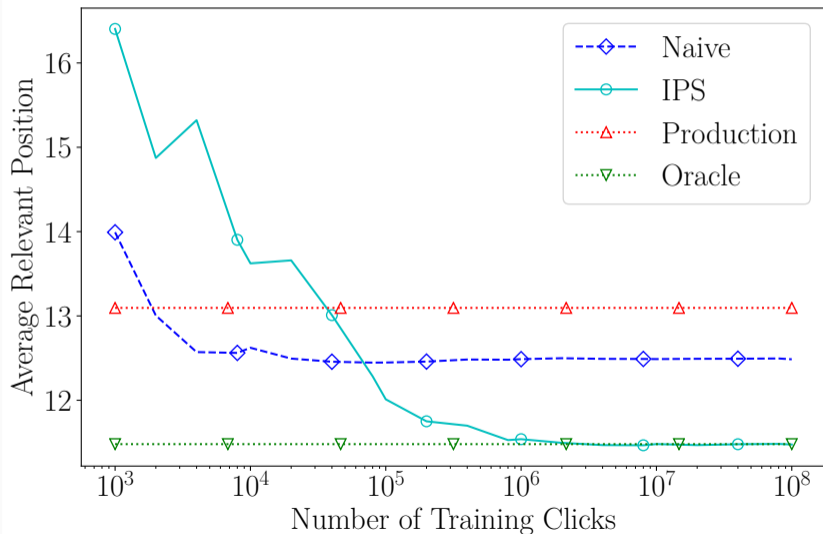
- **Update the model** f_{θ} by adding/subtracting the gradient.

Unbiased LTR methods are commonly **evaluated** through **semi-synthetic experiments** (Joachims, 2002; Agarwal et al., 2019a; Jagerman et al., 2019).

The experimental setup:

- Traditional LTR dataset, e.g., Yahoo! Webscope (Chapelle and Chang, 2011).
- Create rankings according to a baseline ranking model.
- Simulate clicks by modelling:
 - **Click Noise**, e.g., 10% chance of clicking on a non-relevant document.
 - **Position Bias**, e.g., $P(o_i = 1 \mid R, d_i) = \frac{1}{\text{rank}(d \mid R)}$.
- Hyper-parameter tuning by unbiased evaluation methods.

Propensity-weighted LTR: Results



Estimating Position Bias

So far we have seen how to:

- Perform **Counterfactual Evaluation** with **unbiased estimators**.
- Perform **Counterfactual LTR** by optimizing **unbiased estimators**.

At the core of these methods is the propensity score: $P(o_i = 1 \mid R, d_i)$, which helps to remove bias from user interactions.

In this section, we will show how this **propensity score** can be **estimated** for **position bias**.

Recall that position bias is a form of bias where higher positioned results are more likely to be observed and therefore clicked.

Assumption: The **observation probability** only depends on the rank of a document:


$$P(o_i = 1 \mid i).$$

The objective is now to **estimate**, for each rank i , the propensity $P(o_i = 1 \mid i)$.

This user model was first formalized by Craswell et al. (2008).

Estimating Position Bias

RandTop- n Algorithm:



Estimating Position Bias

RandTop- n Algorithm:

Document d_1	Document d_3
Document d_2	Document d_4
Document d_3	Document d_1
Document d_4	Document d_2

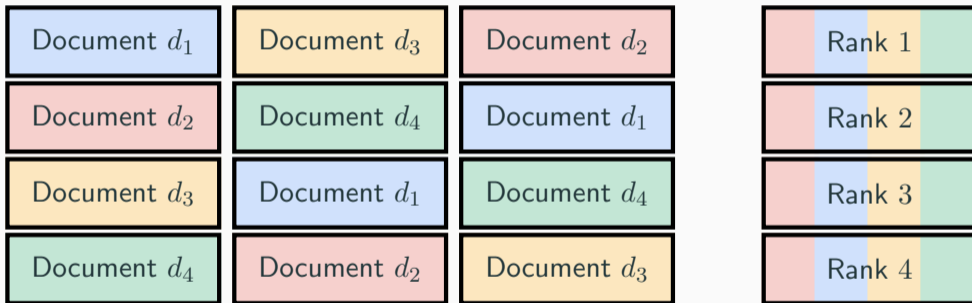
Estimating Position Bias

RandTop- n Algorithm:

Document d_1	Document d_3	Document d_2
Document d_2	Document d_4	Document d_1
Document d_3	Document d_1	Document d_4
Document d_4	Document d_2	Document d_3

Estimating Position Bias

RandTop- n Algorithm:



RandTop- n Algorithm:

- ① Repeat:
 - Randomly shuffle the top n items
 - Record clicks
- ② Aggregate clicks per rank
- ③ Normalize to obtain propensities $p_i \propto P(o_i | i)$


Note: we only need propensities proportional to the true observation probability for learning.

Uniformly **randomizing** the top n results may negatively impacts users during data logging.

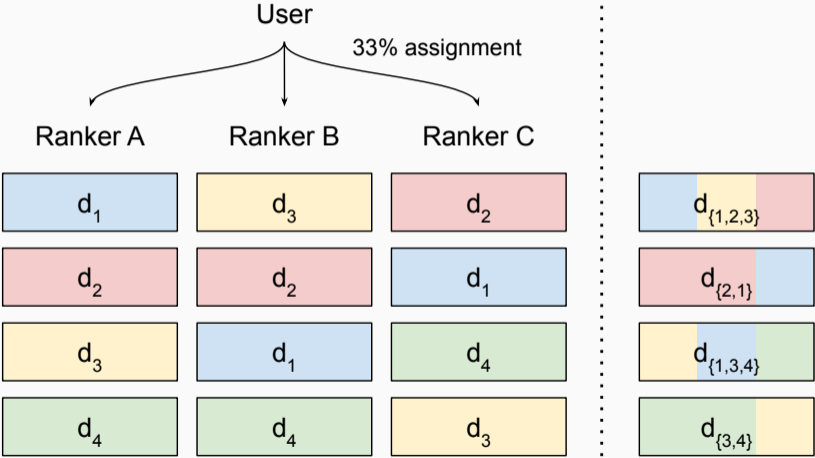
There are various methods that minimize the impact to the user:

- **RandPair:** Choose a pivot rank k and only swap a random other document with the document at this pivot rank (Joachims et al., 2017).
- **Interventional Sets:** Exploit inherent “randomness” in data coming from multiple rankers (e.g., A/B tests in production logs) (Agarwal et al., 2017).

Intervention Harvesting



Intervention Harvesting



Jointly Learning and Estimating

In the previous sections we have seen:

- Counterfactual ranker evaluation with unbiased estimators.
- Counterfactual LTR by optimizing unbiased estimators.
- Estimating propensity scores through randomization.

Instead of treating **propensity estimation** and **unbiased learning to rank** as two separate tasks, recent work has explored **jointly learning rankings and estimating propensities**.

Recall that the probability of a click can be decomposed as:

$$\underbrace{P(c_i = 1 \wedge o_i = 1 \mid y(d_i), R)}_{\text{click probability}} = \underbrace{P(c_i = 1 \mid o_i = 1, y(d_i))}_{\text{relevance probability}} \cdot \underbrace{P(o_i \mid R, d_i)}_{\text{observation probability}}.$$

In the previous sections we have seen that, if the **observation probability** is known, we can find an unbiased estimate of relevance via IPS.

It is possible to **jointly learn and estimate** by iterating two steps:

- 1 Learn an optimal ranker given a correct propensity model:

$$\underbrace{P(c_i = 1 \mid o_i = 1, y(d_i))}_{\text{relevance probability}} = \frac{P(c_i = 1 \wedge o_i = 1 \mid y(d_i), R)}{P(o_i = 1 \mid R, d_i)}.$$

- 2 Learn an optimal propensity model given a correct ranker:

$$\underbrace{P(o_i = 1 \mid R, d_i)}_{\text{observation probability}} = \frac{P(c_i = 1 \wedge o_i = 1 \mid y(d_i), R)}{P(c_i = 1 \mid o_i = 1, y(d_i))}.$$

Given an accurate **model of relevance**, it is possible to find an accurate **propensity model**, and vice versa.

This approach requires **no randomization**.

Recent work has solved this via either an **Expectation-Maximization approach** (Wang et al. (2018a)) or a **Dual Learning Objective** (Ai et al. (2018)).

Practical Considerations

Practitioners of counterfactual LTR systems will run into the problem of **high variance**.

High variance can be due to many factors:

- Not enough training data.
- Extreme position bias and very small propensities.
- Large amounts of noisy clicks on documents with small propensities.

The usual suspect is **one or a few data points** with extremely small propensity that **overpower** the rest of the data set.

A typical solution to **high variance** is to apply **propensity clipping**.


Propensity clipping: Bound the *propensity*, to prevent any single sample from overpowering the rest of the data set:

$$\hat{\Delta}_{Clipped-IPS}(f_{\theta}, D, c) = \sum_{d_i \in D} \frac{\lambda(\text{rank}(d_i | f_{\theta}, D))}{\max\{\tau, P(o_i = 1 | R, d_i)\}} \cdot c_i.$$

This solution trades off bias for variance: it will introduce some amount of bias but can substantially reduce variance.

Note that when $\tau = 1$, we obtain the biased naive estimator.

Practical Considerations



Comparison to Supervised LTR

Supervised LTR:

- Uses **manually annotated labels**:
 - expensive to create,
 - impossible in many settings,
 - often misaligned with actual user preferences.
- Optimization is widely studied and very effective w.r.t. evaluation on annotated labels.
- Often unavailable for practitioners.

Counterfactual LTR:

- Uses **click logs**:
 - available in abundant quantities,
 - effectively no cost,
 - contains **noise** and **biases**.
- **Noise**: amortized over large numbers of clicks.
- **Biases**:
 - position bias mitigated with inverse propensity scoring.
 - other biases are an active area of research.

Conclusion

Today we discussed:

- **User interactions** with rankings are **very biased**.
- **Counterfactual Learning to Rank:**
 - Correct for position bias with inverse propensity scoring.
 - Requires an explicit user model.
- Unbiased learning from **historical** interaction logs.

The unbiased learning to rank field is very active:

- *Addressing Trust Bias for Unbiased Learning-to-Rank* (Agarwal et al., 2019b).

Small Selection of Recent Work in the Field

The unbiased learning to rank field is very active:

- *Addressing Trust Bias for Unbiased Learning-to-Rank* (Agarwal et al., 2019b).
- *Fair Learning-to-Rank from Implicit Feedback* (Yadav et al., 2019).

The unbiased learning to rank field is very active:

- *Addressing Trust Bias for Unbiased Learning-to-Rank* (Agarwal et al., 2019b).
- *Fair Learning-to-Rank from Implicit Feedback* (Yadav et al., 2019).
- *Correcting for Selection Bias in Learning-to-rank Systems* (Ovaisi et al., 2020).
- *Policy-Aware Unbiased Learning to Rank for Top-k Rankings* (Oosterhuis and de Rijke, 2020).

The unbiased learning to rank field is very active:

- *Addressing Trust Bias for Unbiased Learning-to-Rank* (Agarwal et al., 2019b).
- *Fair Learning-to-Rank from Implicit Feedback* (Yadav et al., 2019).
- *Correcting for Selection Bias in Learning-to-rank Systems* (Ovaisi et al., 2020).
- *Policy-Aware Unbiased Learning to Rank for Top-k Rankings* (Oosterhuis and de Rijke, 2020).
- *When Inverse Propensity Scoring does not Work: Affine Corrections for Unbiased Learning to Rank* (Vardasbi et al., 2020).

Thank you for your attention!

Notation

Definition	Notation	Example
Query	q	–
Candidate documents	D	–
Document	$d \in D$	–
Ranking	R	(R_1, R_2, \dots, R_n)
Document at rank i	R_i	$R_i = d$
Relevance	$y : D \rightarrow \mathbb{N}$	$y(d) = 2$
Ranker model with weights θ	$f_\theta : D \rightarrow \mathbb{R}$	$f_\theta(d) = 0.75$
Click	$c_i \in \{0, 1\}$	–
Observation	$o_i \in \{0, 1\}$	–
Rank of d when f_θ ranks D	$\text{rank}(d \mid f_\theta, D)$	$\text{rank}(d \mid f_\theta, D) = 4$

Differentiable upper bound on $\text{rank}(d, f_\theta, D)$	$\overline{\text{rank}}(d, f_\theta, D)$	–
Average Relevant Position metric	ARP	–
Discounted Cumulative Gain metric	DCG	–
Precision at k metric	$Prec@k$	–
A performance measure or estimator	Δ	–

- Tensorflow Learning to Rank, allows for inverse propensity scoring:
<https://github.com/tensorflow/ranking>
- Inverse Propensity Score Rank-SVM:
https://www.cs.cornell.edu/people/tj/svm_light/svm_proprank.html
- Pairwise Differentiable Gradient Descent and Multileave Gradient Descent:
<https://github.com/Harrie0/OnlineLearningToRank>
- Data and code for comparing counterfactual and online learning to rank
<http://github.com/rjagerman/sigir2019-user-interactions>
- An older online learning to rank framework: Lerot
<https://bitbucket.org/ilps/lerot/>

- A. Agarwal, S. Basu, T. Schnabel, and T. Joachims. Effective evaluation using logged bandit feedback from multiple loggers. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 687–696. ACM, 2017.
- A. Agarwal, K. Takatsu, I. Zaitsev, and T. Joachims. A general framework for counterfactual learning-to-rank. In *42nd International ACM SIGIR Conference on Research & Development in Information Retrieval*, page (to appear). ACM, 2019a.
- A. Agarwal, X. Wang, C. Li, M. Bendersky, and M. Najork. Addressing trust bias for unbiased learning-to-rank. In *The World Wide Web Conference*, pages 4–14. ACM, 2019b.
- A. Agarwal, I. Zaitsev, X. Wang, C. Li, M. Najork, and T. Joachims. Estimating position bias without intrusive interventions. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining*, pages 474–482, 2019c.
- Q. Ai, K. Bi, C. Luo, J. Guo, and W. B. Croft. Unbiased learning to rank with unbiased propensity estimation. In *Proceedings of the 41st International ACM SIGIR conference on Research and Development in Information Retrieval*, pages 385–394. ACM, 2018.

- C. J. Burges. From ranknet to lambdarank to lambdamart: An overview. *Learning*, 11(23-581):81, 2010.
- O. Chapelle and Y. Chang. Yahoo! Learning to Rank Challenge Overview. *Journal of Machine Learning Research*, 14:1–24, 2011.
- N. Craswell, O. Zoeter, M. Taylor, and B. Ramsey. An experimental comparison of click position-bias models. In *Proceedings of the 2008 international conference on web search and data mining*, pages 87–94, 2008.
- Z. Fang, A. Agarwal, and T. Joachims. Intervention harvesting for context-dependent examination-bias estimation. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 825–834, 2019.
- R. Jagerman and M. de Rijke. Accelerated convergence for counterfactual learning to rank. In *SIGIR*. ACM, 2020.

- R. Jagerman, H. Oosterhuis, and M. de Rijke. To model or to intervene: A comparison of counterfactual and online learning to rank from user interactions. In *42nd International ACM SIGIR Conference on Research & Development in Information Retrieval*, page (to appear). ACM, 2019.
- 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. Training linear svms in linear time. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 217–226. ACM, 2006.
- 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.

- T.-Y. Liu et al. Learning to rank for information retrieval. *Foundations and Trends® in Information Retrieval*, 3(3):225–331, 2009.
- H. Oosterhuis and M. de Rijke. Policy-aware unbiased learning to rank for top-k rankings. In *SIGIR*. ACM, 2020.
- Z. Ovaisi, R. Ahsan, Y. Zhang, K. Vasilaky, and E. Zheleva. Correcting for selection bias in learning-to-rank systems. *arXiv preprint arXiv:2001.11358*, 2020.
- 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.
- A. Vardasbi, H. Oosterhuis, and M. de Rijke. When inverse propensity scoring does not work: Affine corrections for unbiased learning to rank. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 1475–1484, 2020.

- 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.
- X. Wang, N. Golbandi, M. Bendersky, D. Metzler, and M. Najork. Position bias estimation for unbiased learning to rank in personal search. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining*, pages 610–618. ACM, 2018a.
- X. Wang, C. Li, N. Golbandi, M. Bendersky, and M. Najork. The lambda loss framework for ranking metric optimization. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 1313–1322. ACM, 2018b.
- H. Yadav, Z. Du, and T. Joachims. Fair learning-to-rank from implicit feedback. *arXiv preprint arXiv:1911.08054*, 2019.

Acknowledgments



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