Fairness in Ranking & Wrap-Up

CS4780/5780 – Machine Learning
Fall 2019

Nika Haghtalab & Thorsten Joachims
Cornell University

Reading:
None

Ranking in Online Systems

Ranking function π that ranks items for context x.
→ Learning-to-Rank

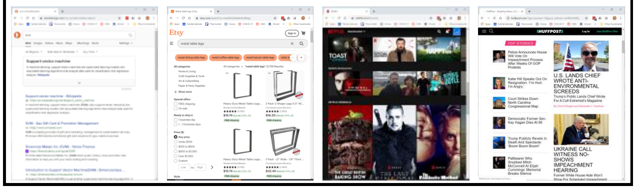

How do we train these systems?

Goal: Maximize utility of rankings to the users.

Probability Ranking Principle [Robertson, 1977]:
- Rank documents by probability of relevance → y^*
- For virtually any measure Δ of ranking quality

$$y^* \coloneqq \text{argmax}_y [\Delta(y|x)]$$

Two-Sided Market

Online Retail
- Utility to Users:
  Customers find products they want
- Utility to Items:
  Sellers get revenue

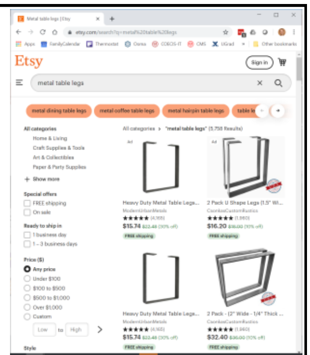

Two-Sided Market

Music Streaming
- Utility to Users:
  Customers find music they enjoy
- Utility to Items:
  Artists get streaming revenue

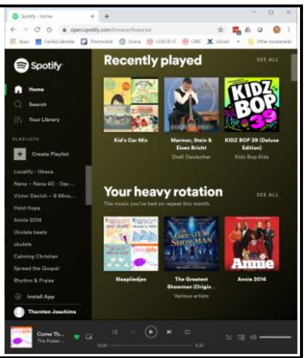

Two-Sided Market

News
- Utility to Users:
  Readers find relevant articles
- Utility to Items:
  Writers get their voice out (and ad revenue)

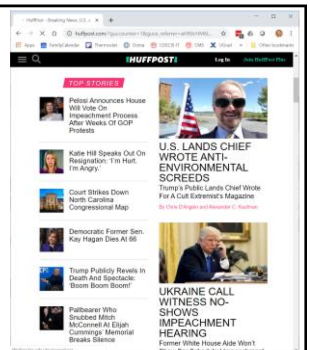

What can go wrong?

Current Learning-to-Rank methods focus only on users, and are oblivious to impact on items.

Fairness of Exposure

Fair ranking policy π allocates exposure to items based on merit.

Endogenous Factors

How to allocate exposure based on merit in order to

- Satisfy legal requirements
- Shape marketplace dynamics (e.g. Spotify, superstar economics)
- Spam, Polarization

Exogenous Factors

How to estimate merit without biases like

- Position bias
- Trust bias
- Uncertainty bias
- Historical actions
- Stereotypes

Are Conventional Methods Fair?

Probability Ranking Principle:

- Rank documents by probability of relevance → y^* [Robertson, 1977]
- For virtually any measure Δ of ranking quality
  $y^* \coloneqq \operatorname{argmax}_y [\Delta(y|x)]$
- Are rankings fair/desirable?

Top News Stories

Rank	Item	P(read)
1	Times 1	50.99
2	Times 2	50.98
3	Times 3	50.97
...	...	...
100	Post 1	49.99
101	Post 2	49.98
102	Post 3	49.97
...	...	...

Position-Based Exposure Model

Definition:

Exposure e_j is the probability a users observes the item at position j.

Rank	Exposure P(observe)
1	e_1
2	e_2
3	e_3
...	...
100	e_{100}
101	e_{101}
102	e_{102}
...	...

How to estimate?

- Eye tracking [Joachims et al. 2007]
- Intervention studies [Joachims et al. 2017]
- Intervention harvesting [Agarwal et al. 2019] [Fang et al. 2019]

Fairness Constraints

$$exposure = f(relevance)$$

- Disparate Exposure:
  – Expected exposure proportional to the expected relevance of the group
- Disparate Impact:
  – Expected revenue (e.g. clicks) proportional to the expected relevance of the group
- Group parity:
  – Expected exposure equal for all groups

Probabilistic Ranking Policies $\pi(y|x)$

Exposure and Quality for $\pi(y|x)$

$$expo(i|x) = \sum_j \mathbb{P}_{i,j}\, e_j$$

$$qual\,(\pi|x) = \sum_i \sum_j e_j\, \mathbb{P}_{i,j}\, rel_i$$

$\mathbb{P}_{i,j}$ = Prob that item i is ranked at position j

e_j = exposure at position j

π

y_1	y_2	y_3	y_4
A	B	A	B
B	A	C	C
C	C	B	A
D	D	D	G
E	E	E	F
F	F	F	E
G	G	G	D
0.52	0.23	0.20	0.05

Disparate Exposure Constraint

Group Exposure and Merit

$$expo(G|P) = \sum_{i \in G} expo\,(i|x) \qquad rel(G|P) = \sum_{i \in G} rel(i|x)$$

Group Fairness Constraint

$$\frac{expo(G_0|x)}{rel(G_0|x)} = \frac{expo(G_1|x)}{rel(G_1|x)}$$

→ Make exposure proportional to relevance

Computing the Best Fair Policy

Goal: Maximize ranking quality while fair to items.

$$\pi^*(y|x) = \operatorname{argmax}_\pi \quad [qual(\pi|x)]$$

$$s.t. \quad \frac{expo(G_0|x)}{rel(G_0|x)} = \frac{expo(G_1|x)}{rel(G_1|x)}$$

→ Computationally hard!

Marginal Rank Distribution $\mathbb{P}$

	y_1	y_2	y_3	y_4
π	A	B	A	B
	B	A	C	C
	C	D	B	A
	D	C	D	G
	E	E	E	F
	F	F	F	E
	G	G	G	D
	0.52	0.23	0.20	0.05

→

$\mathbb{P}$	1	2	3	4	5	6	7
A	0.72	0.23	0.05	0	0	0	0
B	0.28	0.52	0.20	0	0	0	0
C	…						
D				$\mathbb{P}_{i,j}$			
E							
F							
G							

Computing the Best Fair Policy

- Optimal $\mathbb{P}^*$ is solution of linear program

$$\mathbb{P}^* = \operatorname{argmax}_{\mathbb{P}} \quad [rel^T \mathbb{P} e]$$

$$s.t. \quad 1^T \mathbb{P} = 1$$

$$\mathbb{P} 1 = 1$$

$$0 \le \mathbb{P} \le 1$$

$$rel_2 g_1^T \mathbb{P} e = rel_1 g_2^T \mathbb{P} e$$

Quality

P is doubly stochastic

Fairness

Computing π^* from $\mathbb{P}^*$

Birkhoff-von Neumann decomposition

$$\mathbb{P}^* = \theta_1 P_1 + \cdots + \theta_k P_k$$

where $P_1 \ldots P_k$ are permutation matrices and $\theta_i \ge 0$ with $\sum_i \theta_i = 1$.

→ Ranking policy $\pi^*(y|x) = \begin{cases} \theta_i & if \;\; (y = P_i) \\ 0 & else \end{cases}$

Summary of Method

1. Estimate relevances r for query x
2. Define (merit-based) fairness constraint
3. Solve linear program for marginal rank matrix

$$\mathbb{P}^* = \operatorname{argmax}_{\mathbb{P}} \quad [r^T \mathbb{P} q]$$

$$s.t. \quad 1^T \mathbb{P} = 1$$

$$\mathbb{P} 1 = 1$$

$$0 \le \mathbb{P} \le 1$$

$$\mathbb{P} \; is \; fair$$

4. Compute ranking policy π^* from $\mathbb{P}^*$

Example

- Six items, two groups
- Relevances: $\mathrm{rel}(G_1) = \{82\%, 81\%, 80\%\}, \mathrm{rel}(G_2) = \{79\%, 78\%, 77\%\}$

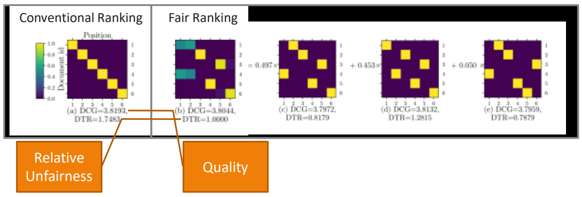

Fairness of Exposure

Fair ranking policy π allocates exposure to items based on merit.

Endogenous Factors

How to allocate exposure based on merit in order to

- Satisfy legal requirements
- Shape marketplace dynamics (e.g. Spotify, superstar economics)
- Spam, Polarization

Exogenous Factors

How to estimate merit without biases like

- Position bias
- Trust bias
- Uncertainty bias
- Historical actions
- Stereotypes

Estimating Merit from Interactions

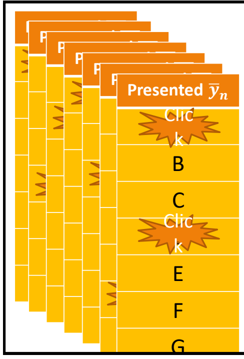

Data

- Query distribution: $x_j \sim \boldsymbol{P(X)}$
- Deployed ranker: $\bar{y}_j = \boldsymbol{\pi_0}(x_j)$
- Feedback: clicks, purchases, plays, reads

→ Feedback is biased!

Modeling Position Bias

- Assume:
  - Click implies observed and relevant:

$$(click_i = 1) \leftrightarrow (obs_i = 1) \wedge (rel_i = 1)$$

- Problem:
  - No click can mean not relevant OR not observed

$$(click_i = 0) \leftrightarrow (obs_i = 0) \vee (rel_i = 0)$$

→ Understand observation mechanism

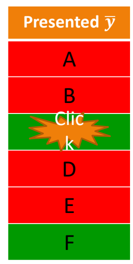

Inverse Propensity Score Estimator

- Observation Propensities
  - $Q(obs_i = 1|x, \bar{y})$
  - Random variable $obs_i \in \{0,1\}$ indicates whether relevance label rel_i is observed.
  - Can use position-based exposure $Q(obs_i = 1|x, \bar{y}) = e_i$
- De-biased Regression via IPS weighting

→ In expectation independent of past rankings!

Presented $\bar{y}$	Q
A	1.0
B	0.8
C	0.5
D	0.2
E	0.2
F	0.2

Counterfactual Policy Learning

- Policy Learning for Contextual Bandits and Ranking
  - Data is biased by past system actions
    - Propensity logging and/or propensity estimation
  - Unbiased learning objective based on causal inference
    - Inverse Propensity Score (IPS) weighting estimators
  - Directly optimize effectiveness of policy
    - Policy-gradient methods like POEM [Swaminathan & Joachims 2015], BanditNet [Joachims et al. 2018], Propensity LTR [Joachims et al. 2017]
- Transforming how industry approaches these problems
  - YouTube recommendations [Chen et al. 2019], Spotify [McInerny et al. 2018], Google Drive [Agarwal et al. 2019], …

Fairness of Exposure

Fair ranking policy π allocates exposure to items based on merit.

Endogenous Factors

How to allocate exposure based on merit in order to

- Satisfy legal requirements
- Shape marketplace dynamics (e.g. Spotify, superstar economics)
- Spam, Polarization

Exogenous Factors

How to estimate merit without biases like

- Position bias
- Trust bias
- Uncertainty bias
- Historical actions
- Stereotypes

Summary and Conclusions

- Take care of the biases in the data
  - Improve quality, solid foundation for decisions
- Shape how system serves all constituencies
  - Fairness, incentives and market dynamics, legal

→ Long term health of the system

http://www.joachims.org

WRAP-UP

Theme: Prediction and Action

- Building intelligent systems vs. analyzing existing systems
  - Prediction
  - Intelligent action
  - Guarantees on prediction/action quality

→ CS4786 Machine Learning for Data Science
→ CS4850 Math Found for the Information Age
→ INFO 6150 Advanced Topic Models

Theme: Overfitting

- Fundamental trade-off in learning
  - Training error vs. prediction error
  - Model capacity
  - Statistical learning theory
  - Empirical risk minimization

Theme: Massive Overparameterization

- The success story of machine learning
  - Regularized linear models
  - Kernels
  - Deep networks
  - → Number of parameters $\gg$ number of examples

Theme: Theoretical Underpinning

- Theory for understanding sake
  - Identify the mechanisms at play in ML
  - Understand model complexity
  - Understand common themes between algorithms

Design Approaches for ML

- Empirical Risk Minimization (ERM)
  - Fixed at training time: class of decision rules h: X → Y, loss, x and y
  - Strategy: minimize training loss
- Conditional Probability Models
  - Fixed at training time: class of models for P(Y|X), x and y
  - Strategy: max conditional likelihood or MAP (or Bayes)
- Generative Models
  - Fixed at training time: class models for P(Y,X)
  - Strategy: max likelihood or MAP (or Bayes)

- Not covered: Bayesian ML perspective → ORIE 6741

Batch Learning for Classification

- ERM
  - Decision Trees
  - Perceptron
  - SVMs
  - Neural Networks
  - Boosting
- Conditional Probability
  - Logistic Regression
  - Conditional Random Fields
  - Ridge Regression
- Generative
  - Multinomial Naïve Bayes
  - Multivariate Naïve Bayes
  - Linear Discriminant
- Other Methods
  - Gaussian Processes
  - Deep Networks
  - Recurrent Networks
  - Parametric (Graphical) Models
  - Matrix factorization
  - Many, many more …

  - *-Regression
  - *-Multiclass

Structured Output Prediction

- ERM
  - Structural SVMs
- Conditional Probability
  - Conditional Random Fields
- Generative
  - Hidden Markov Model
- Other Methods
  - Maximum Margin Markov Networks
  - Markov Random Fields
  - Bayesian Networks
  - Statistical Relational Learning
  - Markov Logic Networks
  - Encoder/Decoder Networks

→ NLP classes

Online Learning

- Expert Setting
  - Halving Algorithm
  - Weighted Majority
  - Randomized WM
- Bandit Setting
  - None
- Other Methods
  - UCB
  - EXP3
  - Follow the Leader
  - Partial Monitoring
  - Contextual Bandits
  - Dueling Bandits
  - Coactive Learning

→ CS6781 Theoretical Foundations of Machine Learning

Unsupervised Learning

- Clustering
  - None
- Other Methods
  - Spectral Clustering
  - Multi-Dimensional Scaling
  - Latent Dirichlet Allocation
  - Semantic Embeddings
  - Deep Auto-Encoders
- Other Tasks
  - Outlier Detection
  - Novelty Detection
  - Dimensionality Reduction
  - Non-Linear Manifold Detection

→ CS4786 Machine Learning for Data Science
→ CS4850 Math Found for the Information Age
→ INFO 6150 Advanced Topic Models

ML in Computer Visions

- Covered
  - Feedforward Neural Networks
- Other
  - Convolutional Networks
  - More Deep Learning
  - Even more Deep Learning

→CS6670 Computer Vision
→CS4670 Intro Computer Vision

Learning to Act

- Covered
  - Off-policy policy learning
  - Contextual Bandits
- Other
  - Reinforcement learning
  - Markov Decision Processes
  - Model-based vs. model-free
  - On policy vs. off policy
  - Policy gradient

→ CS4700 Artificial Intelligence

ML and Causality

- Covered
  - Potential outcomes model
- Other
  - Observational setting
  - Instrumental variables
  - Continuous treatments
  - Longitudinal treatments
  - Causal discovery
  - Parameter inference
  - Causal networks
  - Structural equation models

ML and Fairness

- Covered
  - Privacy
  - Intelligibility
  - Fairness
- Other
  - Accountability
  - Transparency
  - Algorithms and guarantees

→ INFO4270: Ethics and Policy in DS

FINAL EXAM
SUNDAY 7:00PM, BARTON