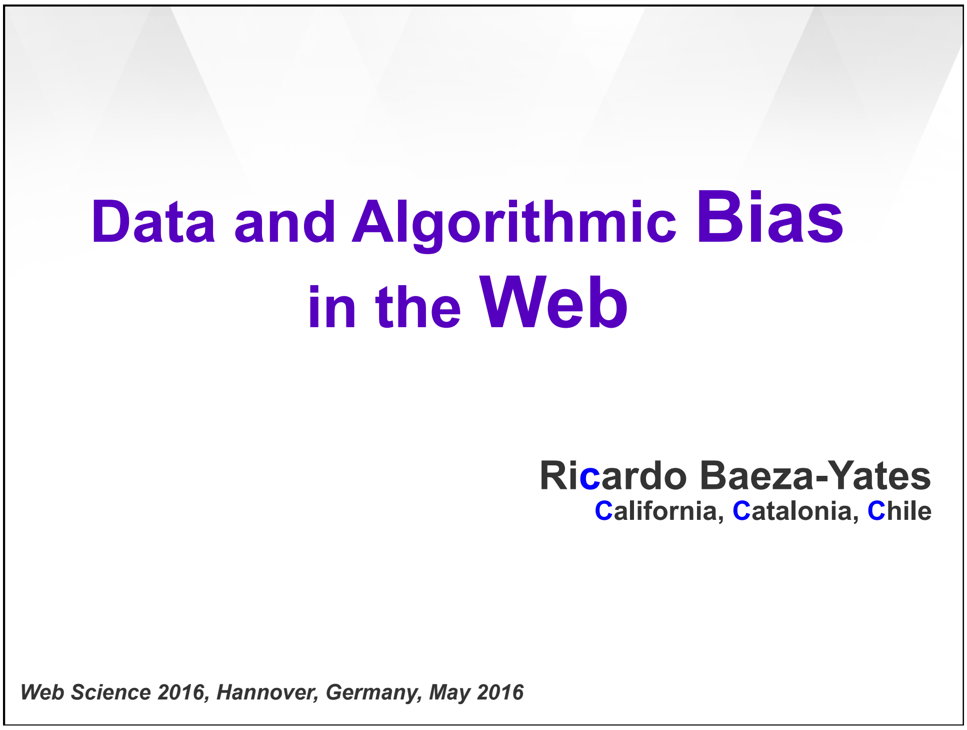

Data and Algorithmic Bias in the Web

Ricardo Baeza-Yates
California, Catalonia, Chile

Web Science 2016, Hannover, Germany, May 2016

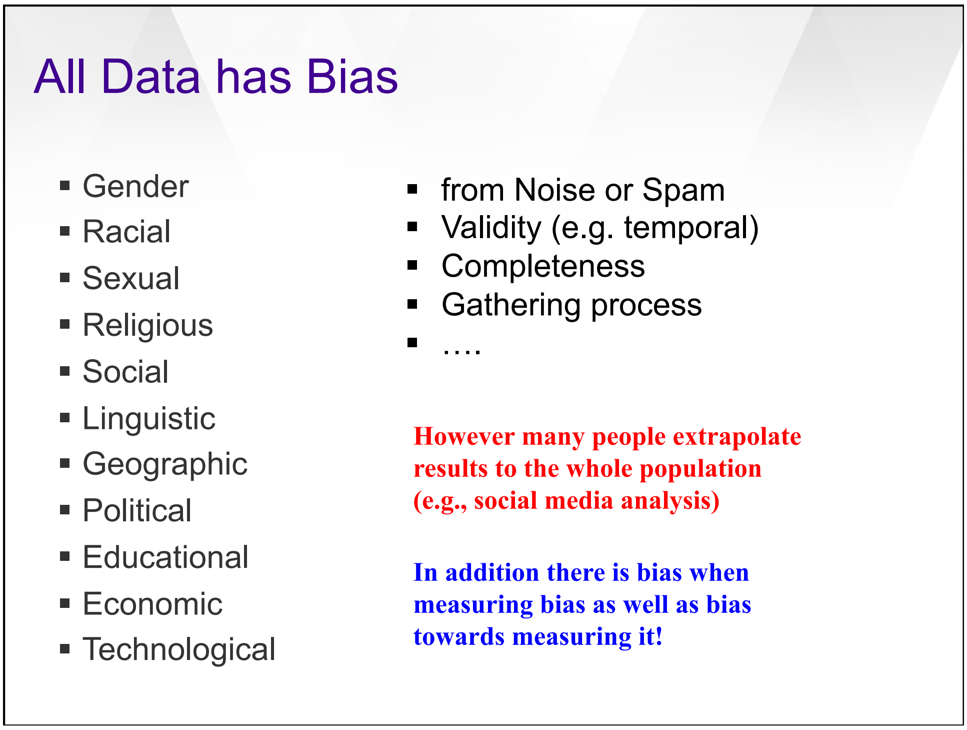

All Data has Bias

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

- from Noise or Spam
- Validity (e.g. temporal)
- Completeness
- Gathering process
- ….

However many people extrapolate results to the whole population (e.g., social media analysis)

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

Yes, We Live in a (Very) Biased World!

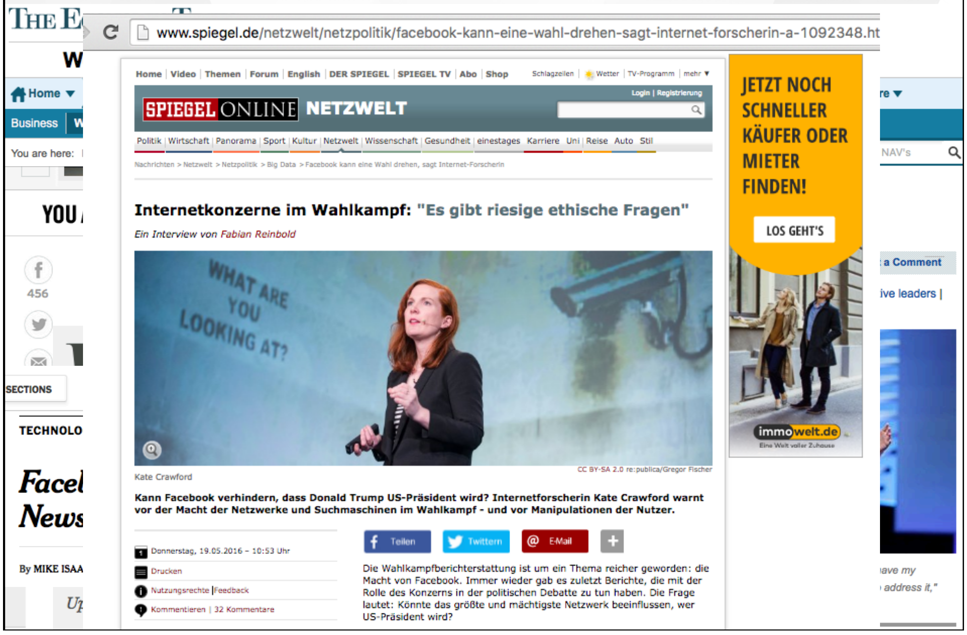
www.spiegel.de/netzwelt/netzpolitik/facebook-kann-eine-wahl-drehen-sagt-internet-forscherin-a-1092348.ht
Home | Video | Themen | Forum | English | DER SPIEGEL | SPIEGEL TV | Abo | Shop
Schlagzeilen | Wetter | TV-Programm | mehr ▾
Login | Registrierung
SPIEGEL ONLINE NETZWELT
Politik | Wirtschaft | Panorama | Sport | Kultur | Netzwelt | Wissenschaft | Gesundheit | einestages Karriere Uni | Reise Auto Stil
Nachrichten > Netzwelt > Netzpolitik > Big Data > Facebook kann eine Wahl drehen, sagt Internet-Forscherin
WHAT ARE YOU LOOKING AT?
Teilen
Twittern
E-Mail
+
JETZT NOCH SCHNELLER KÄUFER ODER MIETER FINDEN!
LOS GEHT'S
immowelt.de
Eine Welt voller Zuhause
THE E
W
Home ▾
Business
You are here:
YOU
456
SECTIONS
TECHNOLO
Facel
News
By MIKE ISAA
a Comment
ive leaders |
NAV's
ave my
address it,"

Internetkonzerne im Wahlkampf: "Es gibt riesige ethische Fragen"

Ein Interview von Fabian Reinbold

Kate Crawford

CC BY-SA 2.0 re:publica/Gregor Fischer

Kann Facebook verhindern, dass Donald Trump US-Präsident wird? Internetforscherin Kate Crawford warnt vor der Macht der Netzwerke und Suchmaschinen im Wahlkampf - und vor Manipulationen der Nutzer.

Donnerstag, 19.05.2016 – 10:53 Uhr

Drucken

Nutzungsrechte | Feedback

Kommentieren | 32 Kommentare

Die Wahlkampfberichterstattung ist um ein Thema reicher geworden: die Macht von Facebook. Immer wieder gab es zuletzt Berichte, die mit der Rolle des Konzerns in der politischen Debatte zu tun haben. Die Frage lautet: Könnte das größte und mächtigste Netzwerk beeinflussen, wer US-Präsident wird?

A Non-Technical Question

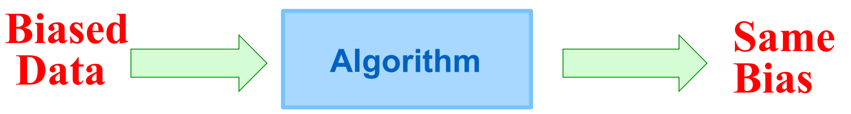

Biased Data

Algorithm

Same Bias

Neutral?

Not Always!

Unbias the data
Tune the algorithm
Unbias the output

Bias awareness!

Big Data and Bias

- The quality of any algorithm is bounded by the quality of the data that uses
- Data bias awareness
- 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 issues

Total number of websites (linear scale)

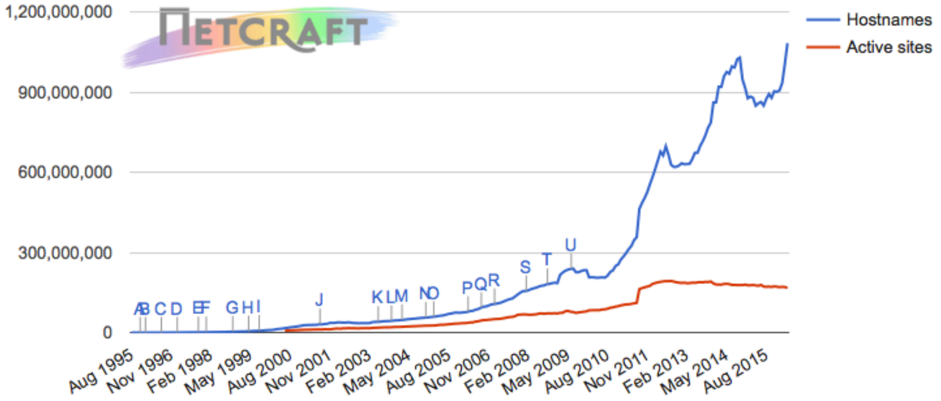

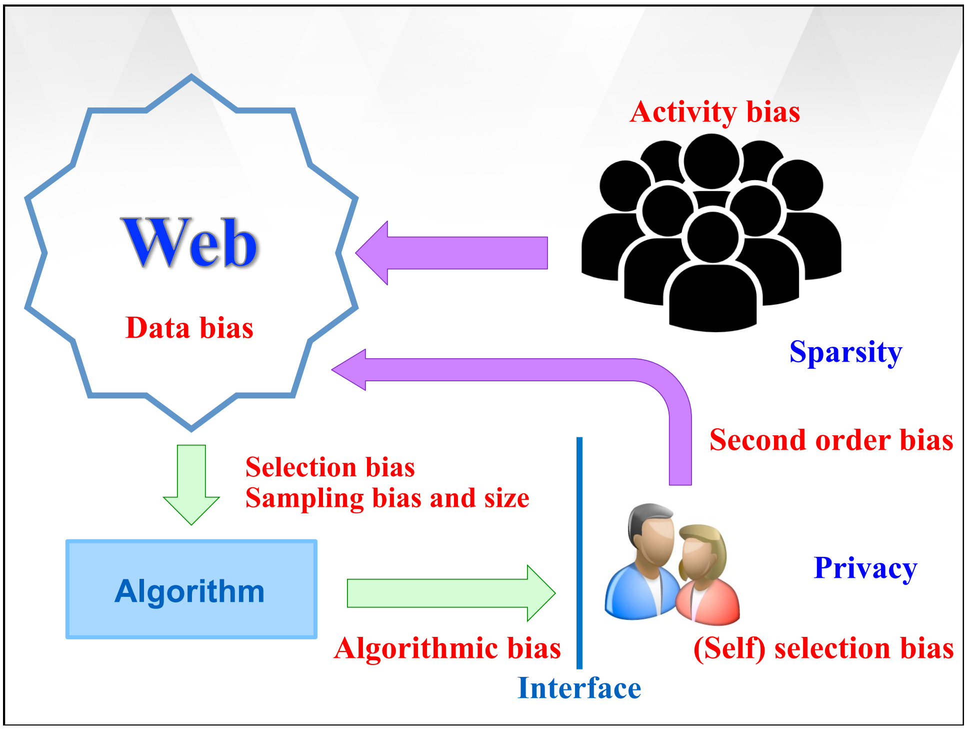

Web

Data bias

Activity bias

Sparsity

Second order bias

Selection bias
Sampling bias and size

Algorithm

Algorithmic bias

Interface

Privacy

(Self) selection bias

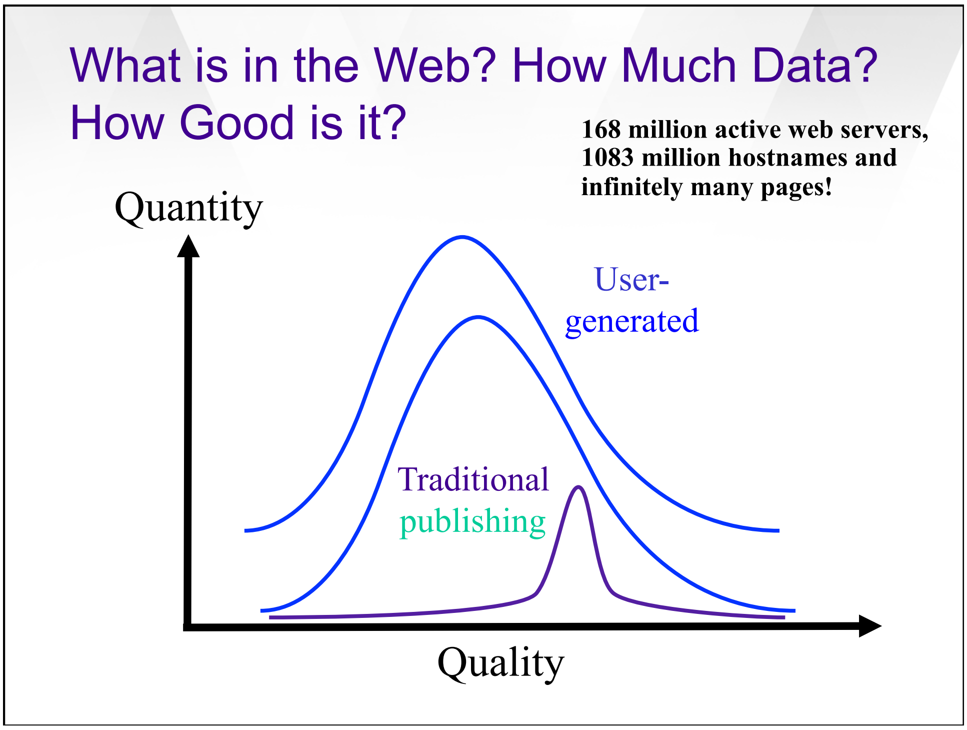

What is in the Web? How Much Data? How Good is it?

168 million active web servers, 1083 million hostnames and infinitely many pages!

What else is in the Web?

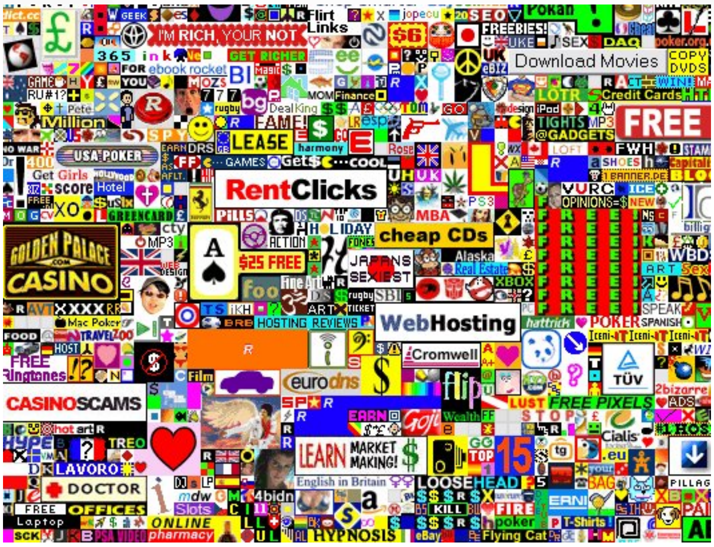

Noise and Spam

- Noise may come from many places:
  - Instruments that measure (e.g., IoT)
  - How we interpret the data
- Spam is everywhere
- Fight both with the wisdom of the crowds

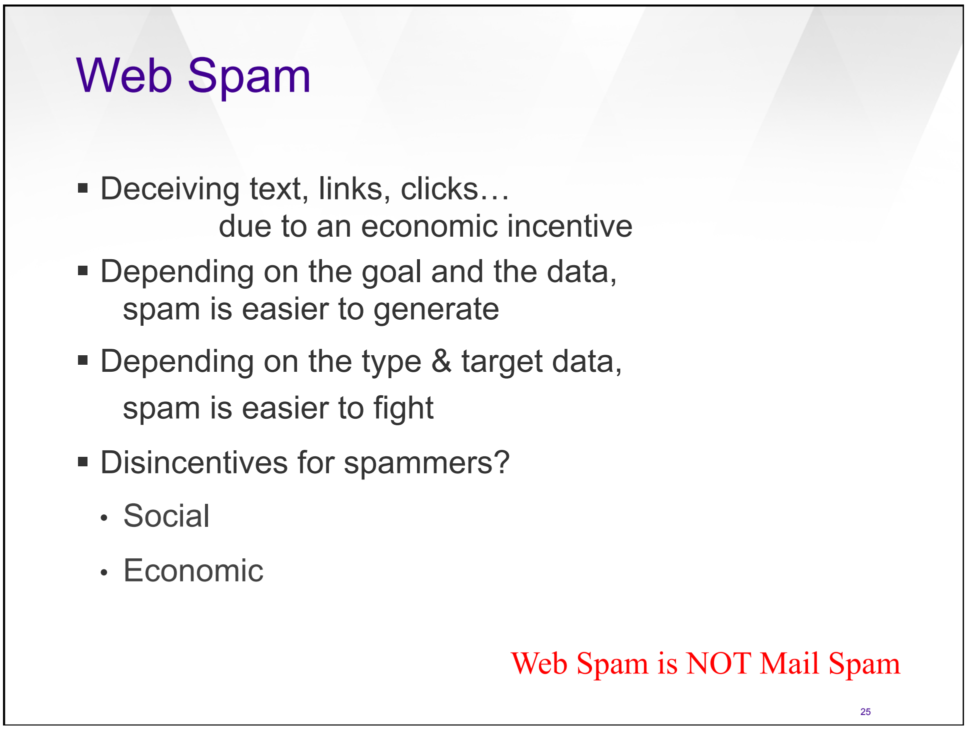

Web Spam

- Deceiving text, links, clicks…
  due to an economic incentive
- Depending on the goal and the data,
  spam is easier to generate
- Depending on the type & target data,
  spam is easier to fight
- Disincentives for spammers?
  - Social
  - Economic

Web Spam is NOT Mail Spam

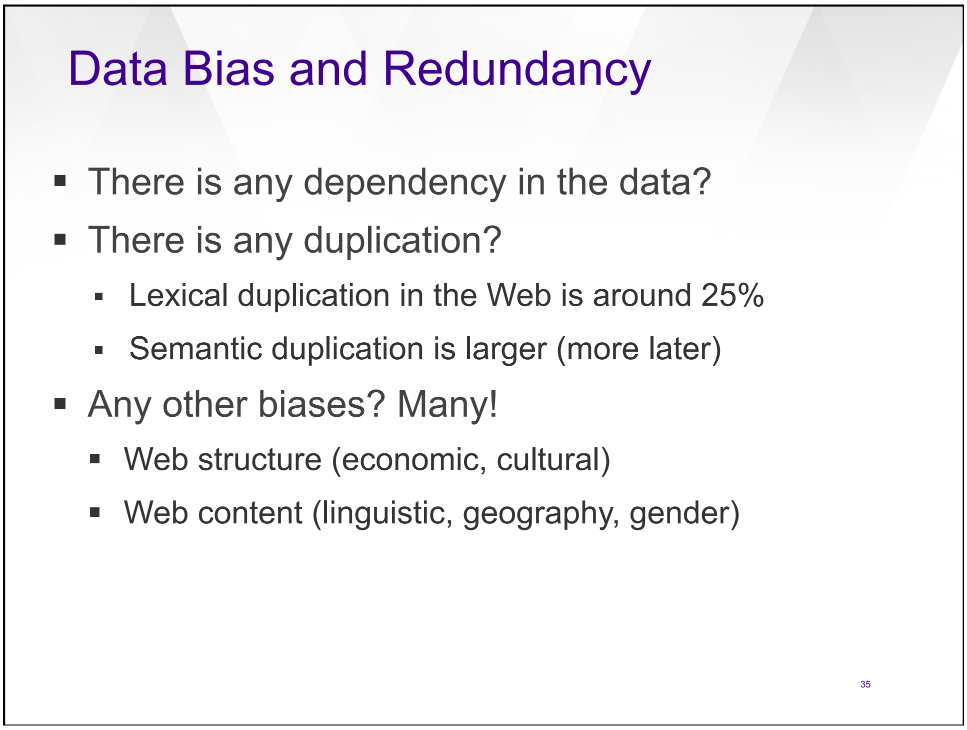

Data Bias and Redundancy

- There is any dependency in the data?
- There is any duplication?
  - Lexical duplication in the Web is around 25%
  - Semantic duplication is larger (more later)
- Any other biases? Many!
  - Web structure (economic, cultural)
  - Web content (linguistic, geography, gender)

Economic Bias in Links

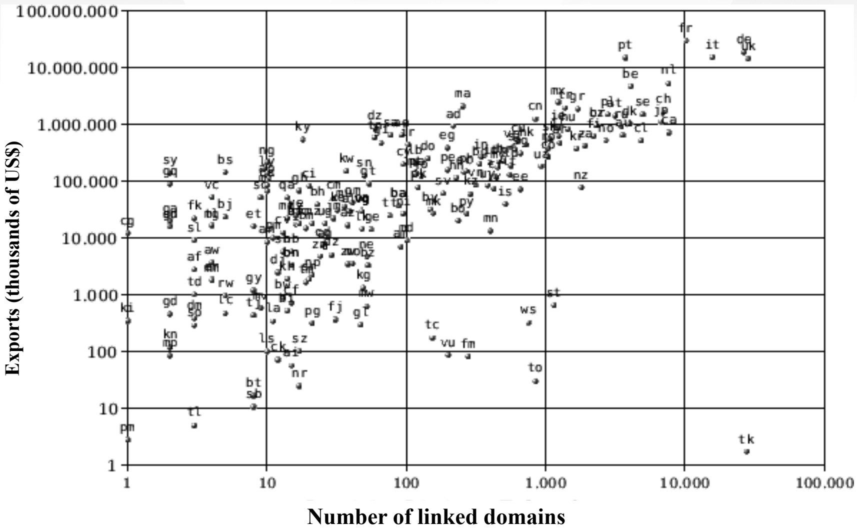

[Baeza-Yates, Castillo & López. Characteristics of the Web of Spain. The Information Professional (Spanish), 2006, vol. 15, n. 1, pp. 6-17]

Exports/Imports vs. Domain Links

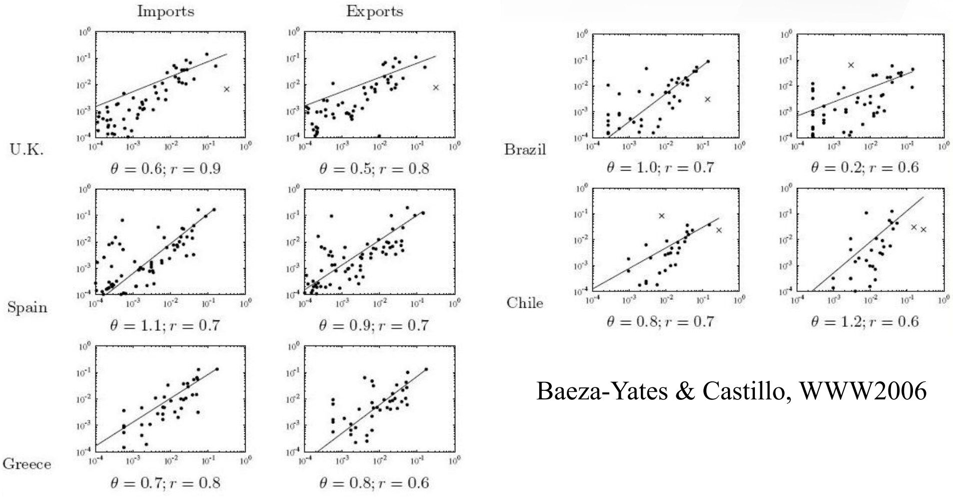

Baeza-Yates & Castillo, WWW2006

Website Structure

Shame

Minimal effort

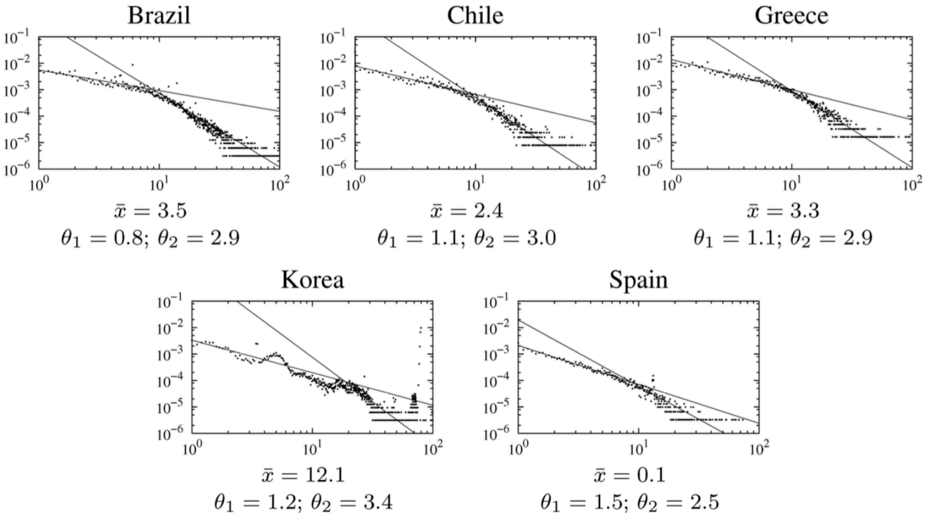

Brazil: $\bar{x} = 3.5$, $\theta_1 = 0.8;\ \theta_2 = 2.9$

Chile: $\bar{x} = 2.4$, $\theta_1 = 1.1;\ \theta_2 = 3.0$

Greece: $\bar{x} = 3.3$, $\theta_1 = 1.1;\ \theta_2 = 2.9$

Korea: $\bar{x} = 12.1$, $\theta_1 = 1.2;\ \theta_2 = 3.4$

Spain: $\bar{x} = 0.1$, $\theta_1 = 1.5;\ \theta_2 = 2.5$

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

Linguistic Bias

Top 25 World Languages

- ■ Chinese, Mandarin
- ■ Spanish
- ■ English
- ■ Hindi

Top Ten Languages in the Internet in millions of users - November 2015

Languages on the Web

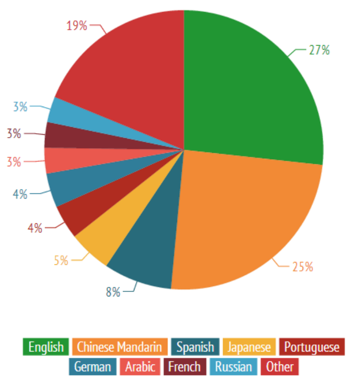

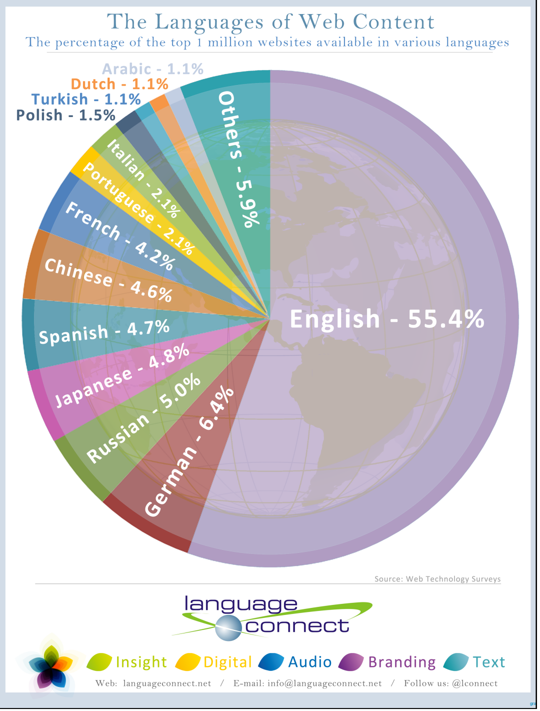

Geographical Bias

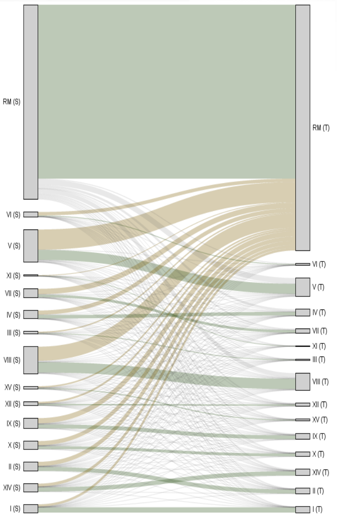

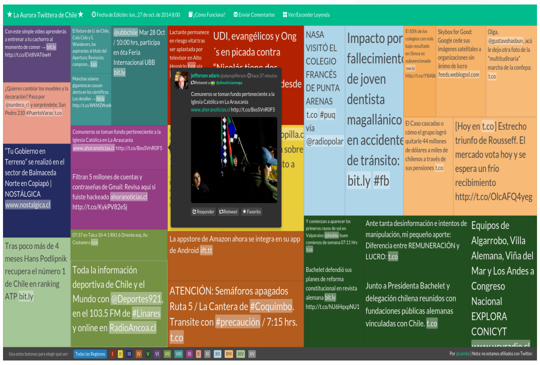
★ La Aurora Twittera de Chile ★
Fecha de Edición: lun., 27 de oct. de 2014 8:00
¿Cómo Funciona?
Enviar Comentarios
Ver/Esconder Leyenda

Con este simple video aprenderás a entrenar a tu cachorro al momento de comer → bit.ly http://t.co/EVdIVAT6wH

¿Quieres cambiar los muebles y la decoración? Pasa por @surdeco_cl y sorpréndete. San Pedro 210 #PuertoVaras t.co

"Tu Gobierno en Terreno" se realizó en el sector de Balmaceda Norte en Copiapó | NOSTÁLGICA www.nostalgica.cl

Tras poco más de 4 meses Hans Podlipnik recupera el número 1 de Chile en ranking ATP bit.ly

El fixture de U. de Chile, Colo Colo y S. Wanderers, los aspirantes al título del Apertura: Revisa los comprom... t.co

Manchas solares gigantescas causan alerta en los científicos. Los detalles → bit.ly http://t.co/WKMZWselk

@ubbchile Mar 28 Oct / 10:00 hrs, participa en 6ta Feria Internacional UBB bit.ly

Comuneros se toman fundo perteneciente a la Iglesia Católica en La Araucanía www.ahoranoticias.cl http://t.co/BxsSVnR0F5

Filtran 5 millones de cuentas y contraseñas de Gmail: Revisa aquí si fuiste hackeado ahoranoticias.cl http://t.co/KykPV82eSj

07:37 en Talca 10-4-1 RX1 6 Oriente esq. Av. Costanera t.co

Toda la información deportiva de Chile y el Mundo con @Deportes921, en el 103.5 FM de #Linares y online en RadioAncoa.cl

Lactante permanece en riesgo vital tras ser aplastada por televisor en Alto Hospicio t.co vía

UDI, evangélicos y Ong ´s en picada contra "Nicolás tiene dos ... desde

jefferson adaro @adarojefferson hace 37 minutos
Retweet a @ahnoticiasmega
Comuneros se toman fundo perteneciente a la Iglesia Católica en La Araucanía www.ahoranoticias.cl http://t.co/BxsSVnR0F5
Responder Retweet Favorito

...opilla.c ... sobre ... to a

La appstore de Amazon ahora se integra en su app de Android ift.tt

ATENCIÓN: Semáforos apagados Ruta 5 / La Cantera de #Coquimbo. Transite con #precaución / 7:15 hrs. t.co

NASA VISITÓ EL COLEGIO FRANCÉS DE PUNTA ARENAS t.co #puq vía @radiopolar

Impacto por fallecimient... de joven dentista magallánico en accidente de tránsito: bit.ly #fb

El 50% de los colegios con más bajo resultado en Simce es subvencionado ow.ly http://t.co/YSIABr

Skybox for Good: Google cede sus imágenes satelitales a organizaciones sin ánimo de lucro feeds.weblogssl.com

Oiga, @gustavohasbun, acá le dejo otra foto de la "multitudinaria" marcha de la confepa: t.co

El Caso cascadas o cómo el grupo logró quitarle 44 millones de dólares a miles de chilenos a través de sus pensiones t.co

[Hoy en t.co] Estrecho triunfo de Rousseff. El mercado vota hoy y se espera un frío recibimiento http://t.co/OlcAFQ4yeg

Y comienzan a aparecer los primeros rayos de sol en Valparaíso @biobio buen comienzo de semana 07:11 Hrs t.co

Bachelet defendió sus planes de reforma constitucional en revista alemana bit.ly http://t.co/hU6HqxpNU1

Ante tanta desinformación e intentos de manipulación, mi pequeño aporte: Diferencia entre REMUNERACIÓN y LUCRO: t.co

Junto a Presidenta Bachelet y delegación chilena reunidos con fundaciones públicas alemanas vinculadas con Chile. t.co

Equipos de Algarrobo, Villa Alemana, Viña del Mar y Los Andes a Congreso Nacional EXPLORA CONICYT www.ucvradio.cl

Usa estos botones para elegir qué ver:
Todas las Regiones
I
II
III
IV
V
VI
VII
VIII
IX
X
XI
XII
RM
XIV
XV
Por @camby | Nota: no estamos afiliados con Twitter.

[E. Graells-Garrido and M. Lalmas, "Balancing diversity to counter-measure geographical centralization in microblogging platforms", ACM Hypertext'14]

Gender Bias

Systemic bias?

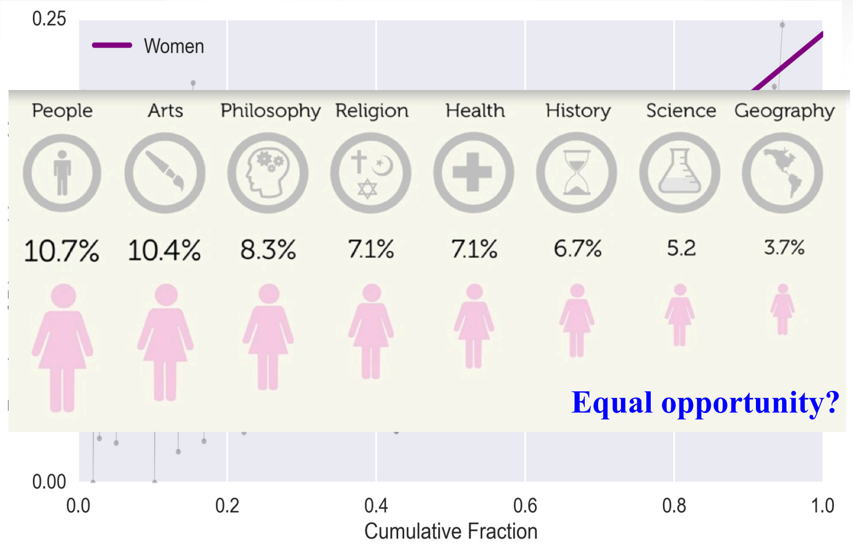

Equal opportunity?

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

Activity Bias: Wisdom of a Few?

- The Web already is influenced by small groups
  - "0.05% of the user population, attract almost 50% of all attention within Twitter" (50K users) [Wu, Hofman, Mason & Watts, WWW 2011]
- We explored this issue further with four different datasets:
  1. a large one from Twitter (2011),
  2. a small one from Facebook (2009),
  3. Amazon reviews (2013), and
  4. Wikipedia editors (2015).
- Digital desert: the content that is never seen

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

Examples

How many users produce most of the content?

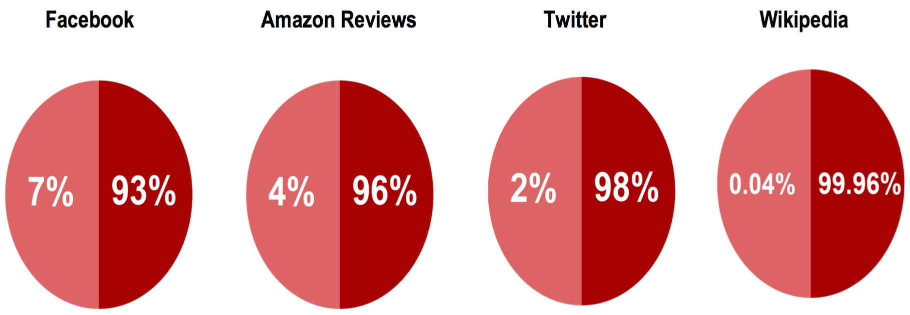

[Baeza-Yates & Saez-Trumper, ACM Hypertext 2015]

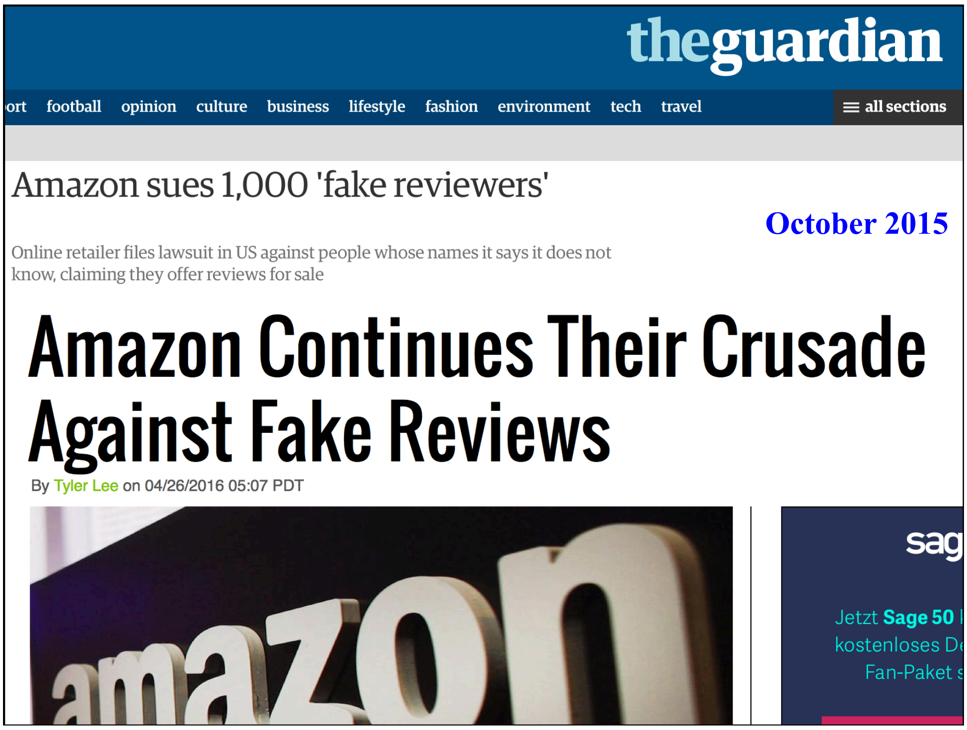
theguardian
ort football opinion culture business lifestyle fashion environment tech travel
all sections
sag
Jetzt Sage 50
kostenloses D
Fan-Paket s

Amazon sues 1,000 'fake reviewers'

October 2015

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

Amazon Continues Their Crusade Against Fake Reviews

By Tyler Lee on 04/26/2016 05:07 PDT

Quality of Content?

- Adding content implies adding wisdom?
- We use Amazon's reviews helpfulness
- We computed the text entropy
- Content-based-wise users
- How many of those users are being paid?

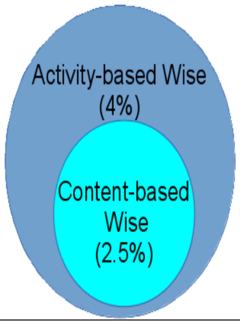

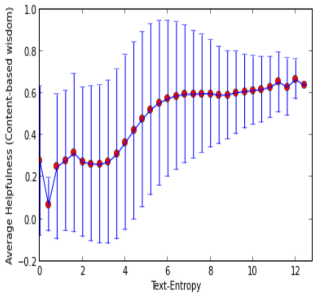

Digital Desert

- 1.1% of the Twitter content is never seen.*
- 31% of articles added/edited in May 2014 in wikipedia, were not visited in June.

Bias in the Interface

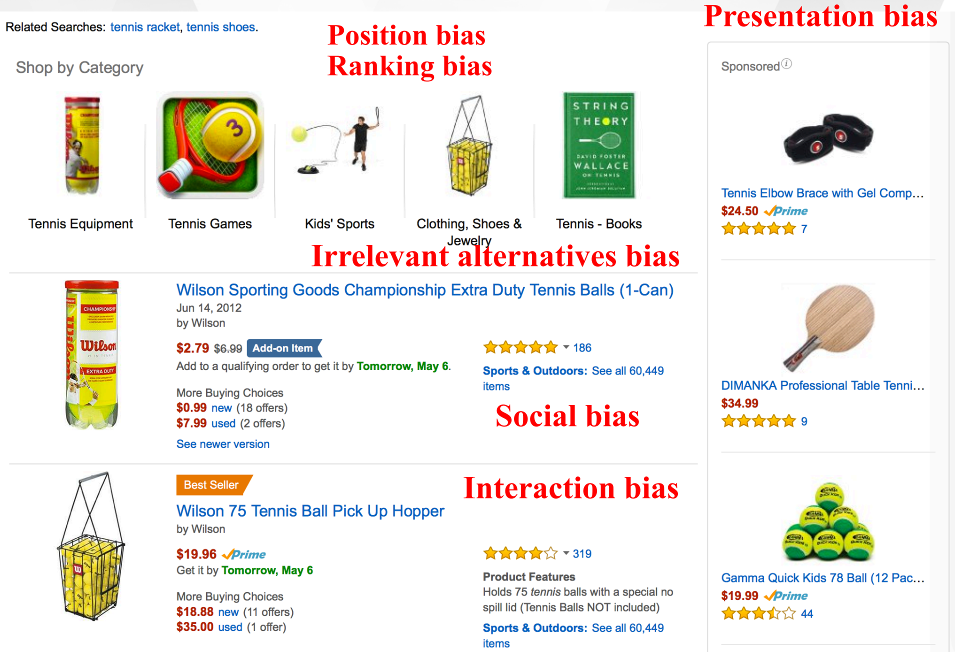

Position bias
Ranking bias

Presentation bias

Irrelevant alternatives bias

Social bias

Interaction bias

Related Searches: tennis racket, tennis shoes.

Shop by Category

Tennis Equipment | Tennis Games | Kids' Sports | Clothing, Shoes & Jewelry | Tennis - Books

Wilson Sporting Goods Championship Extra Duty Tennis Balls (1-Can)
Jun 14, 2012
by Wilson
$2.79 $6.99 Add-on Item
Add to a qualifying order to get it by Tomorrow, May 6.
More Buying Choices
$0.99 new (18 offers)
$7.99 used (2 offers)
See newer version

★★★★★ 186
Sports & Outdoors: See all 60,449 items

Best Seller
Wilson 75 Tennis Ball Pick Up Hopper
by Wilson
$19.96 ✓Prime
Get it by Tomorrow, May 6
More Buying Choices
$18.88 new (11 offers)
$35.00 used (1 offer)

★★★★☆ 319
Product Features
Holds 75 *tennis* balls with a special no spill lid (Tennis Balls NOT included)
Sports & Outdoors: See all 60,449 items

Sponsored

Tennis Elbow Brace with Gel Comp…
$24.50 ✓Prime
★★★★★ 7

DIMANKA Professional Table Tenni…
$34.99
★★★★★ 9

Gamma Quick Kids 78 Ball (12 Pac…
$19.99 ✓Prime
★★★½☆ 44

Presentation Bias

- Interaction data will be biased to what is shown
- In recommender systems, items recommended will get more clicks than items not recommended
- In search systems top ranked results will get more clicks than other results
  - Ranking bias
  - Interaction bias

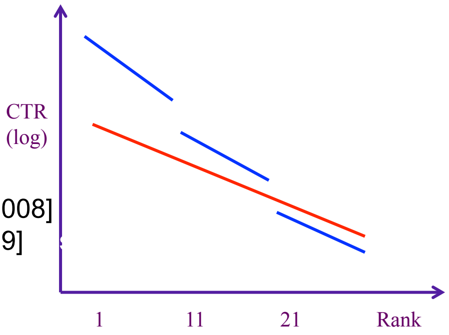

[Dupret & Piwowarski, SIGIR 2008]
[Chapelle & Zhang, WWW 2009]

Social Bias

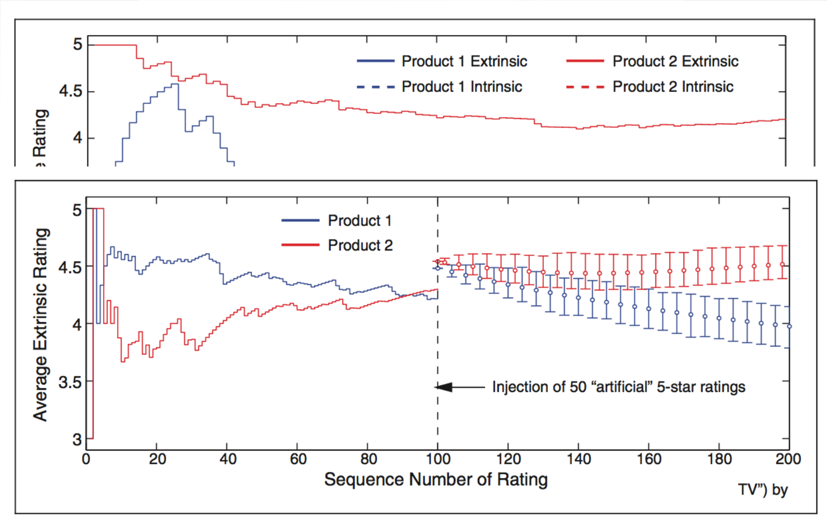

**[WHY AMAZON'S RATINGS MIGHT MISLEAD YOU; The Story of Herding Effects
Ting Wang and Dashun Wang, Big Data, 2014]**

Independence of Irrelevant Alternatives

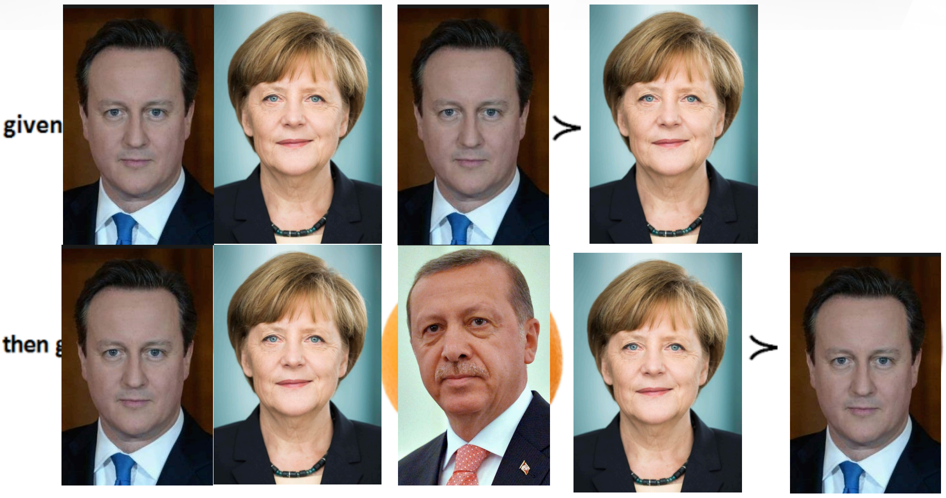

Irrelevant Alternatives (IIA)

- IIA does not always hold
- If that is the case, nested logit should be used instead of multinomial logit
- Optimal quadratic algorithm to recover trees for a nested decision process [Benson et al, WWW 2016]
- Statistical tests to check if a nested model works (95%)

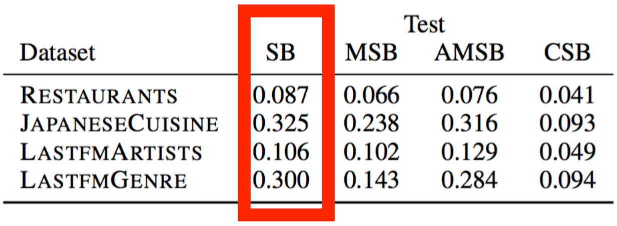

	Test			
Dataset	SB	MSB	AMSB	CSB
RESTAURANTS	0.087	0.066	0.076	0.041
JAPANESECUISINE	0.325	0.238	0.316	0.093
LASTFMARTISTS	0.106	0.102	0.129	0.049
LASTFMGENRE	0.300	0.143	0.284	0.094

- How much can be explained by bias?

Extreme Algorithmic Bias

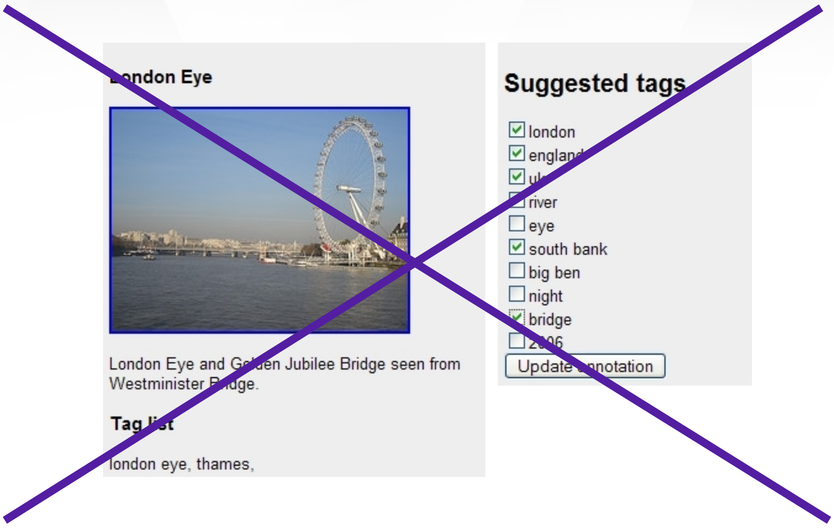

London Eye

London Eye and Golden Jubilee Bridge seen from Westminister Bridge.

Tag list

london eye, thames,

Suggested tags

☑ london
☑ england
☑ uk
□ river
□ eye
☑ south bank
□ big ben
□ night
☑ bridge
□ 2006

Second Order Bias in Web Content

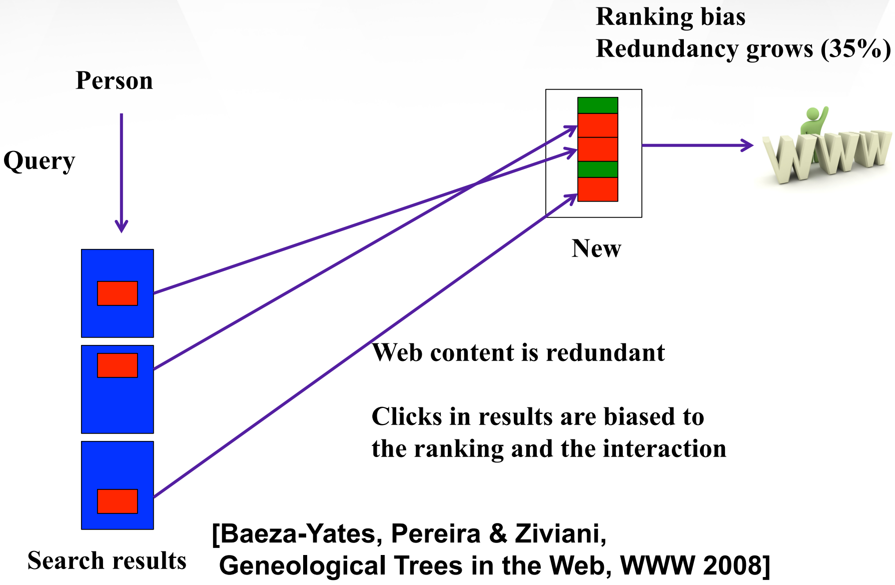

Ranking bias
Redundancy grows (35%)

Person

Query

New

Web content is redundant

Clicks in results are biased to the ranking and the interaction

Search results

[Baeza-Yates, Pereira & Ziviani, Geneological Trees in the Web, WWW 2008]

The Long Tail: Sparsity

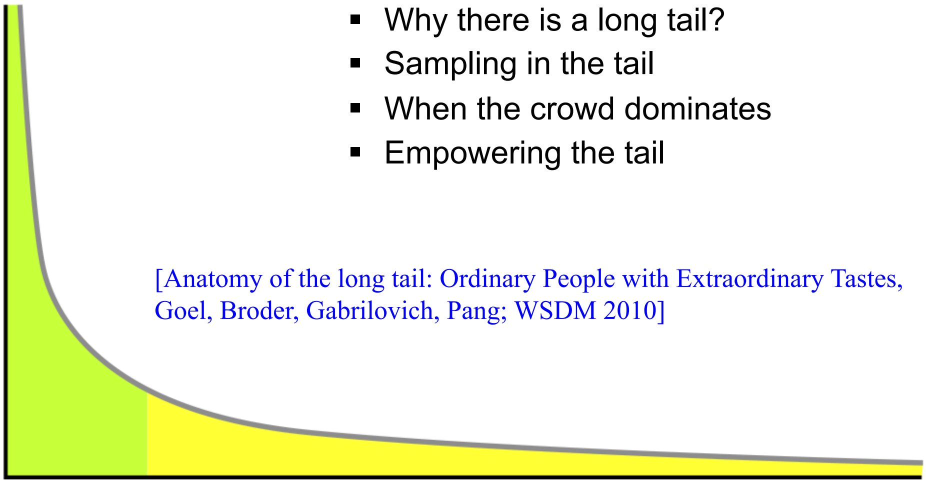

- Why there is a long tail?
- Sampling in the tail
- When the crowd dominates
- Empowering the tail

[Anatomy of the long tail: Ordinary People with Extraordinary Tastes, Goel, Broder, Gabrilovich, Pang; WSDM 2010]

Most measures in the Web follow a power law

Sample Size?

- If we want to estimate the frequency of queries that appear with probability at least *p* with a certain relative error ϵ we can use the standard binomial error formula √(1-p)/*np* which works well for *p* near ½
- Better is the Agresti-Coull technique (also called *take 2*) which gives:

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

- 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 90% and ϵ is 10%, we get *n* = 2342. The standard formula gives *n* = 900.

[Baeza-Yates, SIGIR 2015, Industry track]

Incremental Sampling

- Main goal: make good samples consistent across time
- Simple idea based in stratified sampling: bins + random start point

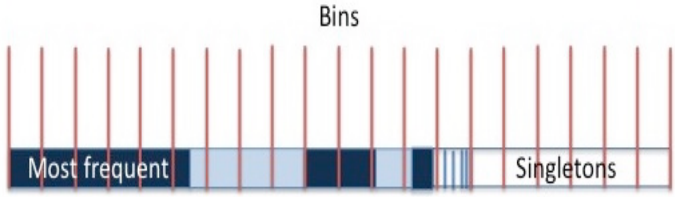

- Bin size can be found by binary search starting with a good approximation if a query frequency model is used (*b* < *V/n*)
- This perfectly mimics the head of the distribution, but not the tail
- Change the bins in the tail to get the right distribution

[Baeza-Yates, SIGIR 2015, Industry track]

Fixing the Tail

- To mimic the tail we change the binning size when we reach a query frequency of *b/*2
- If we want a singleton ratio of $\beta = S/V$ we recalculate the binning size as

$$b' = (1 - \beta)(Q - Q')/(\beta V')$$

- where Q' and V' are the partial vocabulary size and volume before changing the bin size.

[Baeza-Yates, SIGIR 2015, Industry track]

Stratified Sampling Example

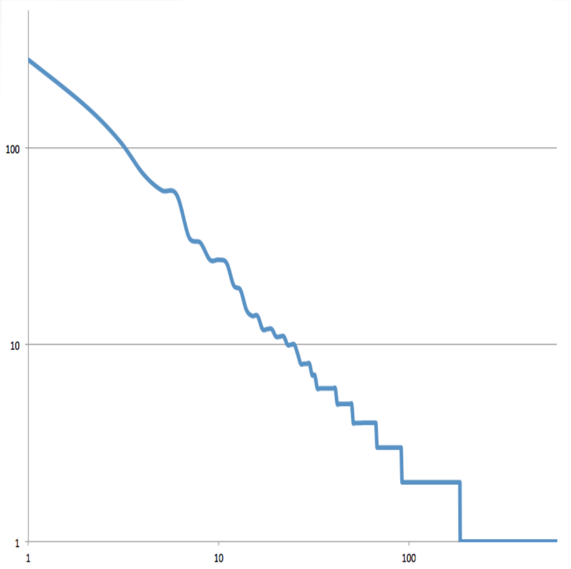

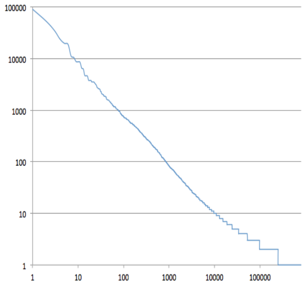

When the Crowd Dominates

Kills the long tail

Personalization "facets":

- Language (not always)
- Location
- Semantic facets per user
  - Query intent prediction in search

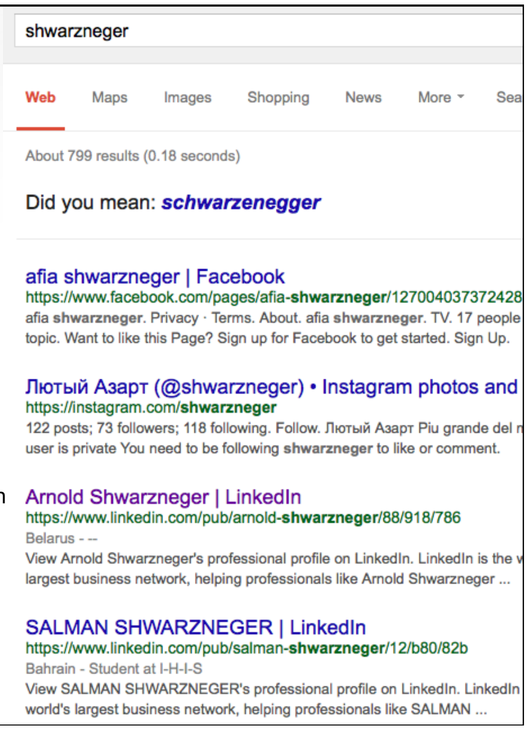

About 799 results (0.18 seconds)

Did you mean: ***schwarzenegger***

afia shwarzneger | Facebook
https://www.facebook.com/pages/afia-**shwarzneger**/127004037372428
afia **shwarzneger**. Privacy · Terms. About. afia **shwarzneger**. TV. 17 people
topic. Want to like this Page? Sign up for Facebook to get started. Sign Up.

Лютый Азарт (@shwarzneger) • Instagram photos and
https://instagram.com/**shwarzneger**
122 posts; 73 followers; 118 following. Follow. Лютый Азарт Piu grande del r
user is private You need to be following **shwarzneger** to like or comment.

Arnold Shwarzneger | LinkedIn
https://www.linkedin.com/pub/arnold-**shwarzneger**/88/918/786
Belarus - --
View Arnold Shwarzneger's professional profile on LinkedIn. LinkedIn is the w
largest business network, helping professionals like Arnold Shwarzneger ...

SALMAN SHWARZNEGER | LinkedIn
https://www.linkedin.com/pub/salman-**shwarzneger**/12/b80/82b
Bahrain - Student at I-H-I-S
View SALMAN SHWARZNEGER's professional profile on LinkedIn. LinkedIn
world's largest business network, helping professionals like SALMAN ...

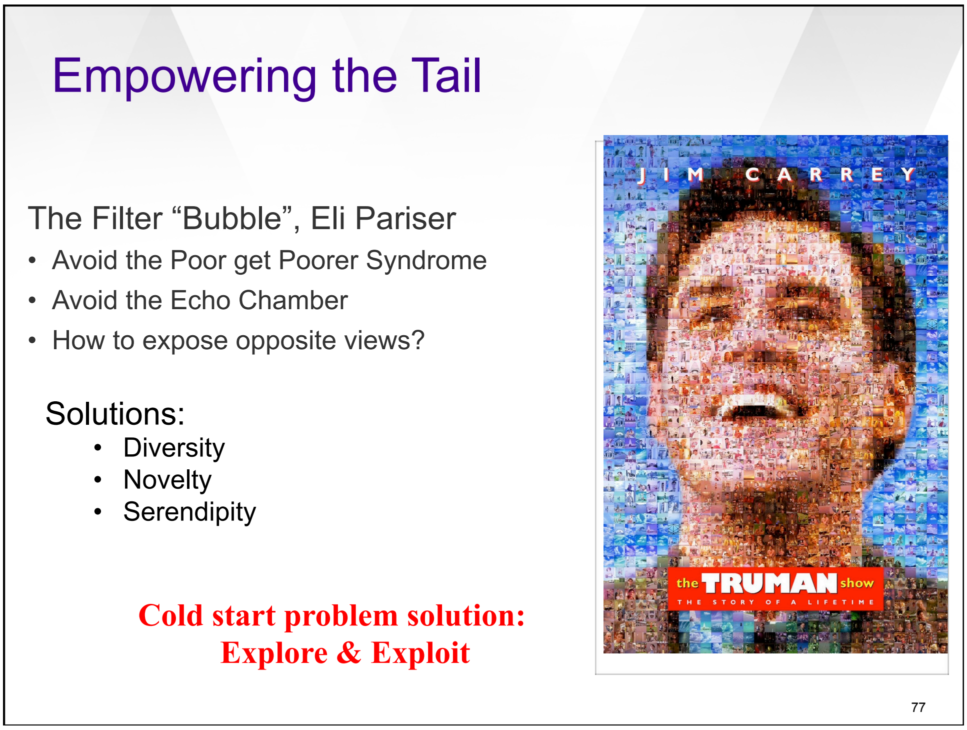

Empowering the Tail

The Filter “Bubble”, Eli Pariser

- Avoid the Poor get Poorer Syndrome
- Avoid the Echo Chamber
- How to expose opposite views?

Solutions:

- Diversity
- Novelty
- Serendipity

Cold start problem solution:
Explore & Exploit

JIM CARREY
the TRUMAN show
THE STORY OF A LIFETIME

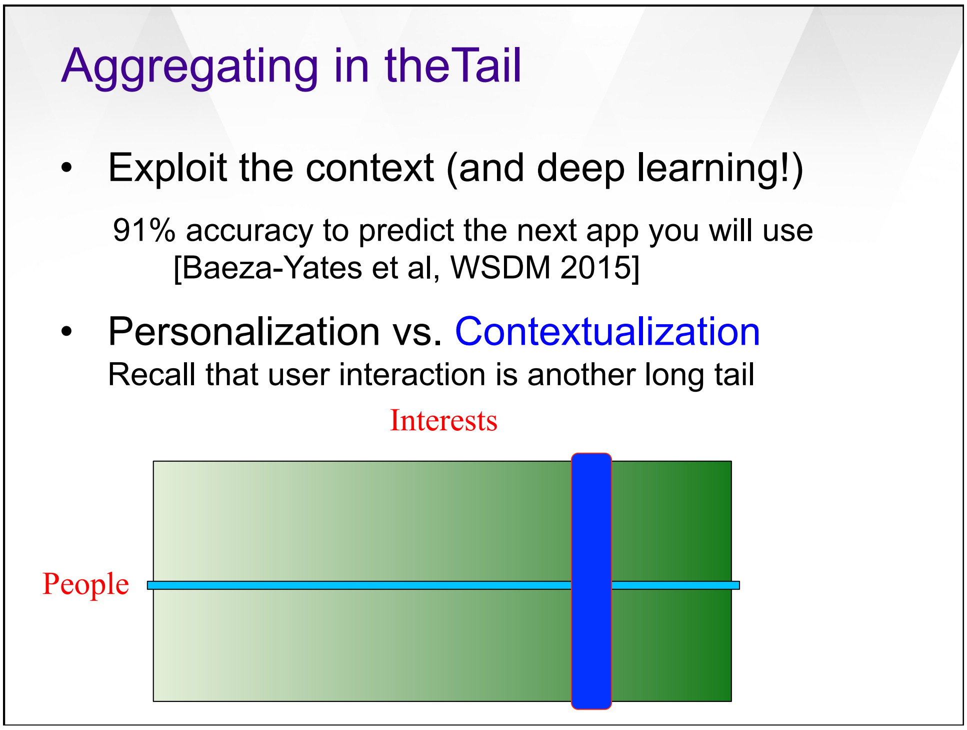

Aggregating in theTail

- 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
  Recall that user interaction is another long tail

Interests
People

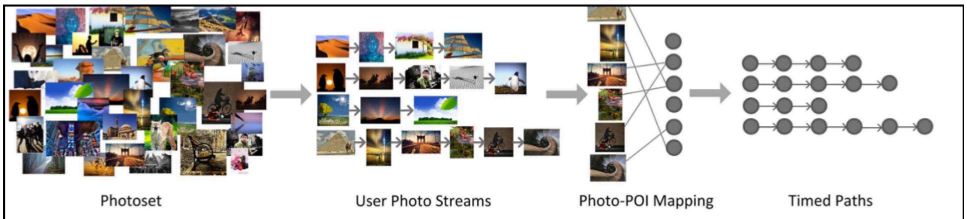

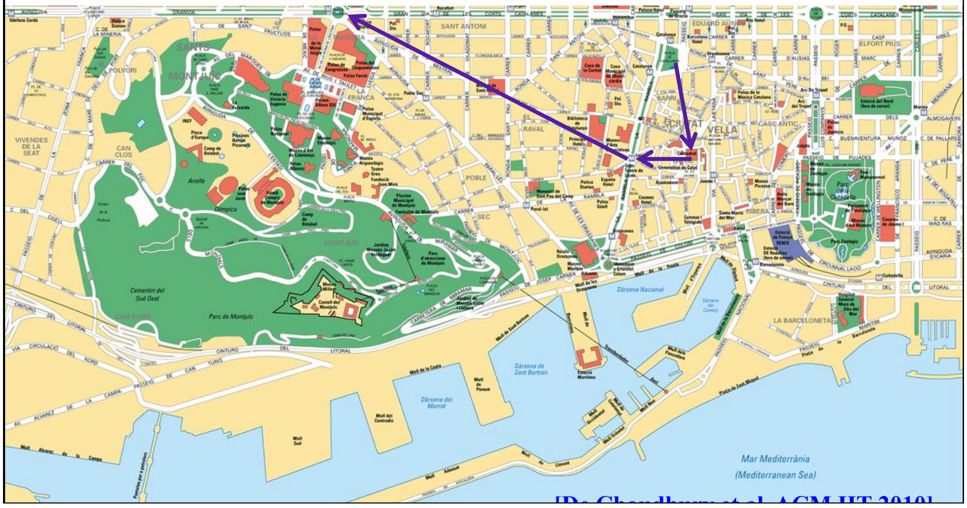

[De Choudhury et al, ACM HT 2010]

Crowdsourcing Data: Good Paths

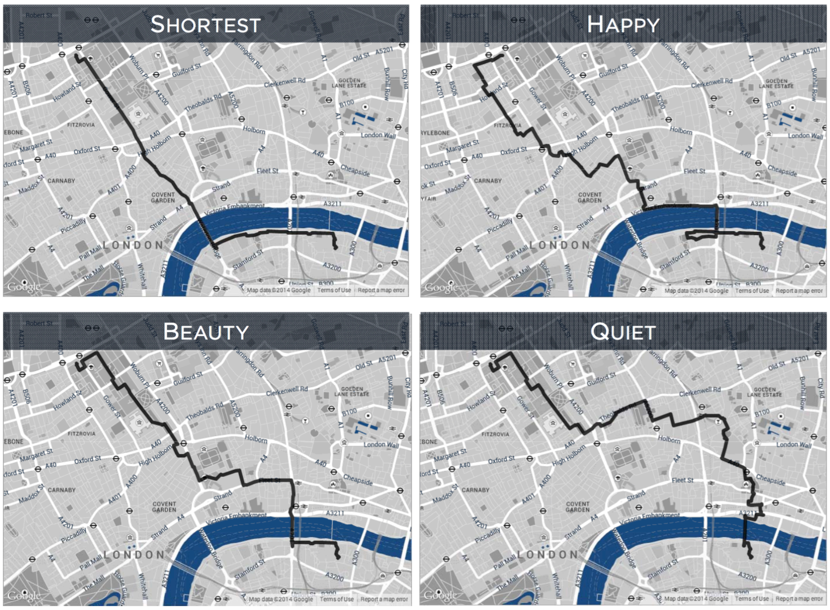

[Quercia et al, ACM HT 2014]

Regions from Pictures

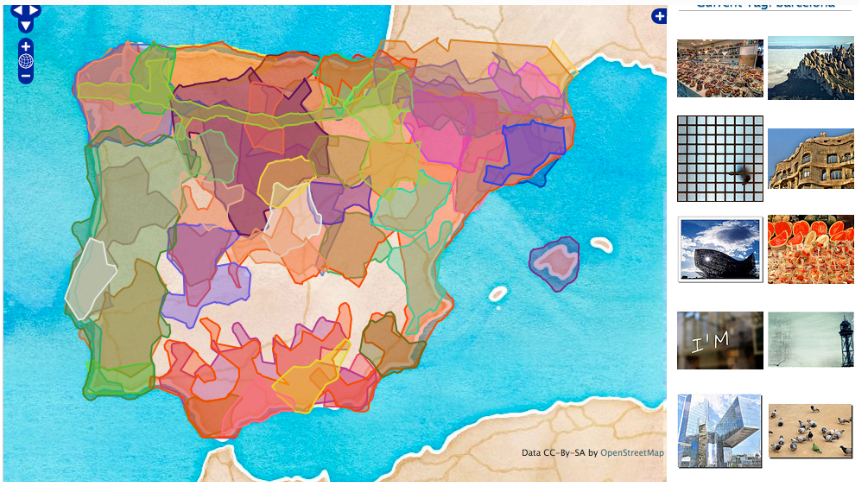

[Thomee et al, Demo at CHI 2014]

Privacy 101:
AOL Query Logs Release Incident

A Face Is Exposed for AOL Searcher No. 4417749,
By MICHAEL BARBARO and TOM ZELLER Jr,
The New York Times, Aug 9 2006

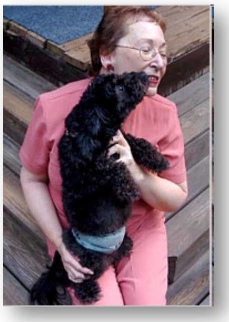

- No. 4417749 conducted hundreds of searches over a three-month period on topics ranging from “numb fingers” to “60 single men”.
- Other queries: “landscapers in Lilburn, Ga,” several people with the last name Arnold and “homes sold in shadow lake subdivision gwinnett county georgia.”
- Data trail led to Thelma Arnold, a 62-year-old widow who lives in Lilburn, Ga., frequently researches her friends’ medical ailments and loves her three dogs.

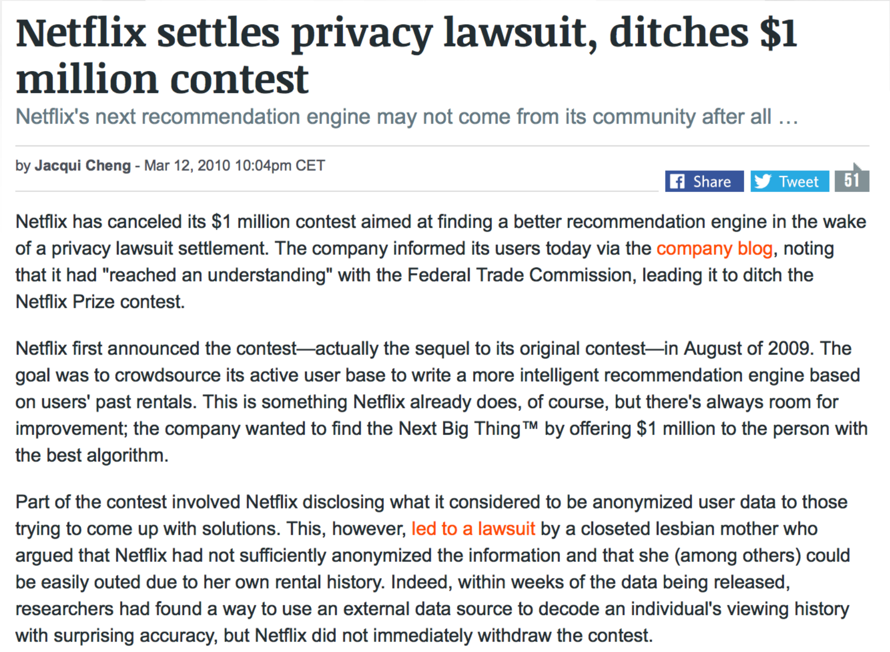

Netflix settles privacy lawsuit, ditches $1 million contest

Netflix's next recommendation engine may not come from its community after all …

by **Jacqui Cheng** - Mar 12, 2010 10:04pm CET

Netflix has canceled its $1 million contest aimed at finding a better recommendation engine in the wake of a privacy lawsuit settlement. The company informed its users today via the company blog, noting that it had "reached an understanding" with the Federal Trade Commission, leading it to ditch the Netflix Prize contest.

Netflix first announced the contest—actually the sequel to its original contest—in August of 2009. The goal was to crowdsource its active user base to write a more intelligent recommendation engine based on users' past rentals. This is something Netflix already does, of course, but there's always room for improvement; the company wanted to find the Next Big Thing™ by offering $1 million to the person with the best algorithm.

Part of the contest involved Netflix disclosing what it considered to be anonymized user data to those trying to come up with solutions. This, however, led to a lawsuit by a closeted lesbian mother who argued that Netflix had not sufficiently anonymized the information and that she (among others) could be easily outed due to her own rental history. Indeed, within weeks of the data being released, researchers had found a way to use an external data source to decode an individual's viewing history with surprising accuracy, but Netflix did not immediately withdraw the contest.

Risks of Privacy in Query Logs

- Profile [Jones, Kumar, Pang, Tompkins, CIKM 2007]
  - Gender: 84%
  - Age (±10): 79%
  - Location (ZIP3): 35%
- Vanity Queries [Jones et al, CIKM 2008]
  - Partial name: 8.9%
  - Complete: 1.2%
- More information:
  - A Survey of query log privacy-enhancing techniques from a policy perspective [Cooper, ACM TWEB 2008]
- A good anonymization technique is still an open problem

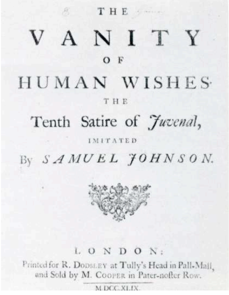

THE
VANITY
OF
HUMAN WISHES
THE
Tenth Satire of *Juvenal*,
IMITATED
By *SAMUEL JOHNSON.*

LONDON:
Printed for R. Dodsley at Tully's Head in Pall-Mall,
and Sold by M. Cooper in Pater-noſter Row.
M.DCC.XLIX.

Privacy Awareness

- How our privacy changes when we change our social network?
- Information gain to predict a private attribute based on public data
- Each user may have a promiscuity score

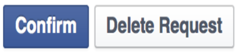

- Example: new friendship request

  Promiscuity(me) > Promiscuity(new) 🟢

  Promiscuity(me) ≥ Promiscuity(new) + max-gain-I-allow 🟡

  Promiscuity(me) < Promiscuity(new) + max-gain-I-allow 🔴

Related work by [Estivill-Castro & Nettleton; Singh, ASONAM 2015]

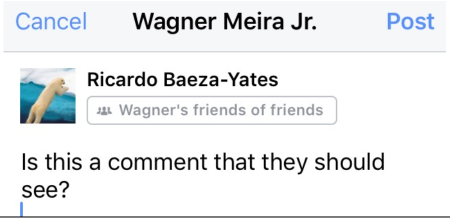
Cancel Wagner Meira Jr. Post

Ricardo Baeza-Yates

Wagner's friends of friends

Is this a comment that they should see?

The Web Works Thanks to Bias!

- Web traffic
  - Local caching
  - Proxy/Akamai caching
- Search engines
  - Answer caching
  - Essential web pages
    - 25% queries can be answered with less than 1% of the URLs [BY,Boldi, Chierichetti, WWW 2015]
- E-Commerce
  - Most revenue comes from few items

Activity bias

(Self) selection bias

Web Data

- A mirror of ourselves, the good, the bad and the ugly
- The web amplifies everything, good or bad, but always leaves traces
- We have to be aware of the biases and contrarrest them
- We have to be aware of our privacy

Big Data of **People** is huge…..
….. but is tiny compared to the future
Big Data of the **Internet of Things (IoT)**

It's Hard to Get Data to Tell the Truth

- The blindness of the averages
  - Look at distributions
- Absolute vs. relative
  - Income per capita vs. Inequality
- Local vs. global optimization
  - Teams competing without knowing, uncorrelated criteria
- You can always see/torture data as you wish
  - 61 analysts, 29 teams: 20 yes and 9 no (Univ. of Virginia, COS)

Same Data, Different Conclusions

Twenty-nine research teams were given the same set of soccer data and asked to determine if referees are more likely to give red cards to dark-skinned players. Each team used a different statistical method, and each found a different relationship between skin color and red cards.

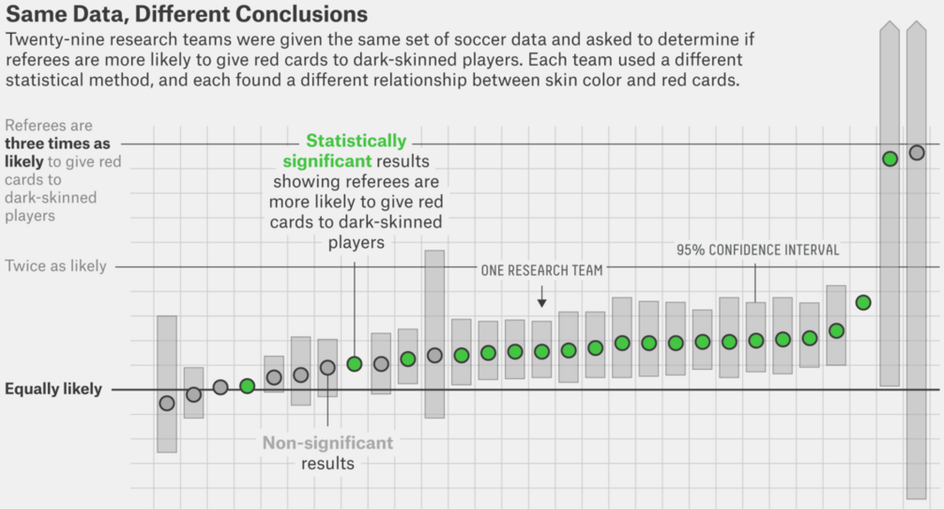

FIVETHIRTYEIGHT

SOURCE: BRIAN NOSEK ET AL.

ASIST 2012 Book of the Year Award

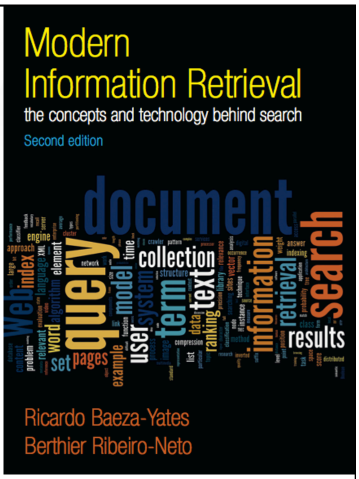
Modern Information Retrieval
the concepts and technology behind search
Second edition
Ricardo Baeza-Yates
Berthier Ribeiro-Neto

Questions?

Contact: rbaeza@acm.org

www.baeza.cl

@polarbearby

Biased Questions?