Bias on Web Systems

Part of this talk appeared in CACM of June 2018

Ricardo Baeza-Yates
Institute for Experiential AI
Northeastern University

ICWE, May 2021

Institute for *Experiential* AI

What do we mean by *Experiential* AI?

- AI with human in the loop
- AI applied to real-world problems yielding pragmatic working solutions

Why we believe is EAI the right direction?

Much evidence that pragmatic working AI solutions have two characteristics:

1. ***Human-in-the-loop:*** ability to bring human decision-making, common sense reasoning into the solution operation
2. ***Strong dependence on Data:*** ML and DS to leverage more quality (big) data: "We don't have better algorithms… we just have more data"

What is Bias?

- Statistical: significant systematic deviation from a prior (unknown) distribution;
- Cultural: interpretations and judgments phenomena acquired through our life;
- Cognitive: systematic pattern of deviation from norm or rationality in judgment;

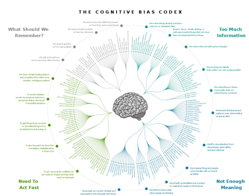

More than 100 cognitive biases!

Motivation: Impact of Bias in Web Systems

- Most web systems are optimized by using implicit user feedback
- However, user data is partly biased to the choices that these systems make
  - Clicks can only be done on things that are shown to us
- As those systems are usually based in ML, they learn to reinforce their own biases, yielding self-fulfilled prophecies and/or sub-optimal solutions
  - For example, personalization and filter bubbles for users
  - but also **echo chambers for (recommender) systems**
- Moreover, sometimes these systems compete among themselves, learning also biases of other systems rather than real user behavior
- Even more, an improvement in one system might be just a degradation in another system that uses a different (even inversely correlated) optimization function
  - For example, user experience vs. monetization

So (Observational) Human Data has Bias

Cultural Biases

- Gender
- Racial
- Sexual
- Age
- Religious
- Social
- Linguistic
- Geographic
- Political
- Educational
- Economic
- Technological

Statistical Biases

- Gathering process
- Sampling process
- Validity (e.g. temporal)
- Completeness
- Noise, spam

Cognitive Biases

Self-selection

Attempt of an unbiased (personal) view on bias in systems

Many people extrapolate results of a sample to the whole population (*e.g.,* social media analysis)

In addition there is bias when measuring bias as well as bias towards measuring it!

What is being fair?

Equality | **Equity** | **Justice**

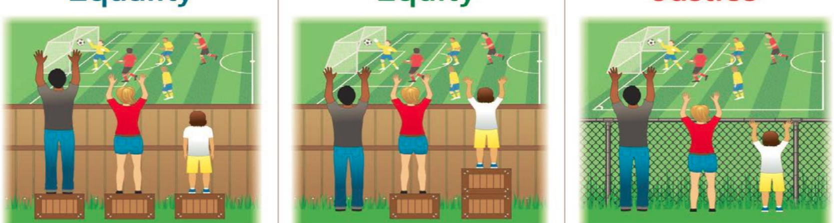

Equality: The assumption is that **everyone benefits from the same supports**. This is equal treatment.

Equity: **Everyone gets the supports they need** (this is the concept of "affirmative action"), thus producing equity.

Justice: All 3 can see the game without supports or accommodations because **the cause(s) of the inequity was addressed**. The systemic barrier has been removed.

A Non-Technical Question

Biased Data → **Algorithm** → **Same or more Bias**

Neutral?
Fair?

Not Always!

Yes, if you harm people

Debias the input
Tune the algorithm
Debias the output

Bias awareness!

ACM US Statement on Algorithm Transparency and Accountability (Jan 2017)

1. Awareness
2. Access and redress
3. Accountability
4. Explanation
5. Data Provenance
6. Auditability
7. Validation and Testing

Systems do not need to be perfect, they just need to be (much?) better than us

Bias in Computing Systems

- The **quality of any algorithm** is bounded by the **quality of the data that uses** (and hence of its users)
- Data bias awareness
  [Gordon & Desjardins; Provost & Buchanan, MLJ 1995]
- Bias in computing systems [Friedman & Nissenbaum 1996]
- Algorithmic fairness
- Key issues for Machine Learning
  - Uniformity of data properties
    - In the Web, distributions resemble a power law
  - Uniformity of error
  - Data sample methodology
    - *E.g.*, sample size to see infrequent events or sampling bias

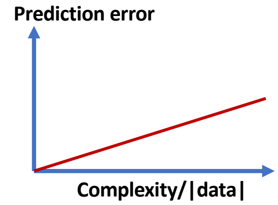

Bias on Web Systems

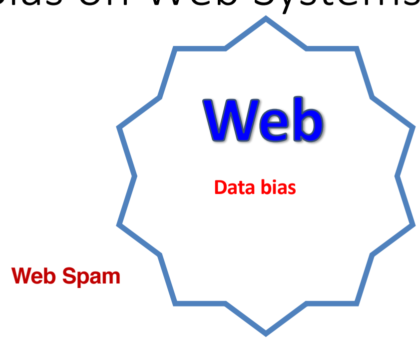

Economic Bias in Links

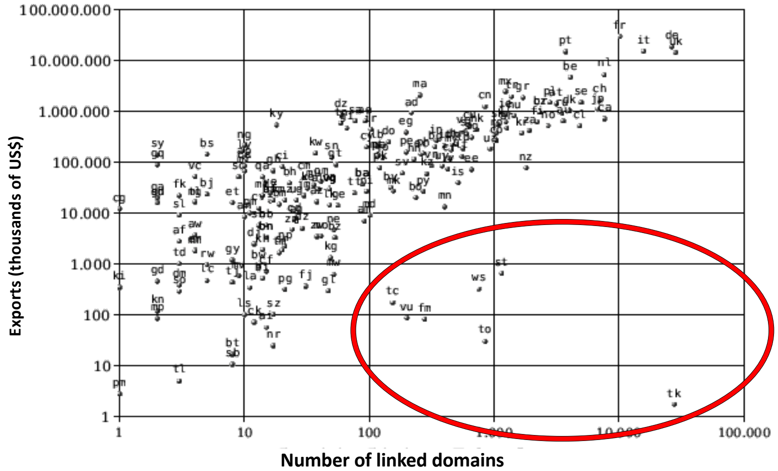

[Baeza-Yates, Castillo & López, 2005]

Economic Bias in Links

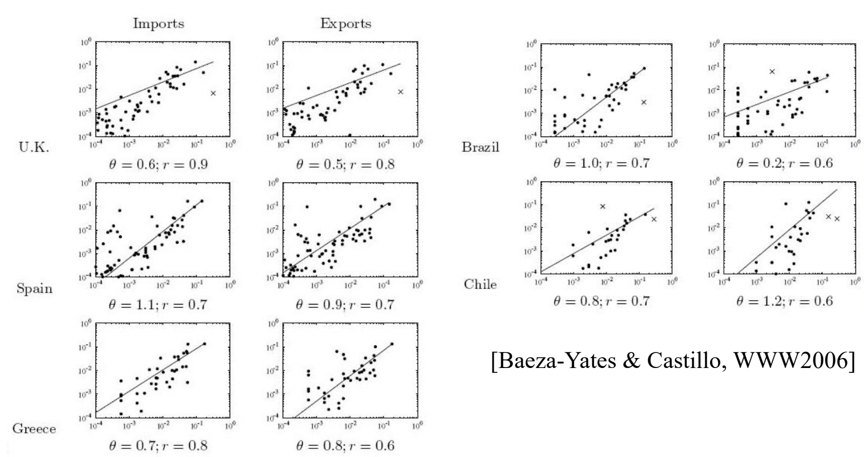

Cultural Bias in Content

Shame

Minimal effort

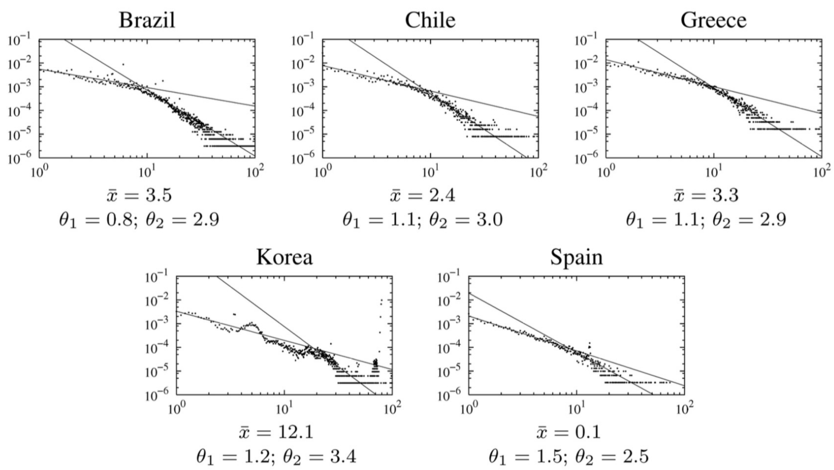

[Baeza-Yates, Castillo, Efthimiadis, TOIT 2007]

Linguistic Bias

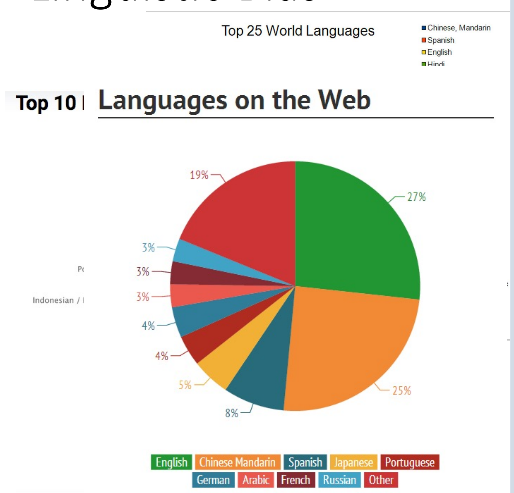

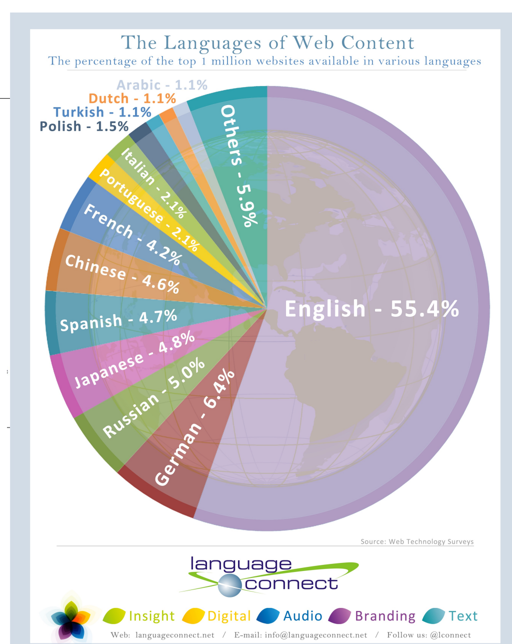

Gender Bias in Content

- Word embedding's in w2vNEWS

Gender stereotype *she-he* analogies.

sewing-carpentry	register-nurse-physician	housewife-shopkeeper
nurse-surgeon	interior designer-architect	softball-baseball
blond-burly	feminism-conservatism	cosmetics-pharmaceuticals
giggle-chuckle	vocalist-guitarist	petite-lanky
sassy-snappy	diva-superstar	charming-affable
volleyball-football	cupcakes-pizzas	hairdresser-barber

Gender appropriate *she-he* analogies.

queen-king	sister-brother	mother-father
waitress-waiter	ovarian cancer-prostate cancer	convent-monastery

Most journalists are men?

[Bolukbasi at al, NIPS 2016]

Yes, about 60 to 70% at work although at college is the inverse

Gender Bias in Translation

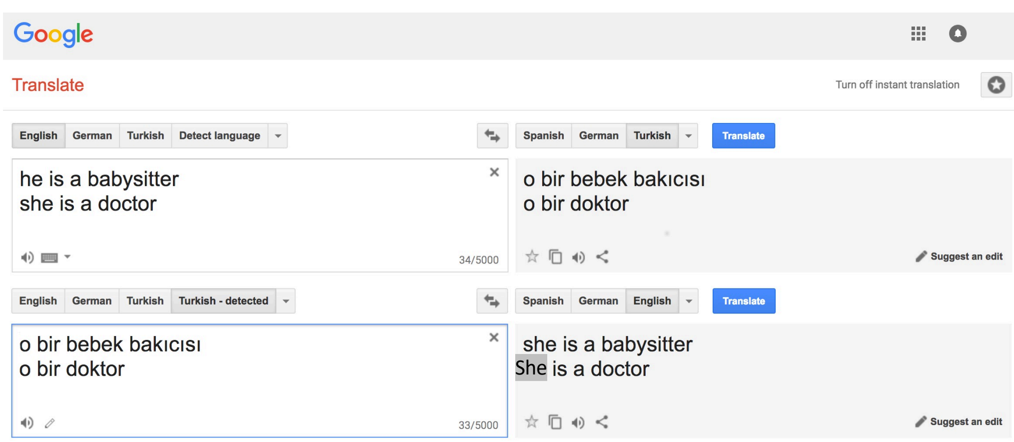

Gender Bias in Content

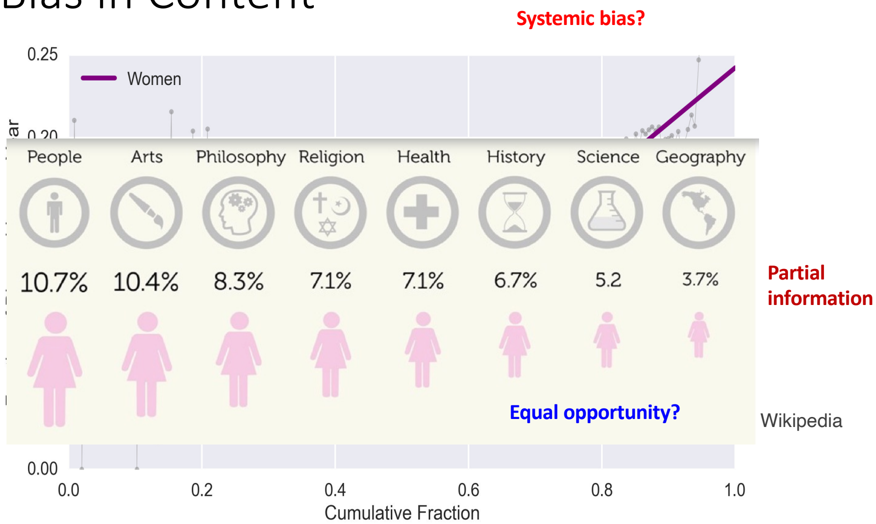

[E. Graells-Garrido et al,. "First Women, Second Sex: Gender Bias in Wikipedia", ACM Hypertext'15]

Engagement or Activity Bias

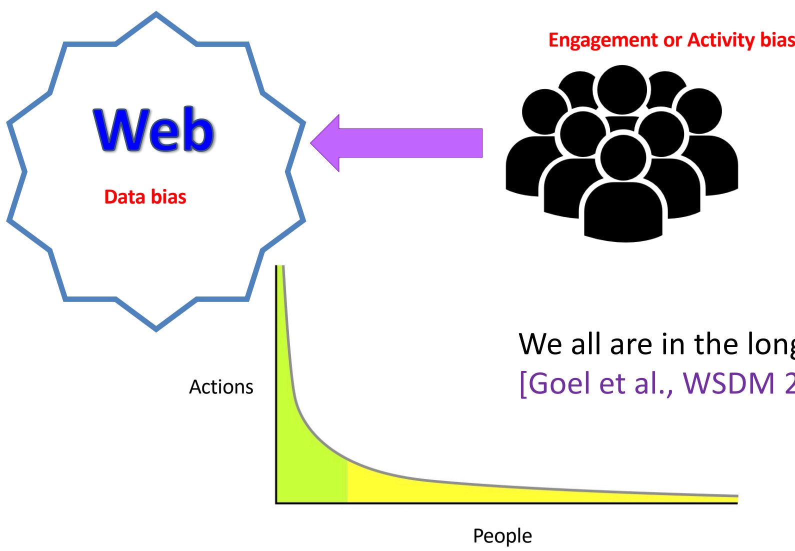

We all are in the long tail!
[Goel et al., WSDM 2010]

Courtesy of Pablo Castells

Popularity Bias in Recommender Systems

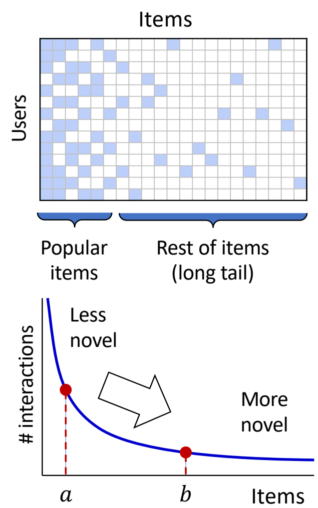

- Take care to recommended items that are not too popular
- Metrics

$$nov(i) = 1 - \frac{\#\ \text{ratings of } i}{\#\ \text{users}}$$

- Novelty enhancement
- Problem solved! …really?

[Vargas & Castells, RecSys 2011]

Courtesy of Pablo Castells

A self-fulfilling prophecy?

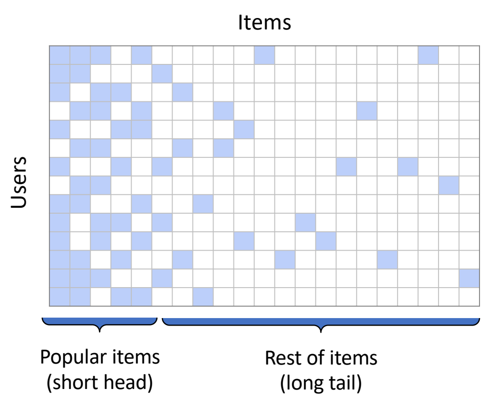

Observed user-item interaction

Unobserved preference

Ratings are missing not at random (MNAR)

Courtesy of Pablo Castells

A self-fulfilling prophecy?

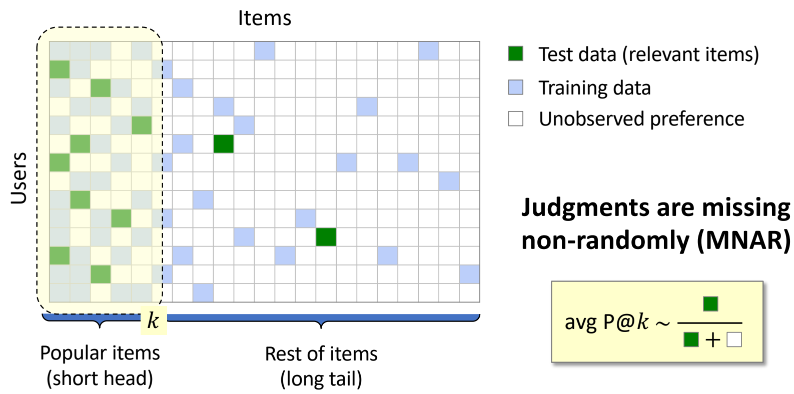

Worse yet: user-system reinforcement loop (more later)

[Marlin et al., RecSys 2010]
[Steck, RecSys 2010, 2011]
[Fleder & Hossanagar, Management Sciences 2009]

Courtesy of Pablo Castells

A problem for IR evaluation methodology!

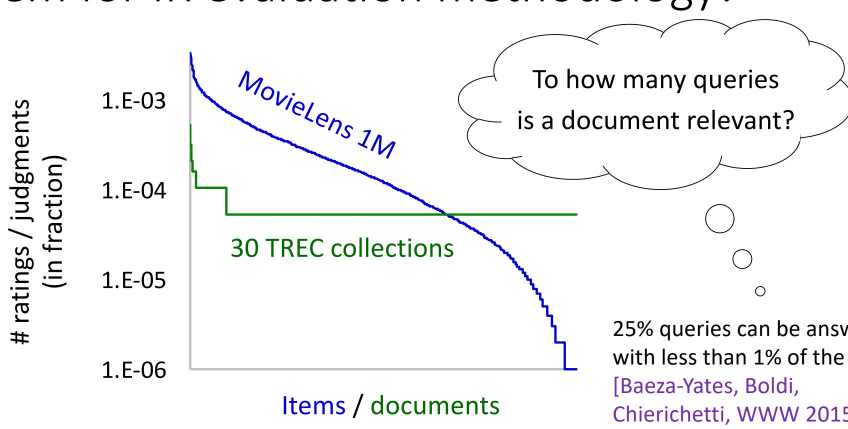

25% queries can be answered with less than 1% of the URLs!
[Baeza-Yates, Boldi, Chierichetti, WWW 2015]

[Bellogín, Castells & Cantador IRJ 2017]

How different or similar are we to each other?

People are very different from each other

People are partly similar, partly different

People are identical to each other

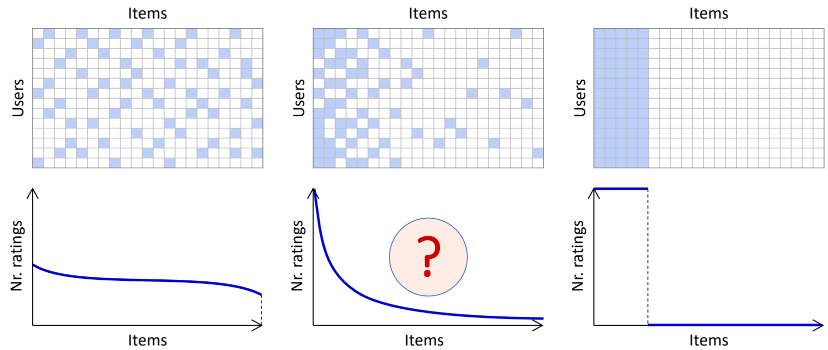

Popularity does not help at all

Popularity is optimal
Personalization is not useful

Courtesy of Pablo Castells

Get rid of the popularity bias!

- In the rating split
  [Bellogín, Castells & Cantador, IRJ 2017]
- In the metrics
  - Stratified recall
    [Steck, RecSys 2011]
  - Importance propensity scoring
    [Yang et al., RecSys 2018]
- In the algorithms
  [Steck, RecSys 2011]
  [Lobato et al., ICML 2014]
  [Jannach et al., UMUAI 2015]
  [Cañamares & Castells, SIGIR 2018, **best paper award**]

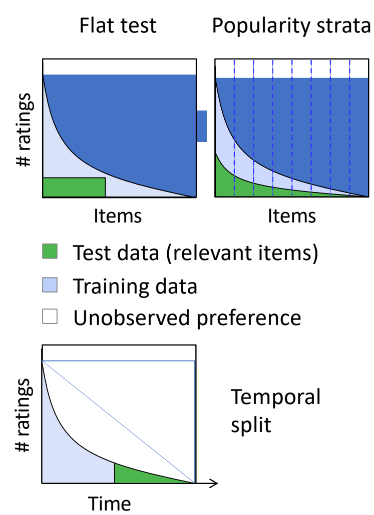

Recommending within the Long Tail

- Exploit the context (and deep learning!)

  91% accuracy to predict the next app you will use
  [Baeza-Yates et al, WSDM 2015]

- Personalization vs. Contextualization
  Break the filter bubble! (more later)

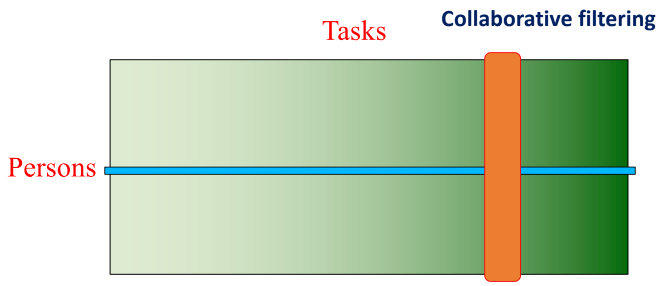

Activity Bias also Affects Content

Most users are passive (*i.e.,* more than 90%) – wisdom of crowds is a partial illusion
Hence, which percentage of **active** users produce 50% of the content?

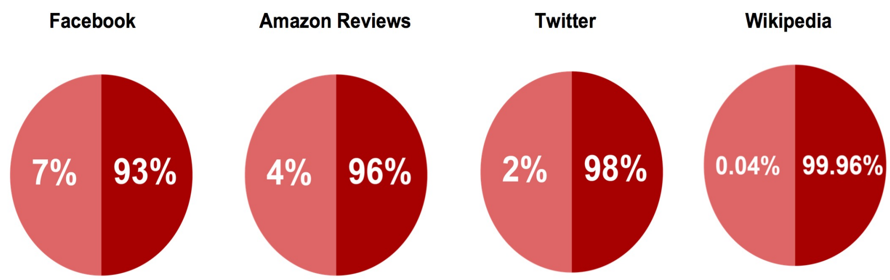

[Baeza-Yates & Saez-Trumper, ACM Hypertext 2015]

Social Bias

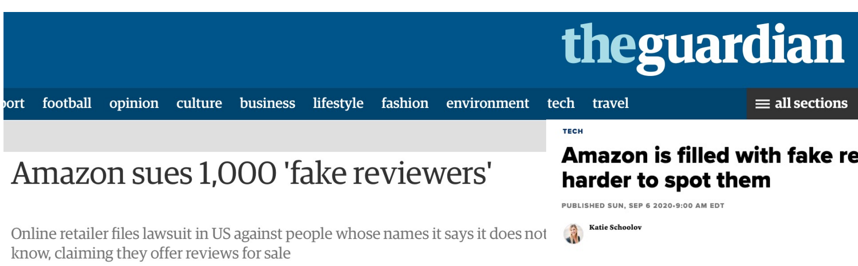

theguardian

ort football opinion culture business lifestyle fashion environment tech travel all sections

Amazon sues 1,000 'fake reviewers'

Online retailer files lawsuit in US against people whose names it says it does not know, claiming they offer reviews for sale

October 2015

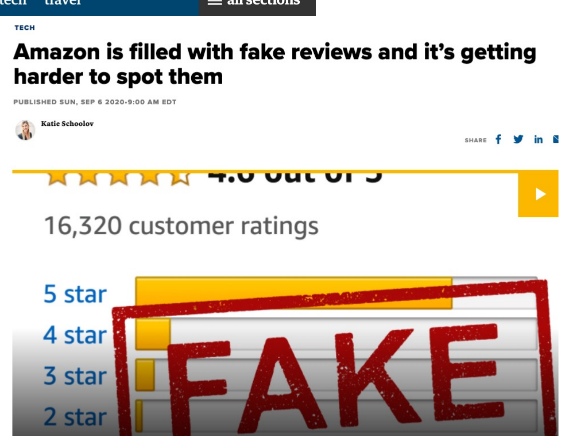

TECH

Amazon is filled with fake reviews and it's getting harder to spot them

PUBLISHED SUN, SEP 6 2020•9:00 AM EDT

Katie Schoolov

SHARE

Quality of Content?

- Adding content ⇛ Adding Wisdom ?
- We use Amazon's Reviews helpfulness
- Content-based-wise users

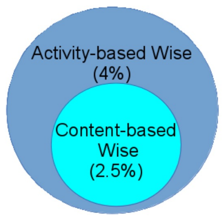

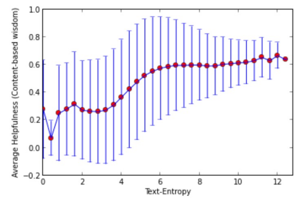

[Baeza-Yates & Saez-Trumper, ACM Hypertext 2015]

Wisdom of a Few?

Which percentage of **active** users produce 50% of the content?
Similar to the 90-9-1 rule of Internet participation [Nielsen 2006]

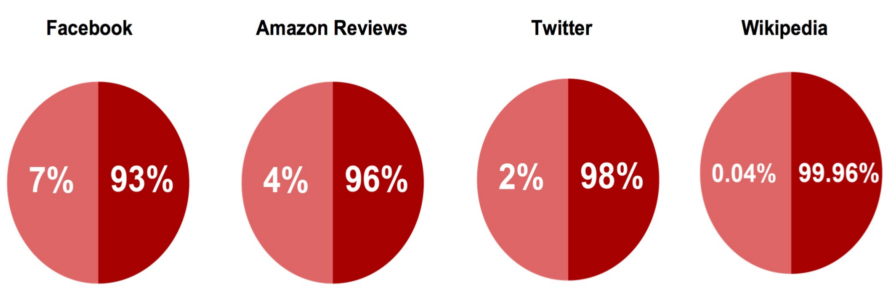

[Baeza-Yates & Saez-Trumper, ACM Hypertext 2015]

Attention Bias: The Digital Desert

[Baeza-Yates & Saez-Trumper, ACM Hypertext 2015]

Bias in the Web

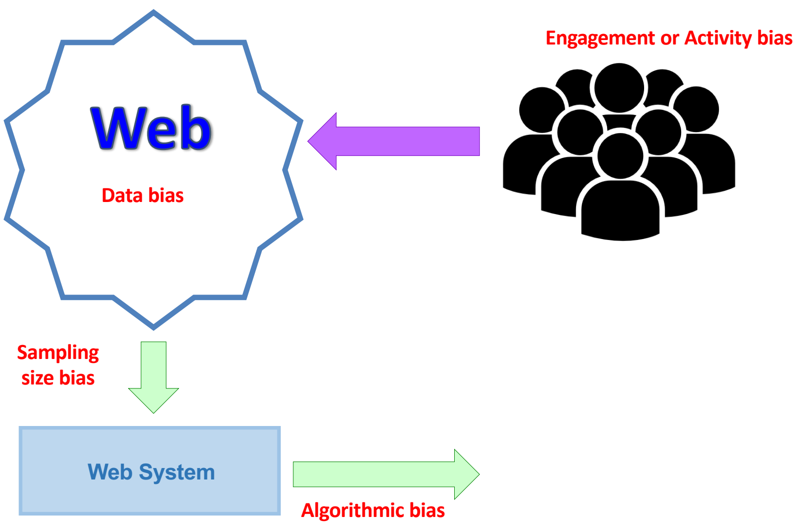

Sample Size?

- If we want to estimate the frequencv of queries that appear with probability at least p with a certain relative error ϵ we can use the standard binomial error formula $\sqrt{(1-p)/np}$ which works well for p near ½ **but not for p near 0**
- Better is the Agresti-Coull technique (also called *take 2*) which gives:

$$n \geq Z^2_{1-\alpha/2}\left(\frac{p'(1-p')}{\epsilon^2} - 1\right)$$

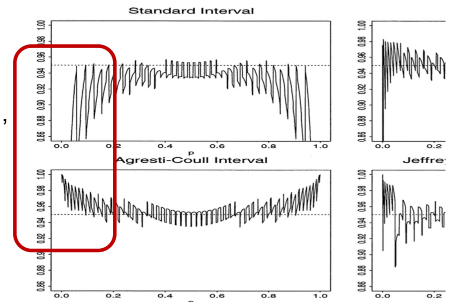

where Z is the inverse of the standard normal distribution, $1-\alpha$ is the confidence interval and $p' = p + Z^2/2$

- If p = 0.1, $1-\alpha$ is 80% and ϵ is 10%, we get n = 2342. The standard formula gives n = 900!

[Brown, Cai & DasGupta, Statistical Science, 2001]
[Baeza-Yates, SIGIR 2015, Industry track]

Sampling Techniques

- Standard technique:

$$p_q \approx \widehat{p}_q(\mathcal{S}) = \frac{f_q(\mathcal{S})}{\sum_{q' \in \mathcal{S}} f_{q'}(\mathcal{S})}$$

- A good sample should cover well all the items distribution but this does not work with very skewed distributions.

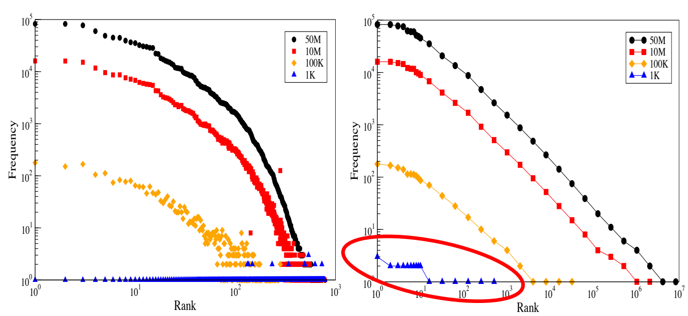

[Zaragoza et al, CIKM 2010]

Stratified Sampling Example

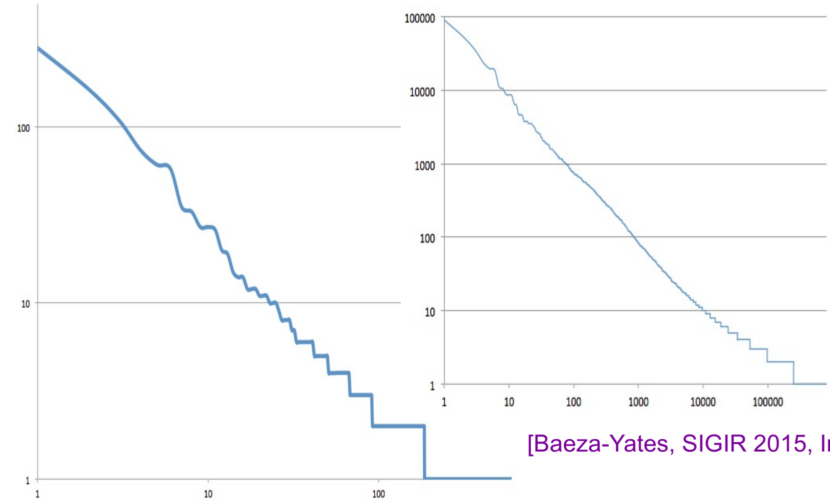

[Baeza-Yates, SIGIR 2015, Industry track]

Extreme Algorithmic Bias

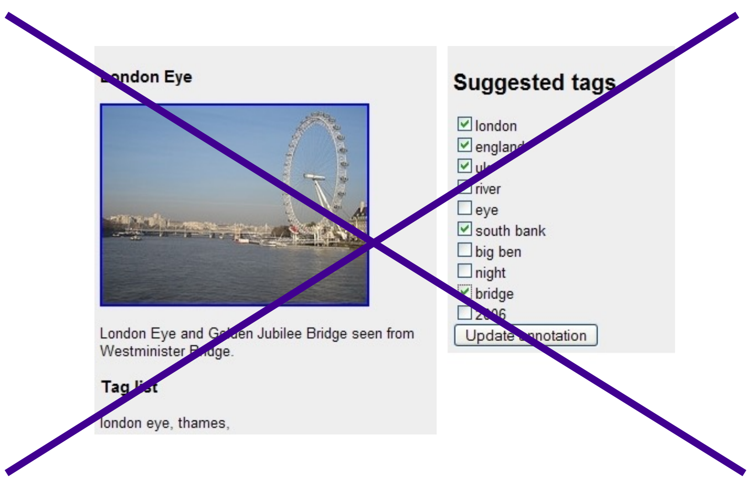

Bias in the User Interaction

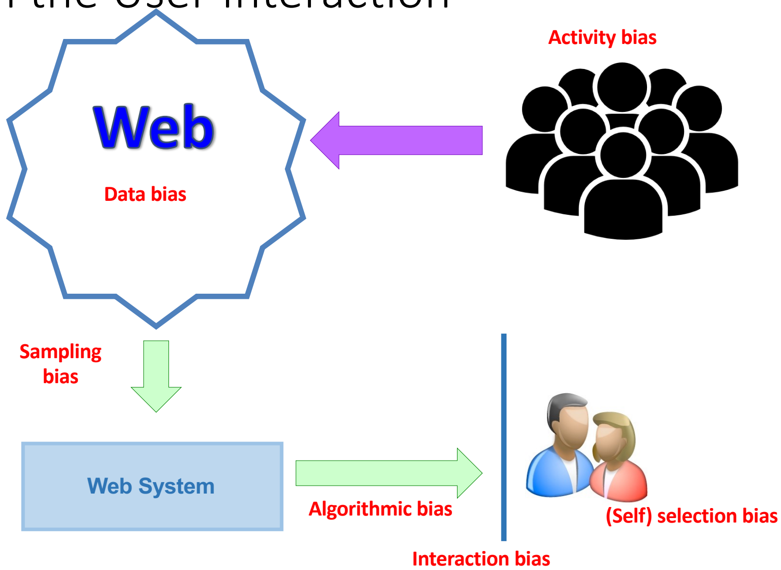

Most surely you will read this at the end

You will read this first

Then you will read this

Bias in the Interaction

Exposure or Presentation bias

Position bias
Ranking bias

Social bias

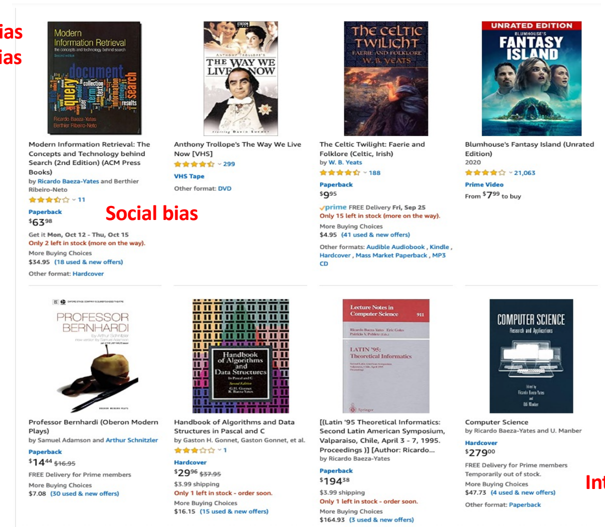

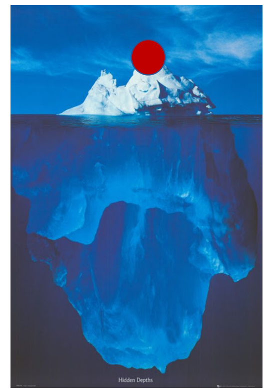

Interaction bias

Dependencies: A Cascade of Biases!

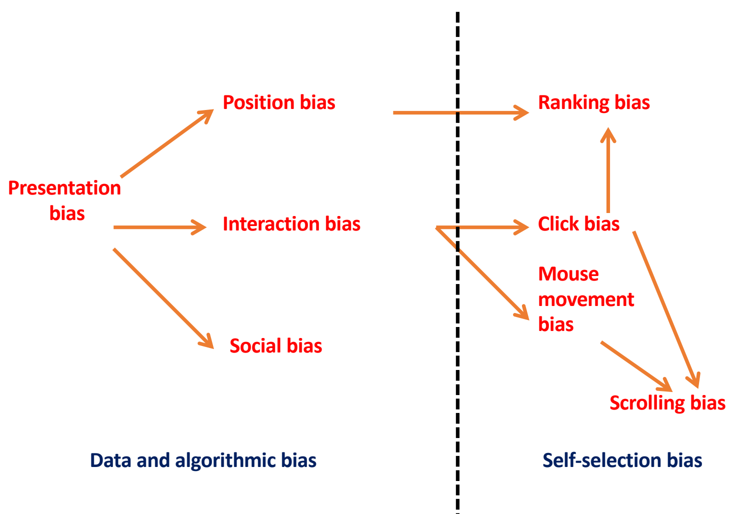

Ranking Bias in Web Search

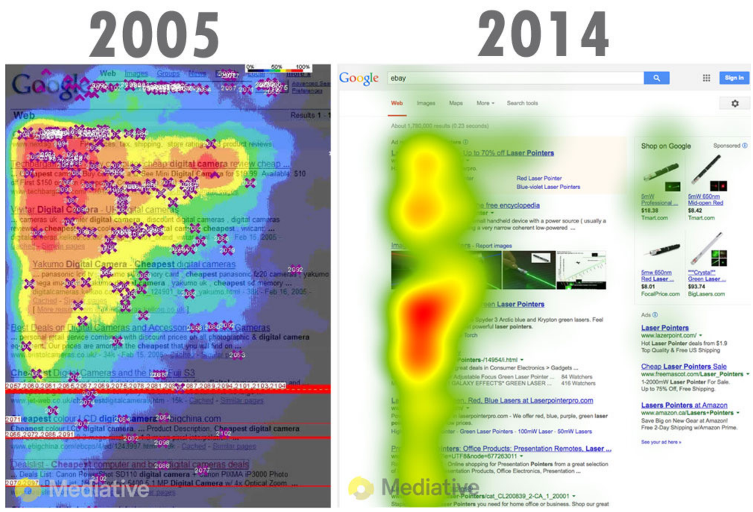

[Mediative Study, 2014]

Ranking Bias: Click Bias in Web Search

- Ranking & next page bias

Debiasing Search Clicks and Other Biases

Clicks as implicit positive user feedback

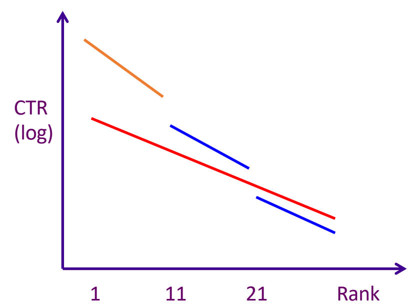

[Dupret & Piwowarski, SIGIR 2008]
[Chapelle & Zhang, WWW 2009]
[Dupret & Liao, WSDM 2010, **best paper**]

Debias the input

Learning to Rank with bias
[Joachims et al., WSDM 2017, **best paper**]
+ many other papers

Tune the algorithm

Fair rankings
[Zehlike et al., CIKM 2017]
+ many other papers

Debias the output

Vicious Cycle of Bias

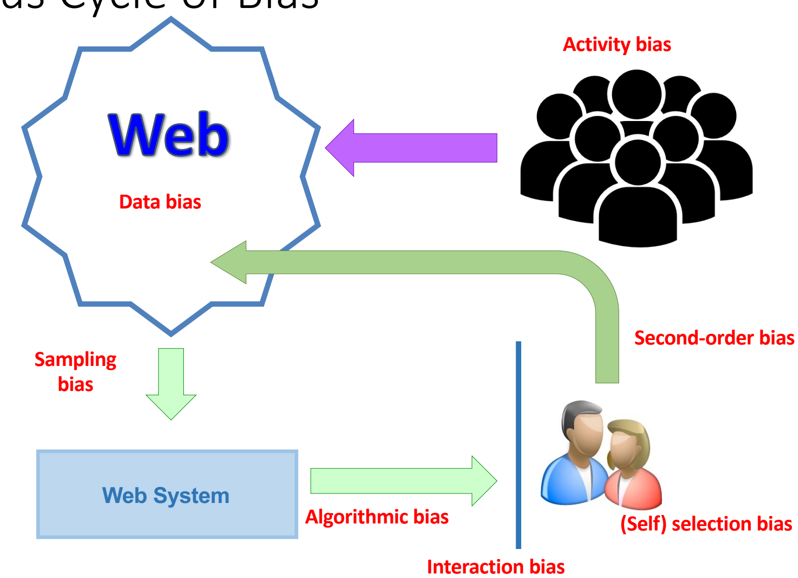

Second Order Bias in Web Content

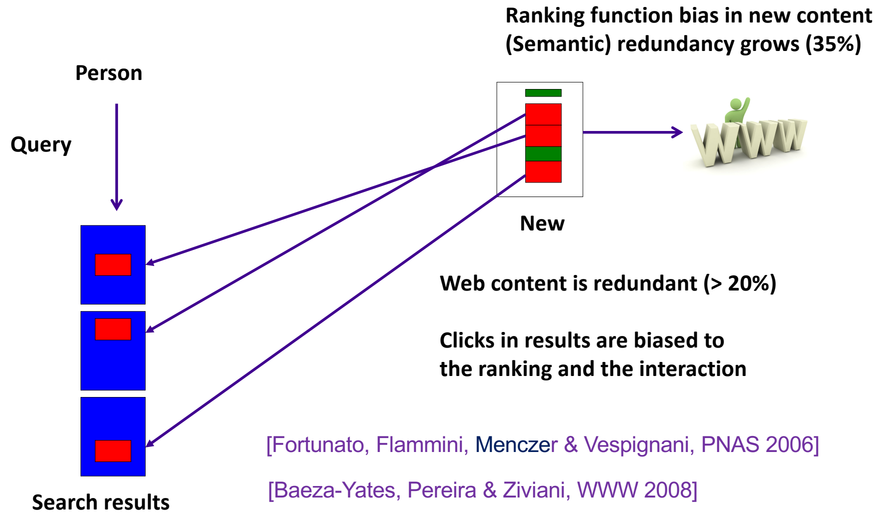

Bias in the Feedback Loop

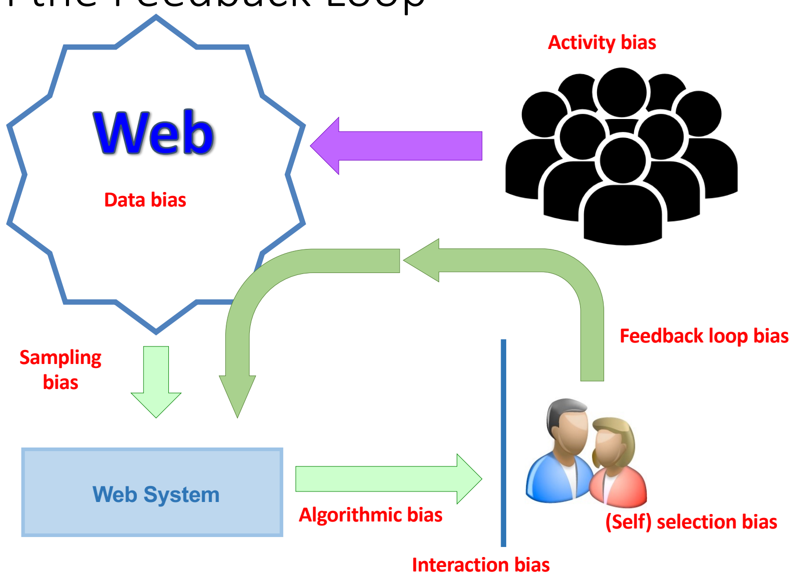

Bias due to Personalization

- Partially the effect of self-selection bias
- Avoid the rich get richer and
  poor get poorer effect
- Avoid the echo chamber by
  empowering the tail

Partial solutions:
- Diversity
- Novelty
- Serendipity
- My dark side

Cold start problem solution: Explore & Exploit

**How much exploration is needed
to counteract exposure bias?**

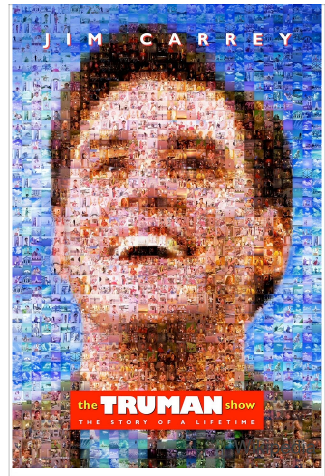

[Eli Pariser, The Filter "Bubble", 2011]

Echo Chambers in Feedback Loops

- For users
  - Filter bubbles
  - Degenerate feedback loops (e.g., YouTube autoplay)
- For systems
  - Short-term greedy optimization
  - The system is partly writing its own future (exposure bias)
  - Partial knowledge of the world if not enough exploration/traffic
  - The **system itself is also in a bubble**!

Users' Echo Chambers in Feedback Loops

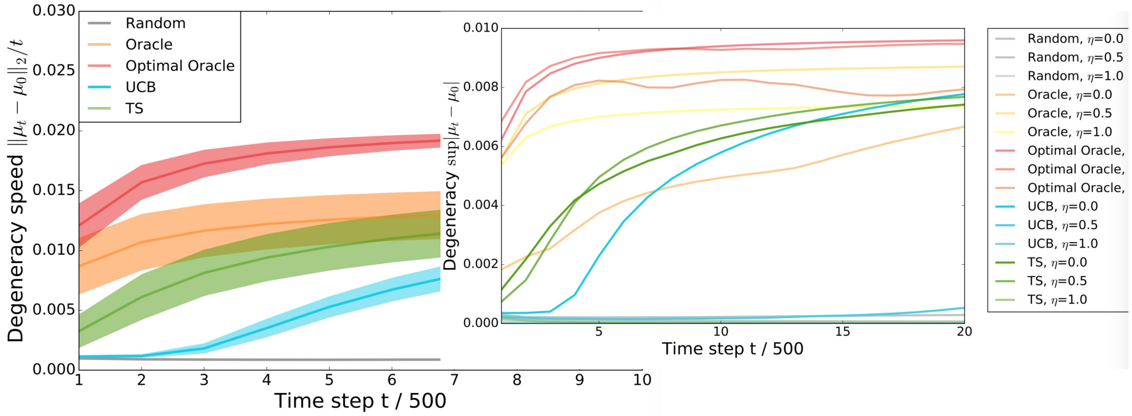

[Jiang et al. Degenerate Feedback Loops in Recommendation Systems, AAAI 2019]

Echo Chamber of the Recommender System

- Short-term greedy optimization, partial knowledge of the world
- Long-term revenue optimization is not achieved
- **Disparate impact**: unfair ecommerce/information markets

- Can we do better?
- Yes, if the amount of new traffic allows **enough** exploration for new items or any other changes in your world

$$\Delta Traffic \geq \text{Approximate}(\Delta World)$$

- Otherwise we will live in a sup-optimal solution

Fairness and Ethics

- Consumers & long tail items/players are discriminated
- Matthew effect again: rich get richer, poor get poorer
- Unfair markets are unhealthy and hence less stable in the long term

- Internet Companies Antitrust - Advertising Transparency
  - Amazon's Antitrust Paradox [Khan, 2017]
- Should marketplaces sell in their own marketplace?
  - Yes, but with regulations [Hagiu, Teh & Smith, 2020]
  - Is data asymmetry ethical? (not new, but gets amplified in e-commerce)
- Fair markets could be better revenue wise
  - Fairness trade-offs [Mehrotra et al., 2018]

Our Professional Biases

- Problems
  - Our **big data and deep learning bias**: **small data** is more frequent & harder [Baeza-Yates, KD Nuggets, 2018]
- Design and Implementation
  - Do systems reflect the characteristics of the designers?
  - Do systems reflect the characteristics of the coders?
- Evaluation [Silberzahn et al., COS, Univ. of Virginia, 2015] [Johansen et al., Norway, 2020]
  - Choose the right experiment
  - Choose the right test data
    - Pool bias in search test collections [Lipani et al., SIGIR 2015, CIKM 2016]
  - Choose the right metric(s)
  - Choose the **right baseline(s)**
  - Julio Gonzalo's talk: http://tiny.cc/ESSIR2019-juliogonzalo

What we can do?

- Data
  - Analyze for known and unknown biases, debias/mitigate when possible/needed
  - Recollect more data for sparse regions of the problem
  - Do not use attributes associated directly/indirectly with harmful bias
- Design and Implementation
  - Make sure that the model is **aware** of the biases all the time
  - Let experts/colleagues/users contest every step of the process
- Interaction
  - Make sure that the user is **aware** of the biases all the time
  - Give more control to the user
- Evaluation
  - Do not fool yourself!

The Web Works Thanks to Bias!

- Web traffic
  - Local caching
  - Proxy/network caching
- Search engines
  - Answer caching
  - Essential web pages
    - 25% queries can be answered with less than 1% of the URLs!
      [Baeza-Yates, Boldi, Chierichetti, WWW 2015]
- E-Commerce
  - Large fraction of revenue comes from few popular items
  - But a large fraction of revenue goes to the marketplace owner

Activity bias

(Self) selection bias

Final Take-Home Messages

- Systems are a mirror of us, the good, the bad and the ugly
- The Web amplifies everything, but always leaves traces

- We need to be aware of our **own biases**!
- We have to be aware of the biases and contrarrest them to stop the **vicious bias cycle**
- We should be **fair**
- **Plenty** of open research problems! (in **small data** even more!)

Questions?

ASIST 2012 Book of the Year Award (Biased Ad)

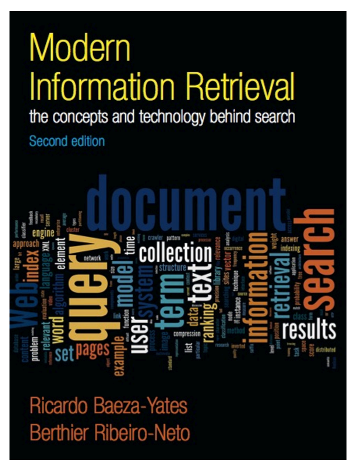

New Conferences that started in 2018:

AAAI/ACM Conference on AI, Ethics, and Society
http://www.aies-conference.com

ACM FAccT: Fairness, Accountability, and Transparency
http://facctconference.org

Contact: rbaeza@acm.org
www.baeza.cl
@polarbearby

Biased Questions?