computational social media

lecture 4: shooting

part 1

daniel gatica-perez

idiap RESEARCH INSTITUTE

EPFL

announcements

assignment #3 will be given today

team project progress presentations: 12:00-13:00

5-minute talks + 3 minutes for Q&A

this lecture

1. a snapshot of the present
    flickr, instagram, snapchat
2. a look into the past
    20th century image sharing practices
3. understanding research on social image systems
    flickr: tags & communities
    computer vision as an enabler
    instagram: selfies & engagement
    snapchat: ephemeral social media

1.
a snapshot of the present

reminder: participation in social media

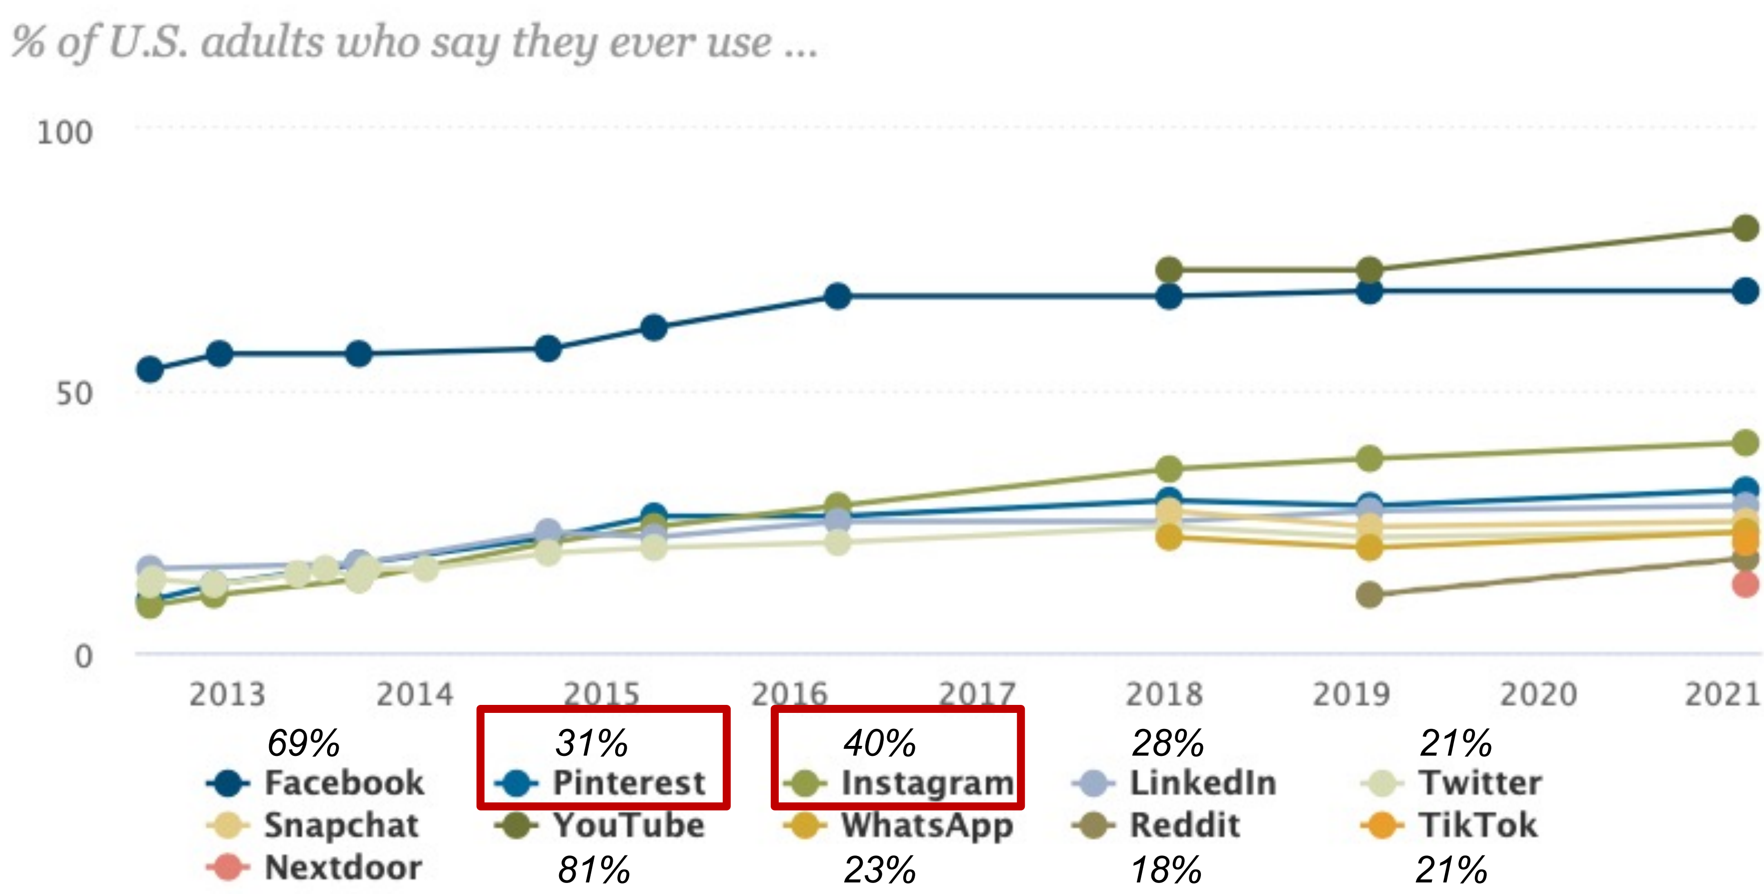

Note: Respondents who did not give an answer are not shown.

Source: Surveys of U.S. adults conducted 2012-2021.

https://www.pewresearch.org/internet/fact-sheet/social-media/

flickr

1. “help people make their photos available to the people who matter to them
2. enable new ways of organizing photos and video”

02.2004	launched
03.2005	bought by Yahoo
08.2011	hosting over 6 billion images
03.2013	87 million users
	3.5 million images per day
04.2018	bought by SmugMug

https://www.flickr.com/about/
http://en.wikipedia.org/wiki/Flickr

instagram

“capture and share the world's moment”

10.2010 launched

04.2012 bought by facebook

	2015	2016	2018	2019
monthly active users	300M	400M	800M	1B
users outside US	70%	75%	n/a	n/a
shared photos	30B	40B	n/a	n/a
uploaded photos per day	70M	80M	n/a	n/a
daily stories	n/a	n/a	300M	500M
likes per day	2.5B	3.5B	n/a	n/a

http://instagram.com/press/ accessed: 04.2016
https://instagram-press.com/our-story/, accessed 04.2019

snapchat

"We are not the sum of everything we have said or done or experienced or published – we are the *result*. We are who we are today, right now."

09.2011	launched
04.2013	80% of users are in US
11.2013	70% users are women
	400 million "snaps" received per day
	declined offers from facebook & google
08.2014	100 million monthly active users

http:// en.wikipedia.org/wiki/Snapchat
http://blogs.wsj.com/digits/2013/11/20/snapchat-ceo-says-70-of-users-are-women/

this lecture

1. a snapshot of the present

    flickr, instagram, snapchat

2. a look into the past

    20^{th} century image sharing practices

3. understanding research on social image systems

    flickr: tags & communities

    computer vision as an enabler

    instagram: selfies & engagement

    snapchat: ephemeral social media

2.
a look into the past

“Everything exists to end up in a book”

Stéphane Mallarmé, 1842-1898

“Everything exists to end up in a photograph”

Susan Sontag, *On Photography*, 1977

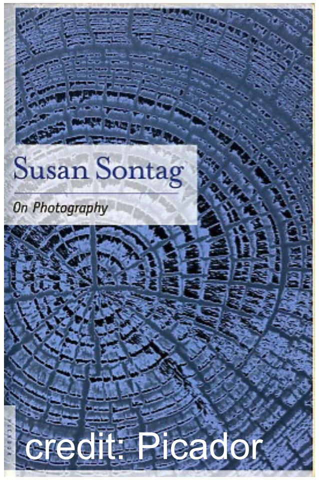

social uses of personal photos (Nancy Van House)

1. memory & identity
2. maintaining relationships
3. self-representation
4. self-expression

N. A. Van House. Flickr and public image-sharing: distant closeness and photo exhibition. In Proc. ACM CHI Extended Abstracts, 2007

credit (cc): Prince Gladson @ flickr
https://www.flickr.com/photos/princegladson

kodak moments (before digital):

few people take **few** pictures

share photos with **family & close friends** as single pics or albums

kodak moments (after digital):

more people take **more** pictures

…but share photos **in the same way as before**

Ofoto (1999) bought by Kodak in 2001

Kodak EasyShare Gallery (2005)

60 million users, billions of images (2008)

2012: Kodak goes bankrupt
Shutterfly buys Kodak Gallery

participatory culture and photo communities
(Henry Jenkins)

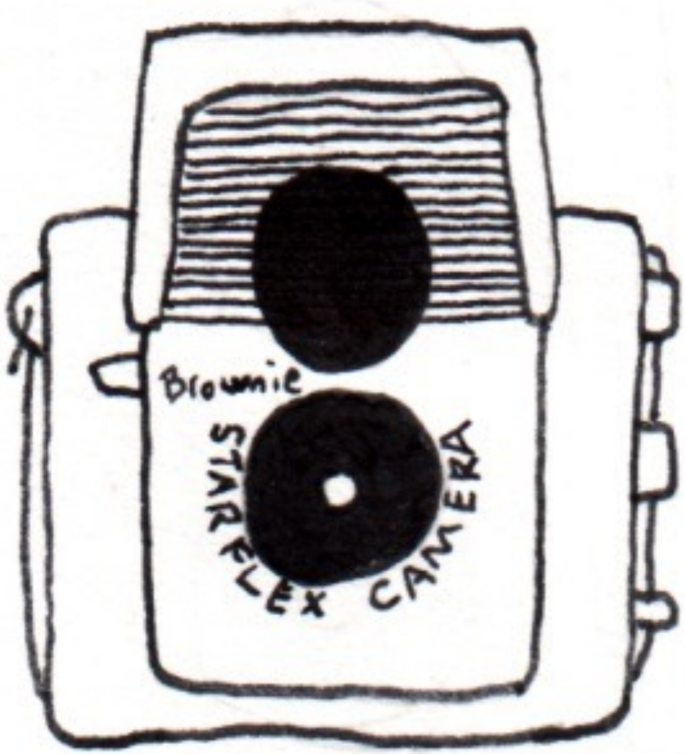

credit (cc): Bradley Stemke @ flickr
https://www.flickr.com/photos/detroitsunrise

a culture in which “fans and other consumers are invited to actively participate in the creation and circulation of new content”

photo communities have existed for as long as photos and cameras have been available

fan clubs, sport fans, birdwatchers, amateur photographers

web 2.0 empowered communities

H. Jenkins, Convergence Culture: Wher Old and New Media Collide, New York University Press, 2006.

stock photography

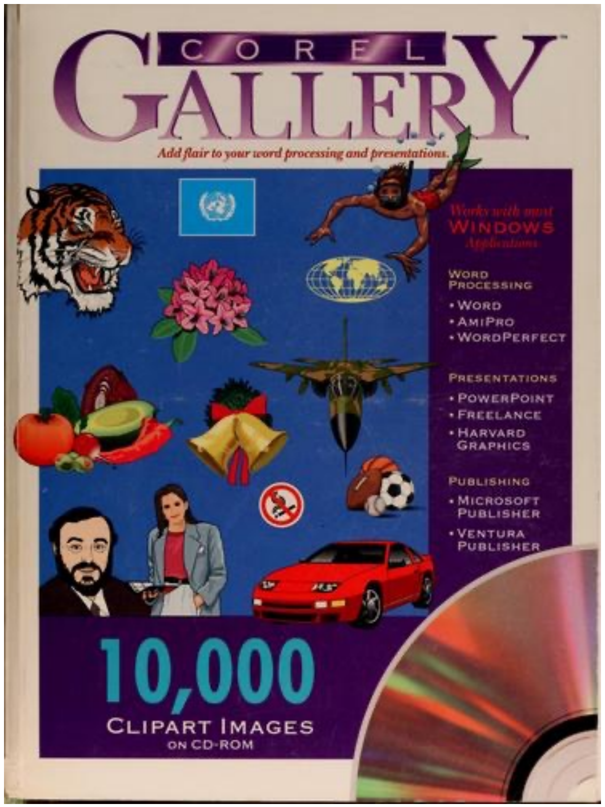

Corel Gallery (1st ed. 1994)

PhotoDisc (2000), now gettyimages

https://openlibrary.org/works/OL16047006W/Corel_gallery

http://en.wikipedia.org/wiki/Getty_images

this lecture

1. a snapshot of the present

    flickr, instagram, snapchat

2. a look into the past

    20th century image sharing practices

3. understanding research on social image systems

    flickr: tags & communities

    computer vision as an enabler

    instagram: selfies & engagement

    snapchat: ephemeral social media

3.
understanding research on social image systems

tags, communities, geo-tagged media

selfies & followers

ephemeral social media

flickr

tags, communities, geo-tags

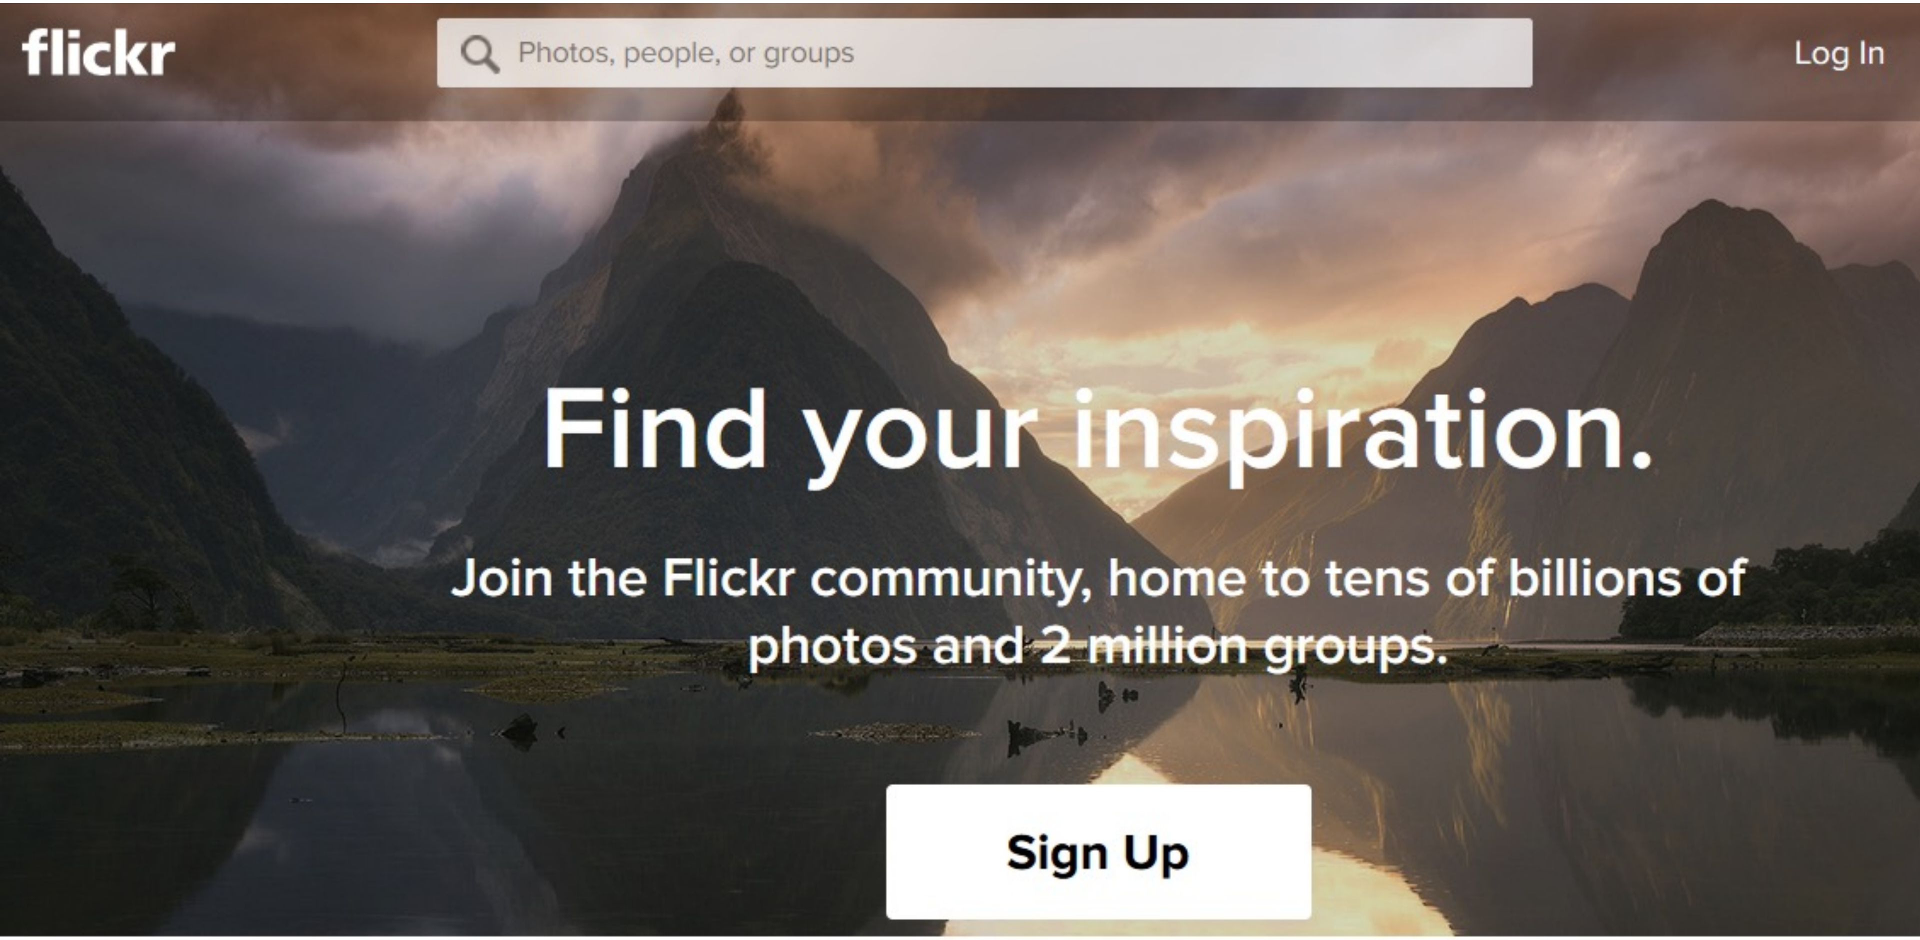

tagging on flickr

flickr tags

https://www.flickr.com/photos/rypscl26/44036709975/in/pool-damndada/

a basic tagging model

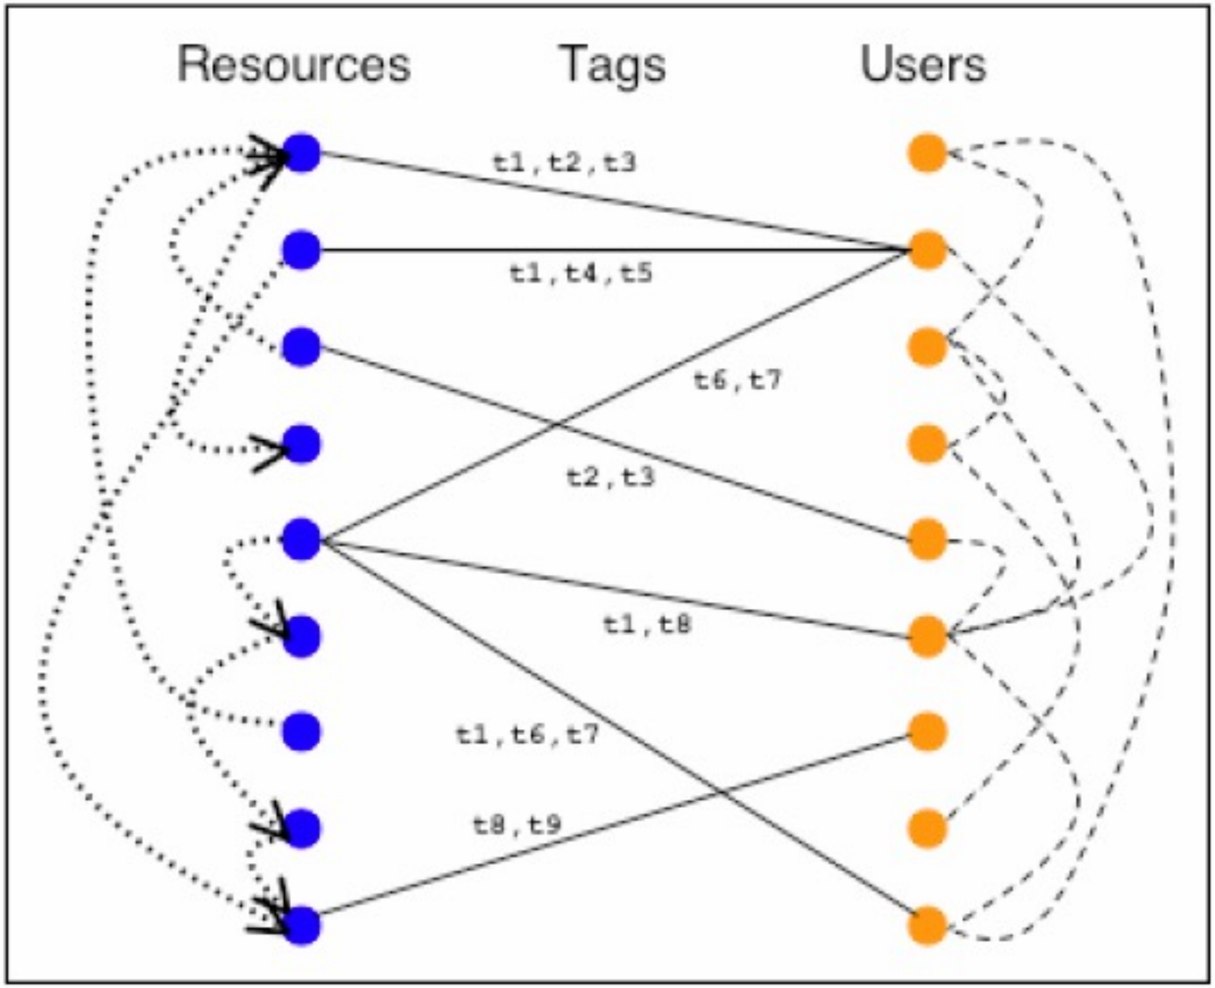

C. Marlow, M. Naaman, D. Boyd, & M. Davis, HT06, tagging paper, taxonomy, Flickr, academic article, to read, in Proc. ACM HT 2006

why do people tag?

		Function	
		Organization	Communication
Sociality	Self	* Retrieval, Directory * Search	* Context for self * Memory
	Social	* Contribution, attention * Ad hoc photo pooling	* Content descriptors * Social Signaling

M. Ames and M. Naaman, Why we tag: motivations for annotation in mobile and online media, in Proc. ACM CHI, 2007

flickr communities

flickr groups

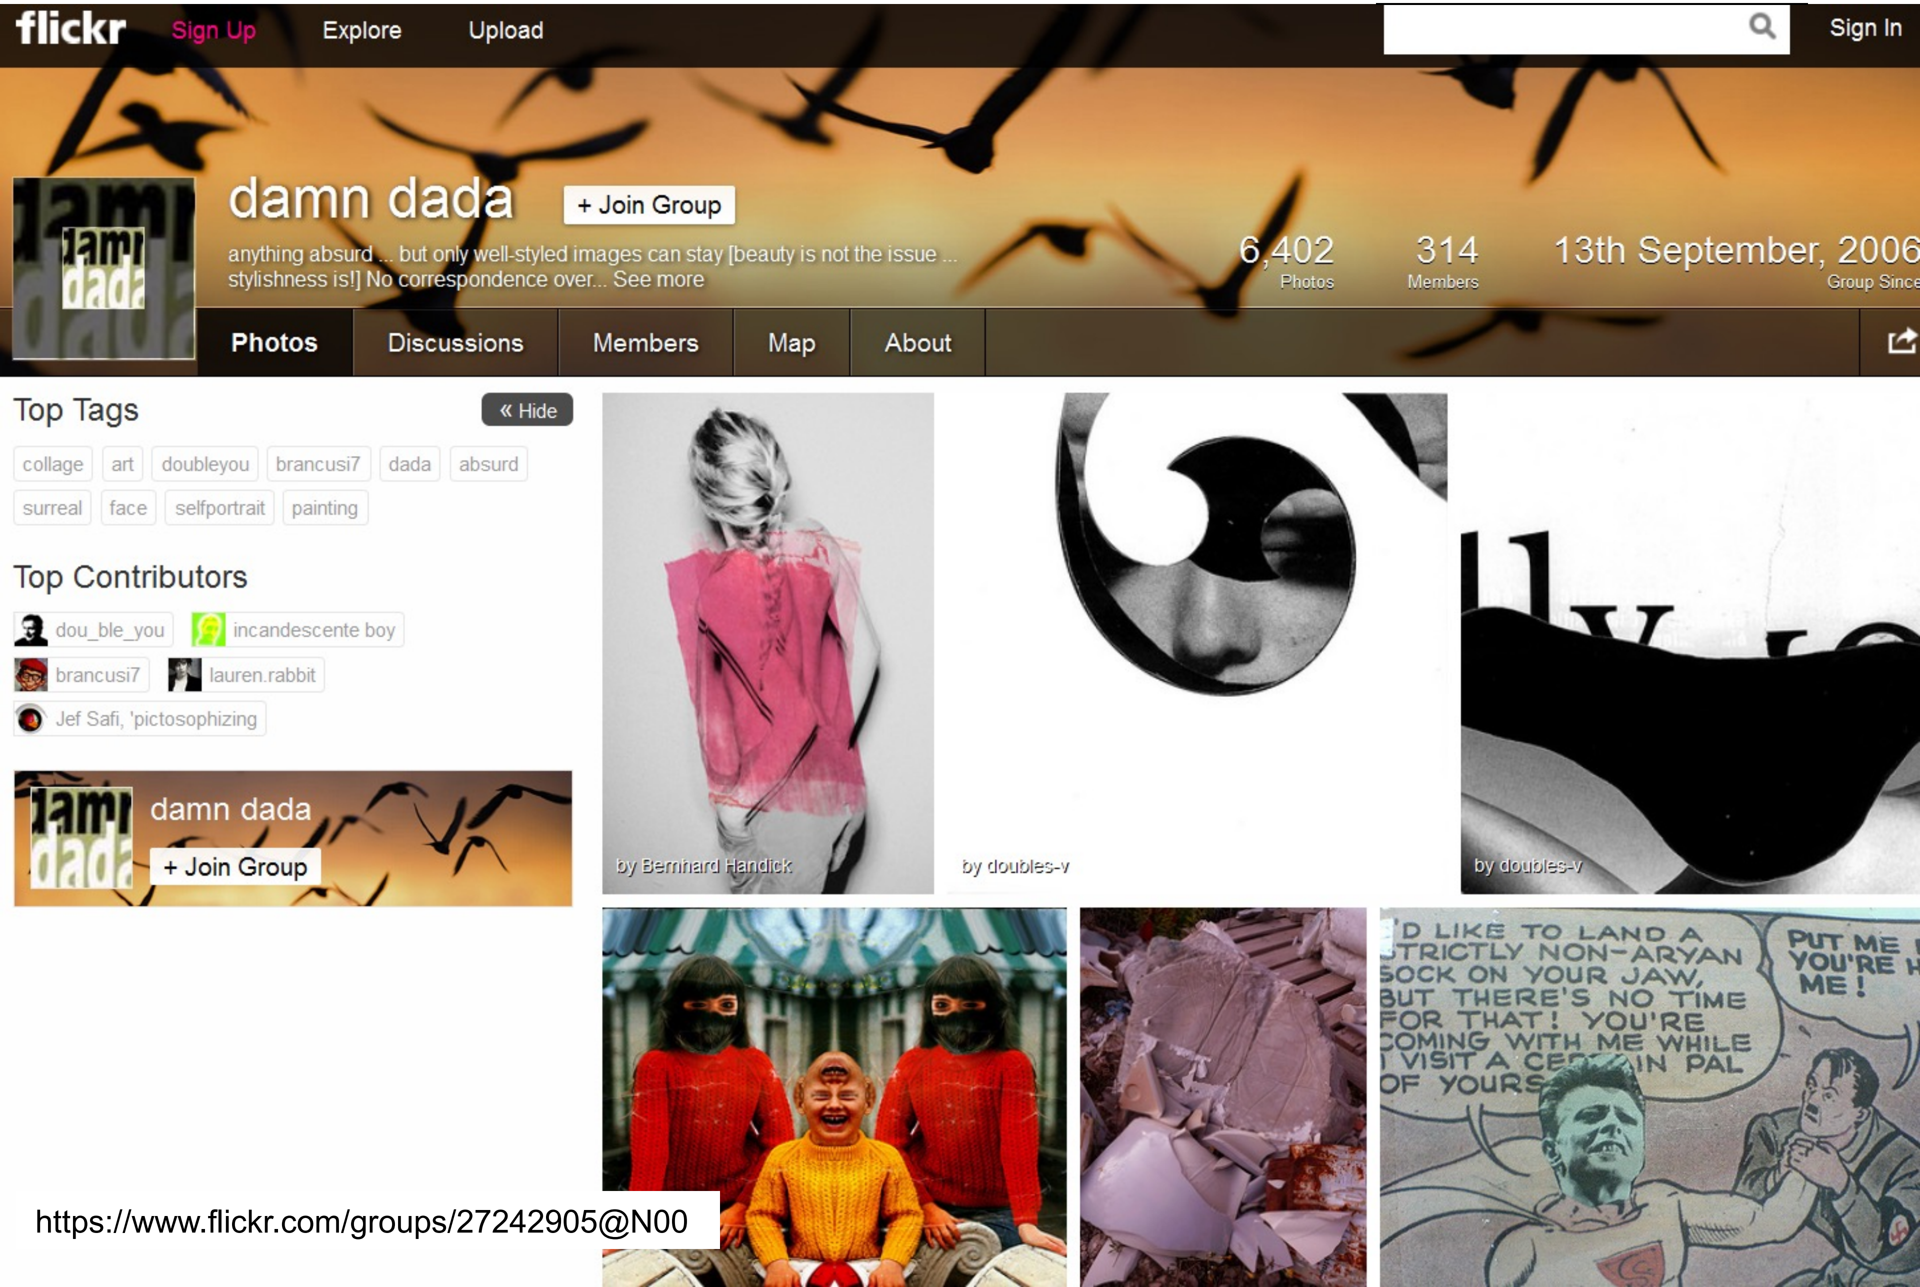

https://www.flickr.com/groups/27242905@N00

group photo pool

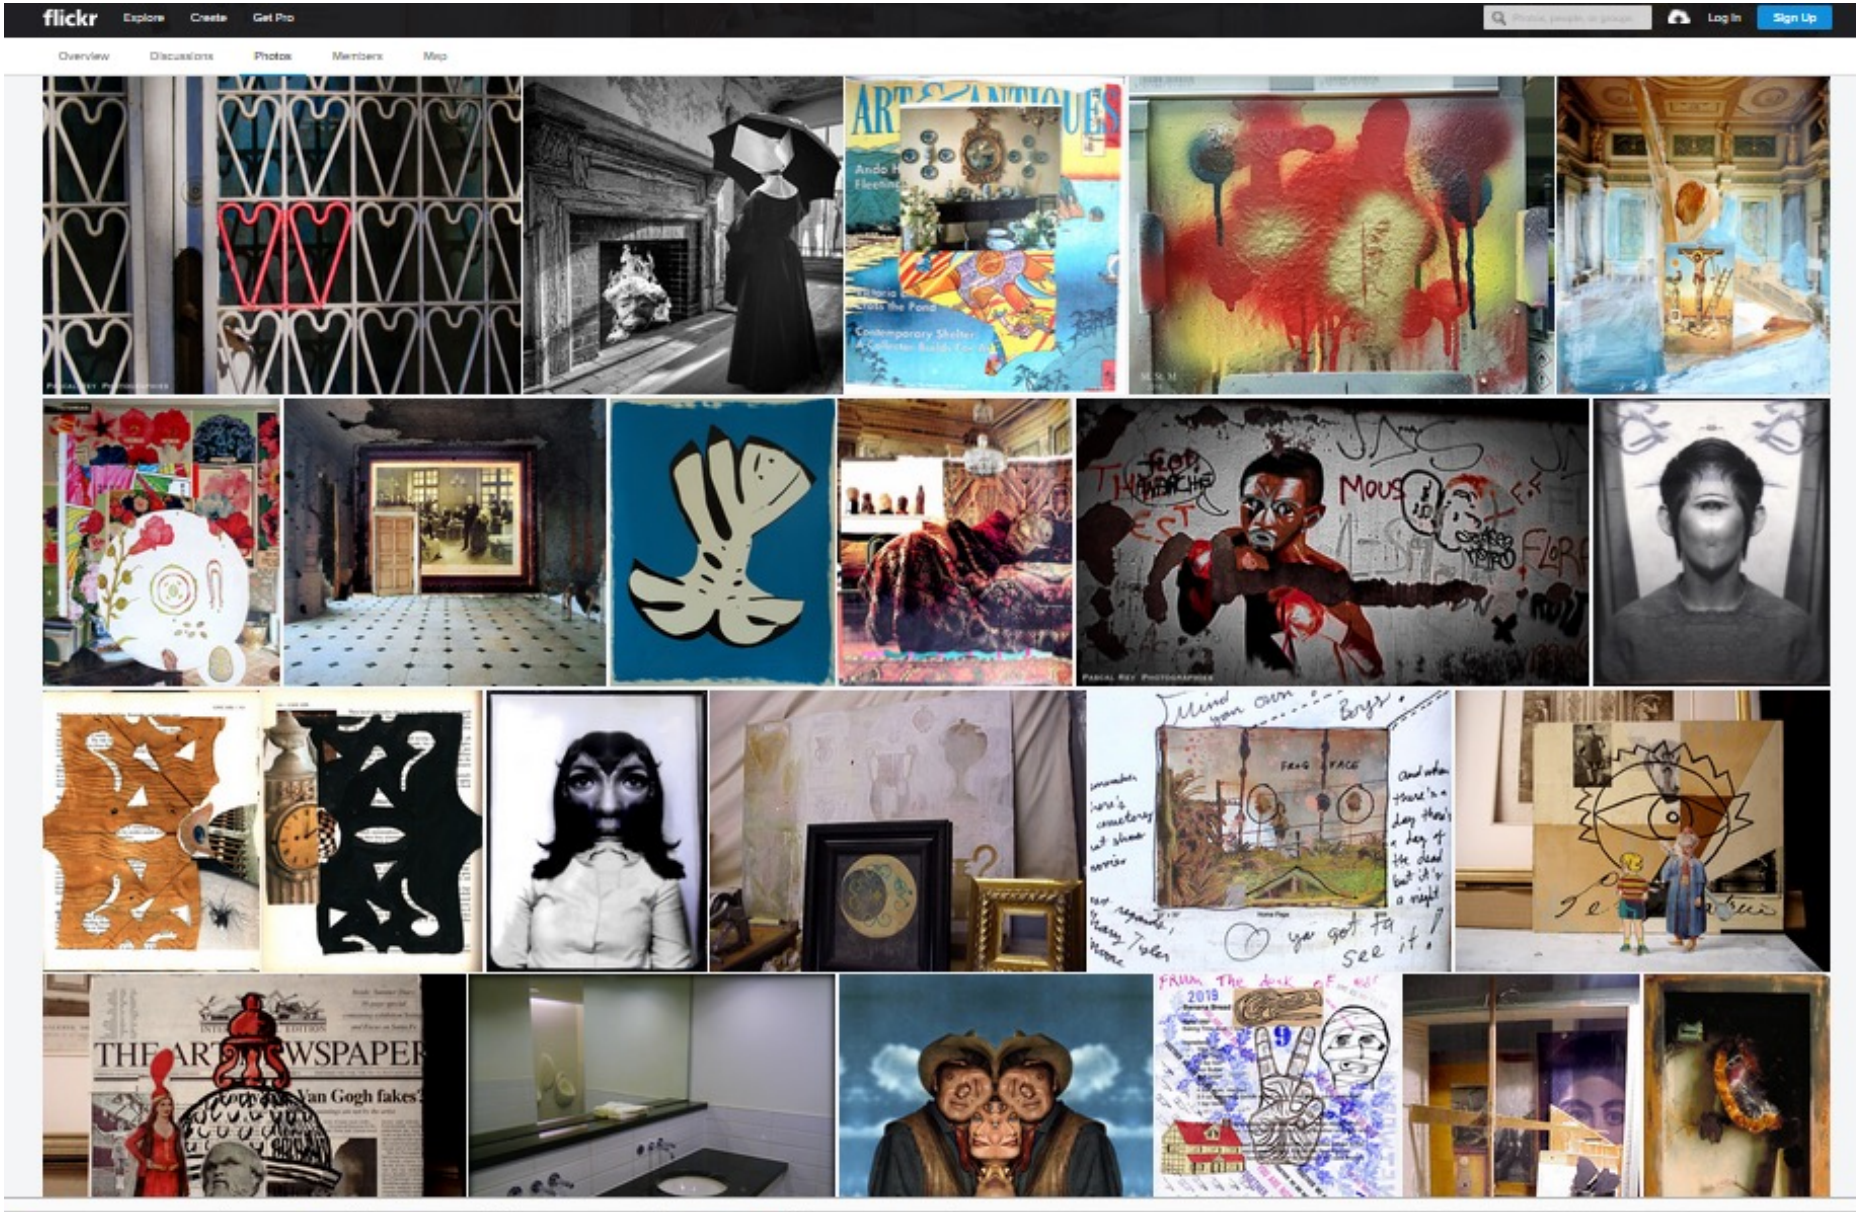

https://www.flickr.com/groups/damndada/pool/

flickr groups: aggregates of images, tags, users

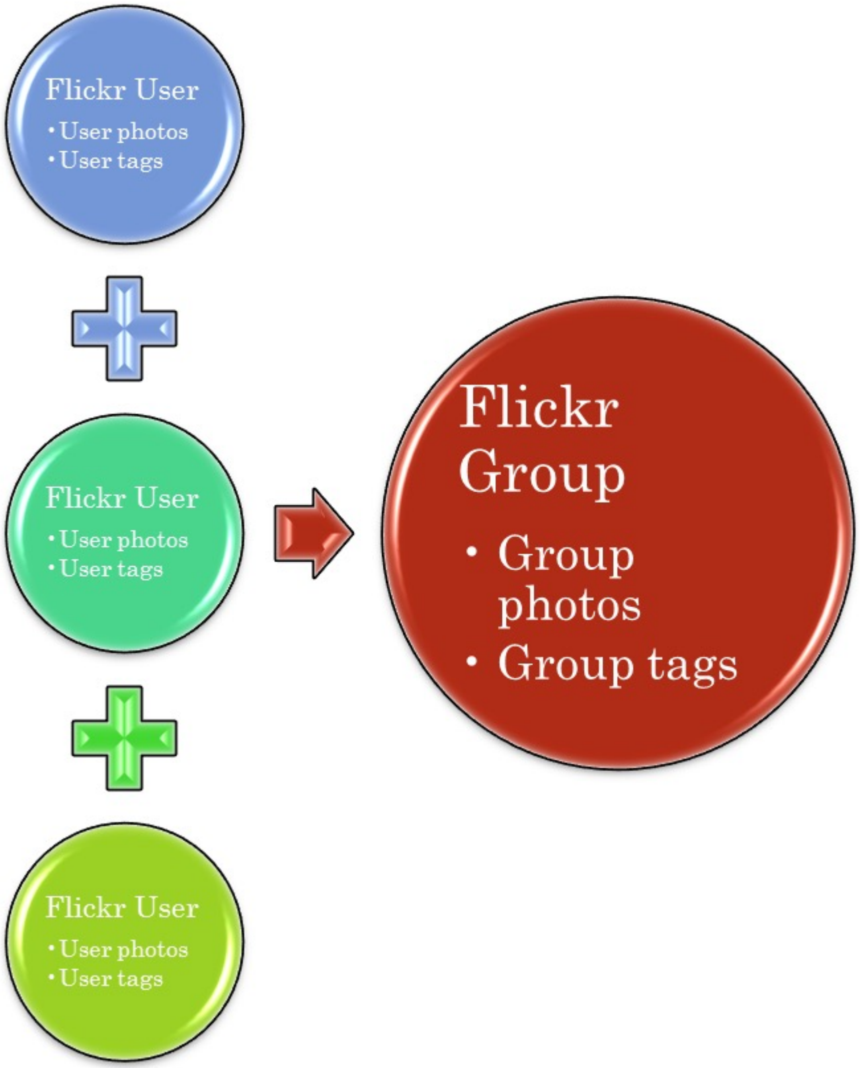

self-organized users with
common interests:
- photographic techniques
- specific content
- geographical locations
- social events
- "just" social interaction

sharing photos with groups creates pools of **photos & tags**

how to discover groups in flickr?
- **topic-model** group representation
- probabilistic unsupervised learning

R. Negoescu and D. Gatica-Perez, Analyzing Flickr Groups, in Proc. ACM CIVR 2008

quick detour: introduction to probabilistic graphical models

Materials taken from C. Bishop, Pattern Recognition and Machine Learning, Springer, 2006.

representing joint distributions with graphs

- **nodes**
  - subsets of random variables (RVs)
  - discrete or continuous
- **vertices**
  - relations between RVs
  - directed or undirected
- **resulting graphs**
  - acyclic (no loops)
  - cyclic

$p(x_1, \ldots, x_K)$ = ?

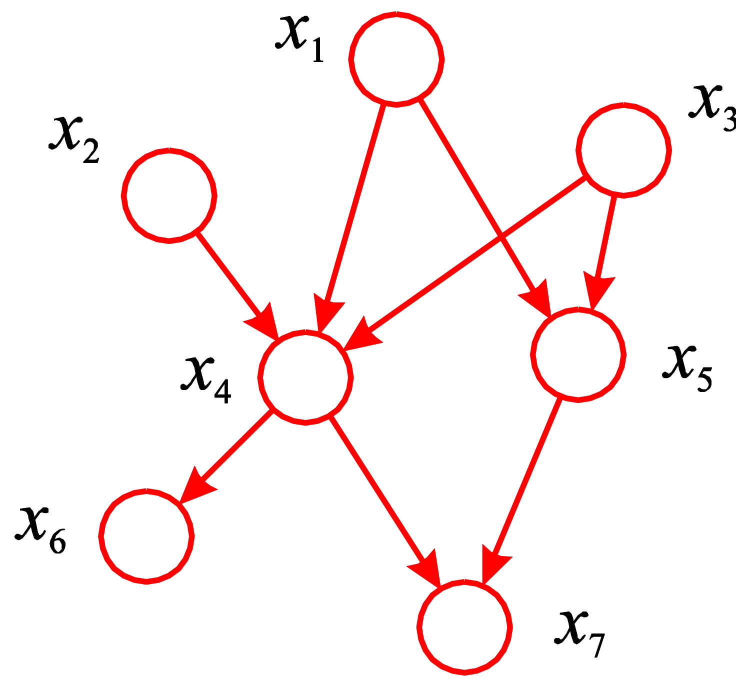

probability theory

this man (with high probability) was not Thomas Bayes!

- **sum rule**

(marginalization)

$$p(\mathbf{x}) = \sum_{\mathbf{y}} p(\mathbf{x}, \mathbf{y})$$

- **product rule**

(conditional probability)

$$p(\mathbf{x}, \mathbf{y}) = p(\mathbf{x}|\mathbf{y})p(\mathbf{y})$$

- **Bayes' theorem**

(posterior $\propto$ likelihood × prior)

$$p(\mathbf{y}|\mathbf{x}) = \frac{p(\mathbf{x}|\mathbf{y})p(\mathbf{y})}{p(\mathbf{x})}$$

$$p(\mathbf{x}) = \sum_{\mathbf{y}} p(\mathbf{x}|\mathbf{y})p(\mathbf{y})$$

bayesian networks (BNs): directed graphical models

- **directed acyclic graphs**
  - no closed paths in the graph
  - we cannot go from node to node along vertices on the direction of the arrows and end up at the original node
  - nodes have 'parents' and 'children'
  - x_4 is a child of $x_1 x_2$, x_3 and is a parent of x_6 and x_7
  - x_1 has no parents

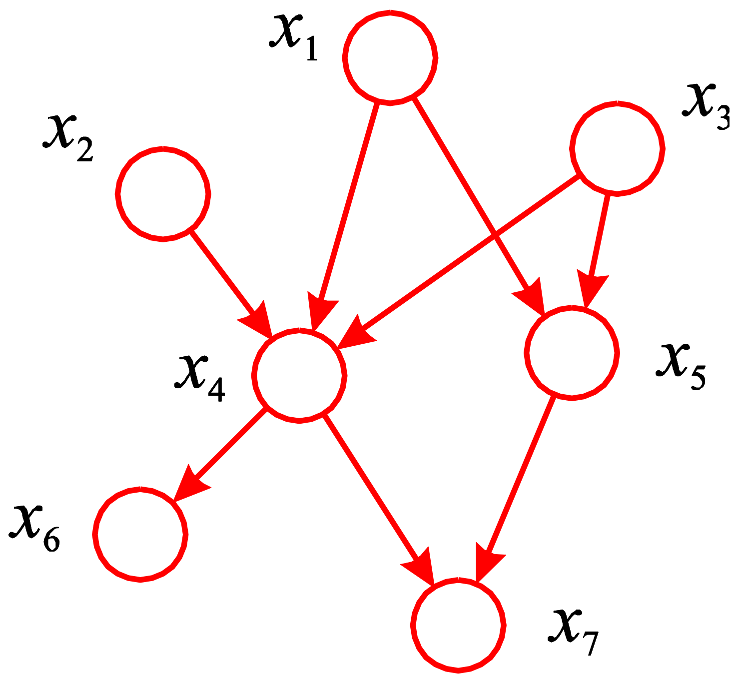

no directed cycles

BNs: the basic equation

- joint distribution for $\mathbf{x} = \{x_1, \ldots, x_K\}$

$$p(\mathbf{x}) = \prod_{k=1}^{K} p(x_k|\mathrm{pa}_k)$$

pa_k : set of parents of x_k

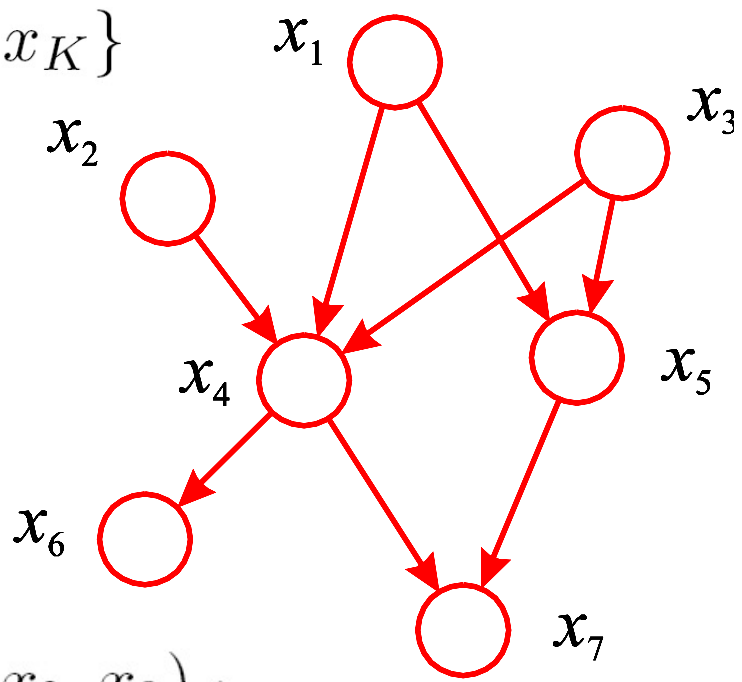

$$p(\mathbf{x}) = p(x_1)p(x_2)p(x_3)p(x_4|x_1,x_2,x_3) \cdot$$
$$p(x_5|x_1,x_3)p(x_6|x_4)p(x_7|x_4,x_5)$$

BNs: the plate notation

$$p(\mathbf{t}, \mathbf{w}) = p(\mathbf{w}) \prod_{n=1}^{N} p(t_n|\mathbf{w}).$$

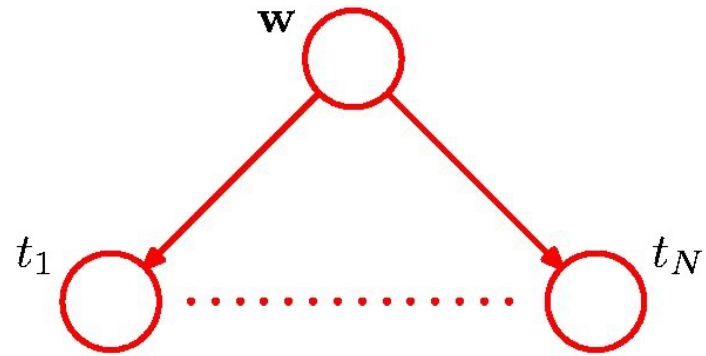

© c. bishop

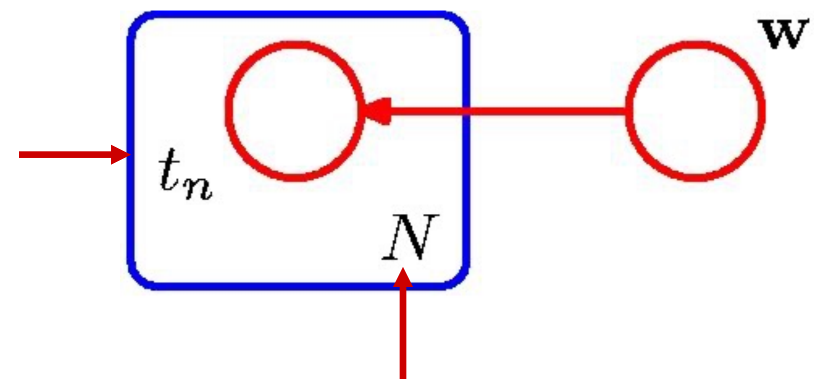

plate

number of nodes of the same kind

BNs: types of variables

- variables may be **hidden** (latent) or **visible** (observed)

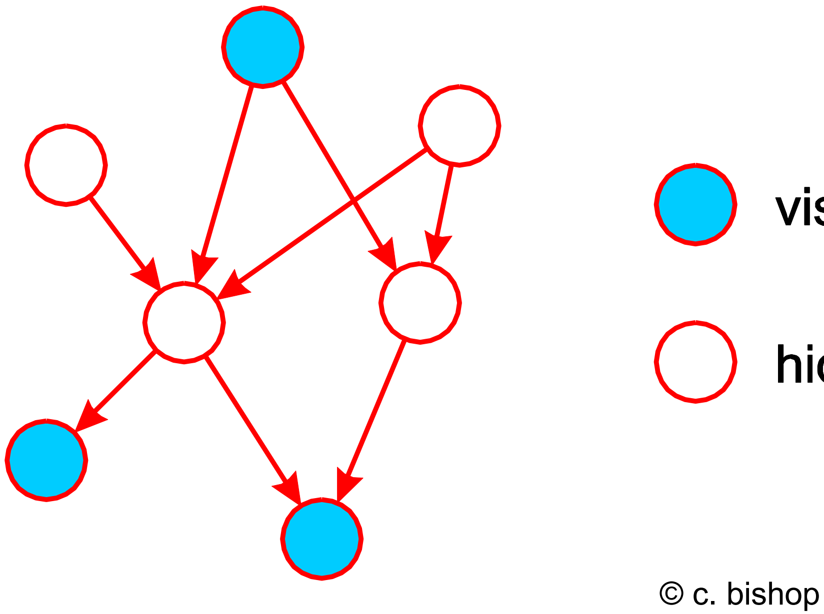

- visible variables
  - physical measurements (text, audio, image)
- latent variables
  - define richer models
  - often have a clear interpretation

BNs: two key problems

- given a parametric form for $$p(\mathbf{x}) = \prod_{k=1}^{K} p(x_k|\mathrm{pa}_k)$$

- learning: given training data, estimate the model parameters

- inference: given a learned model, compute probabilities of hidden nodes

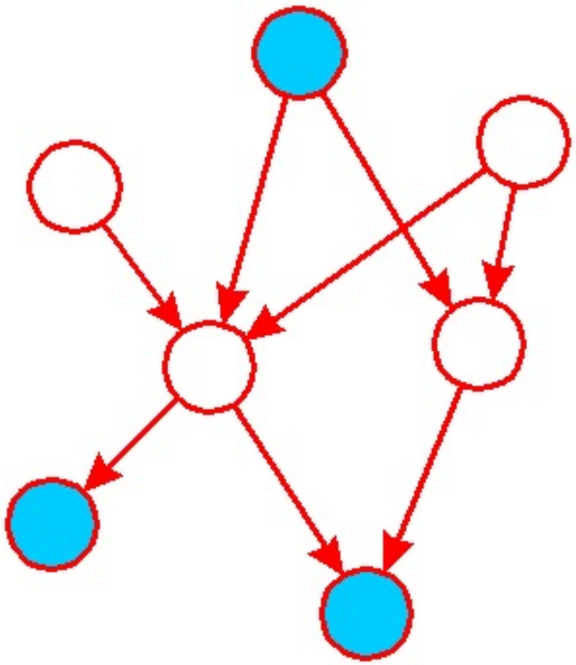

© c. bishop

probabilistic topic models

probabilistic topic models: unsupervised learning for text collections

"Professional football leagues are impacted economically by the coronavirus. Players see their salaries reduced as the leagues are interrupted and clubs lose money. Fans are afraid of getting sick and spreading the disease"

Topics: **Sports** **Economics** **Health**

assumptions:

+ observed documents: bags of words (vectors of word counts)
+ documents are mixtures of topics
+ latent topics: "soft clusters" of words
+ generative models: documents are generated by sampling words from the topics

S. Deerwester, S.Dumais, T. Landauer, G. Furnas, and R. Harshman, Indexing by latent semantic analysis. *J. Am. Soc. Inform. Sci. 41*, 6 (1990), 391–407.
T. Hofmann, Probabilistic latent semantic analysis. In Proc. Uncertainty in Artificial Intelligence (UAI) ,1999.

latent Dirichlet allocation (LDA)

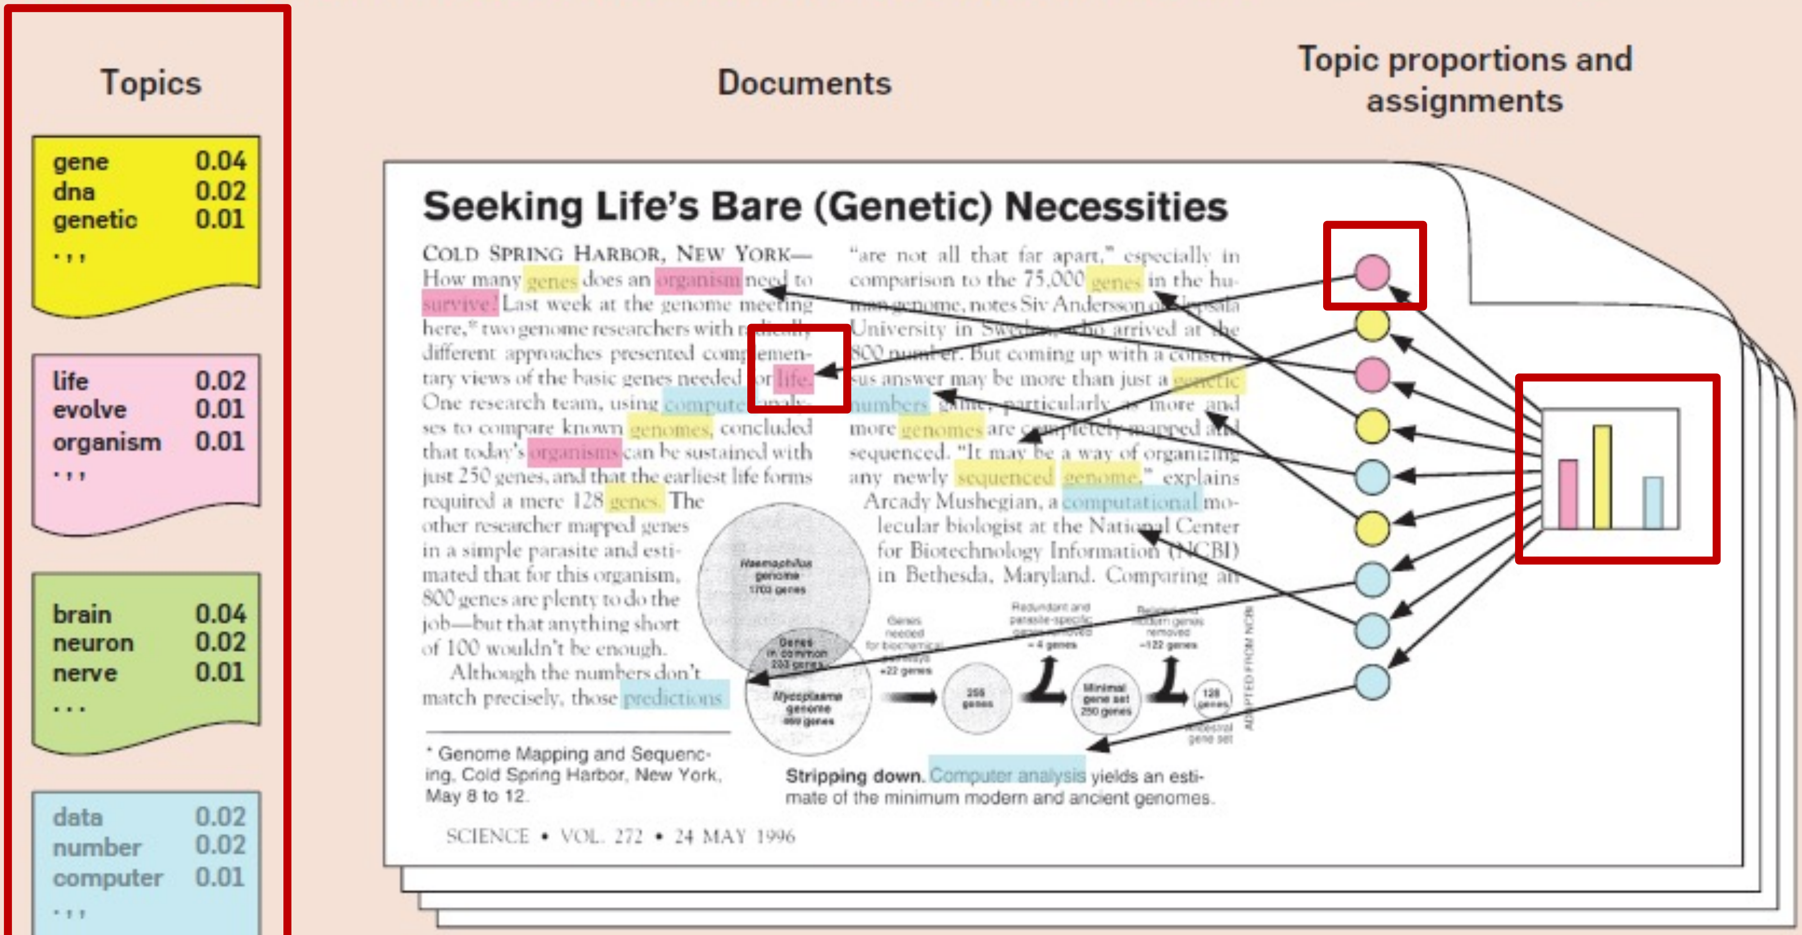

Figure 1. The intuitions behind latent Dirichlet allocation. We assume that some number of "topics," which are distributions over words, exist for the whole collection (far left). Each document is assumed to be generated as follows. First choose a distribution over the topics (the histogram at right); then, for each word, choose a topic assignment (the colored coins) and choose the word from the corresponding topic. The topics and topic assignments in this figure are illustrative—they are not fit from real data.

D. M. Blei, A. Ng, M. Jordan, Latent Dirichlet allocation, J. of Machine Learning Research, 3, Jan. 2003
D. M. Blei. Probabilistic topic models. Communications of the ACM, 55, 4, 77-84, Apr. 2012

latent Dirichlet allocation (2)

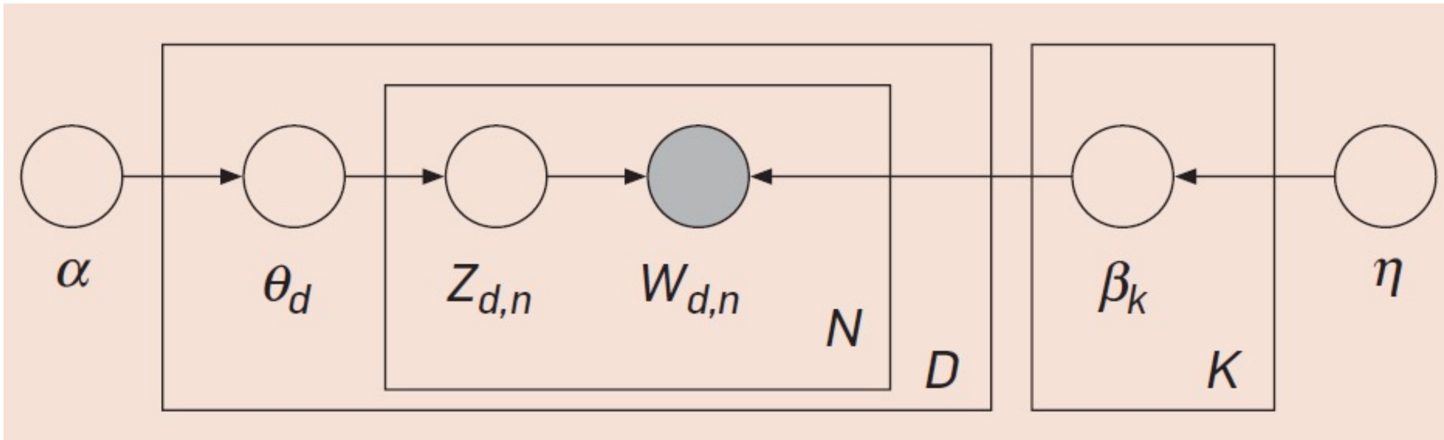

$$p(\beta_{1:K}, \theta_{1:D}, z_{1:D}, w_{1:D}) = \prod_{i=1}^{K} p(\beta_i) \prod_{d=1}^{D} p(\theta_d) \left(\prod_{n=1}^{N} p(z_{d,n} | \theta_d) p(w_{d,n} | \beta_{1:K}, z_{d,n}) \right)$$

Assume D documents, N words per document, and K **latent** topics

Topics are $\beta_{1:K}$,where each β_k is a distribution over the vocabulary

For each document d in the collection D

- Sample topic proportion θ_d ~ Dirichlet(α)
- For each of the N words $w_{d,n}$
  - Sample topic assignment $z_{d,n}$ ~ Multinomial(θ_d)
  - Sample word $w_{d,n} \sim p(w_{d,n} | \beta_{1:K}, z_{d,n})$, a multinomial conditioned on the topic specified by $z_{d,n}$ (out of all possible topics)

D. M. Blei, A. Ng, M. Jordan, Latent Dirichlet allocation, Journal of Machine Learning Research, 3:993-1022, Jan. 2003

D. M. Blei. Probabilistic topic models. Communications of the ACM, 55, 4, 77-84, Apr. 2012

topic modeling for Flickr groups

50k groups, 6.9M photos

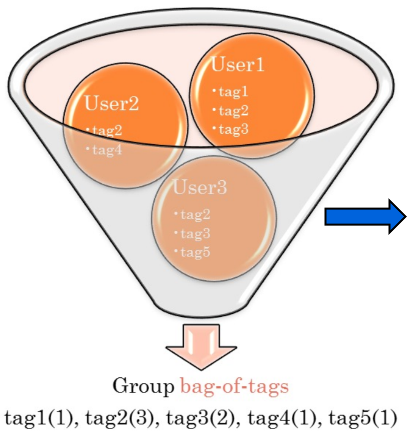

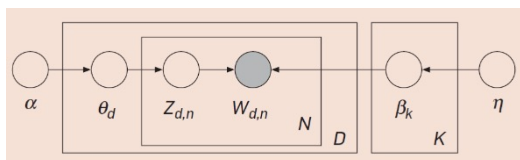

P(t | z)	Tag
0.0766	flower
0.0555	flowers
0.055	nature
0.0431	ilovenature
0.0323	spring
0.0295	garden
0.0243	green
0.0221	yellow
0.0212	macro
0.0204	pink

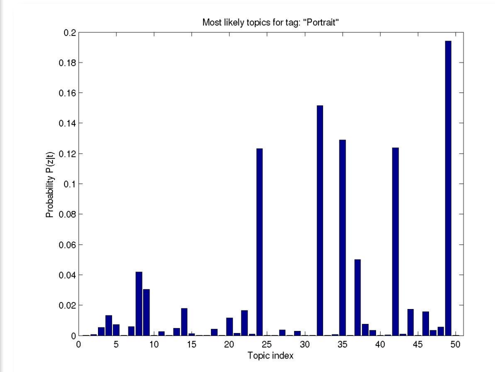

$p(t \mid z) = p(w|z)$

$p(z|G)$

+ groups are **documents (bags of tags)**
+ groups described by the topics they are about
+ inference using Markov Chain Monte Carlo (MCMC)

learned topics: top words and top groups
(using PLSA, precursor to LDA)

Topic 1 top tags	
P(t | z)	Tag
0.0766	flower
0.0555	flowers
0.055	nature
0.0431	ilovenature
0.0323	spring
0.0295	garden
0.0243	green
0.0221	yellow
0.0212	macro
0.0204	pink

Topic 1 top groups	
P(z | G)	Group
0.9715	1-Plants World
0.9456	Flickr Gardens
0.8783	In my garden
0.8718	My Garden
0.8347	Daffodil World
0.8337	What plant is that?
0.8214	Gardening for Fun
0.8102	Garden Flowers
0.7993	grow
0.7377	Backyard Nature

visualizing the top groups for topic 1

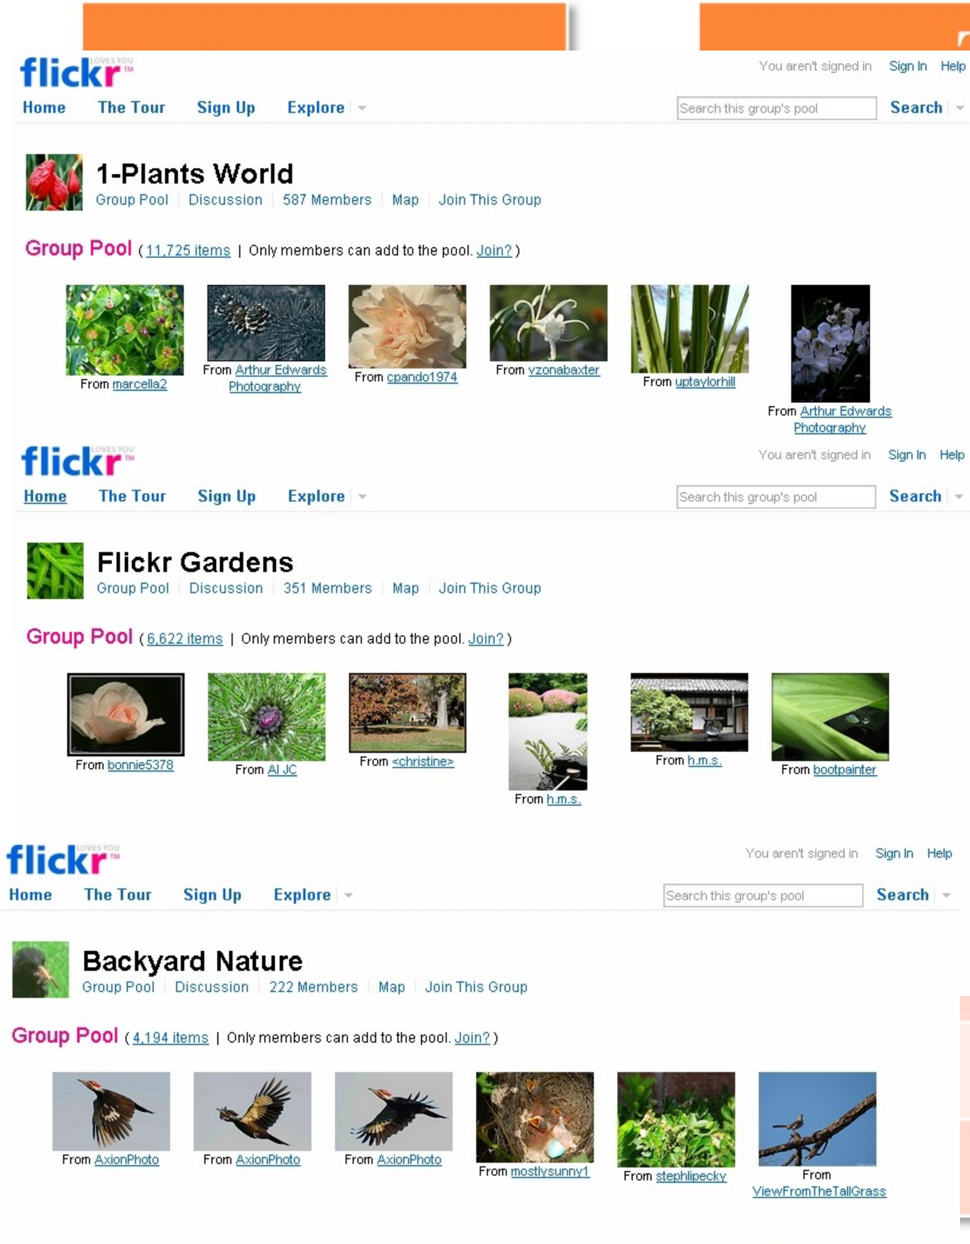

Topic 1 top groups

Group
1-Plants World
Flickr Gardens
In my garden
My Garden
Daffodil World
What plant is that?
Gardening for Fun
Garden Flowers
grow
Backyard Nature

learned topics (2)

Topic 18 top tags	
$P(t \mid z)$	Tag
0.0478	music
0.0175	rock
0.0171	concert
0.0156	live
0.0131	band
0.0127	party
0.0124	florida
0.0123	guitar
0.0104	friends
0.0088	label

Topic 18 top groups	
$P(z \mid G)$	Group
0.9917	**LIVE in CONCERT**
0.9783	Vinyl Junkie
0.973	BUSH-IT Artist
0.9512	REHNQUIST RETIRES THE WAR BEGINS
0.9386	Rock and Roll : live shows only please
0.9307	Concerts
0.9234	Rock in Paris
0.9171	Live Music Photography
0.9135	SINGERS SING! (4 pics at any one time)
0.9088	Concerts!!

learned topics (2)

Topic 18 top tags	
P(t | z)	Tag

Topic 18 top groups	
P(z | G)	Group
0.9917	**LIVE in CONCERT**

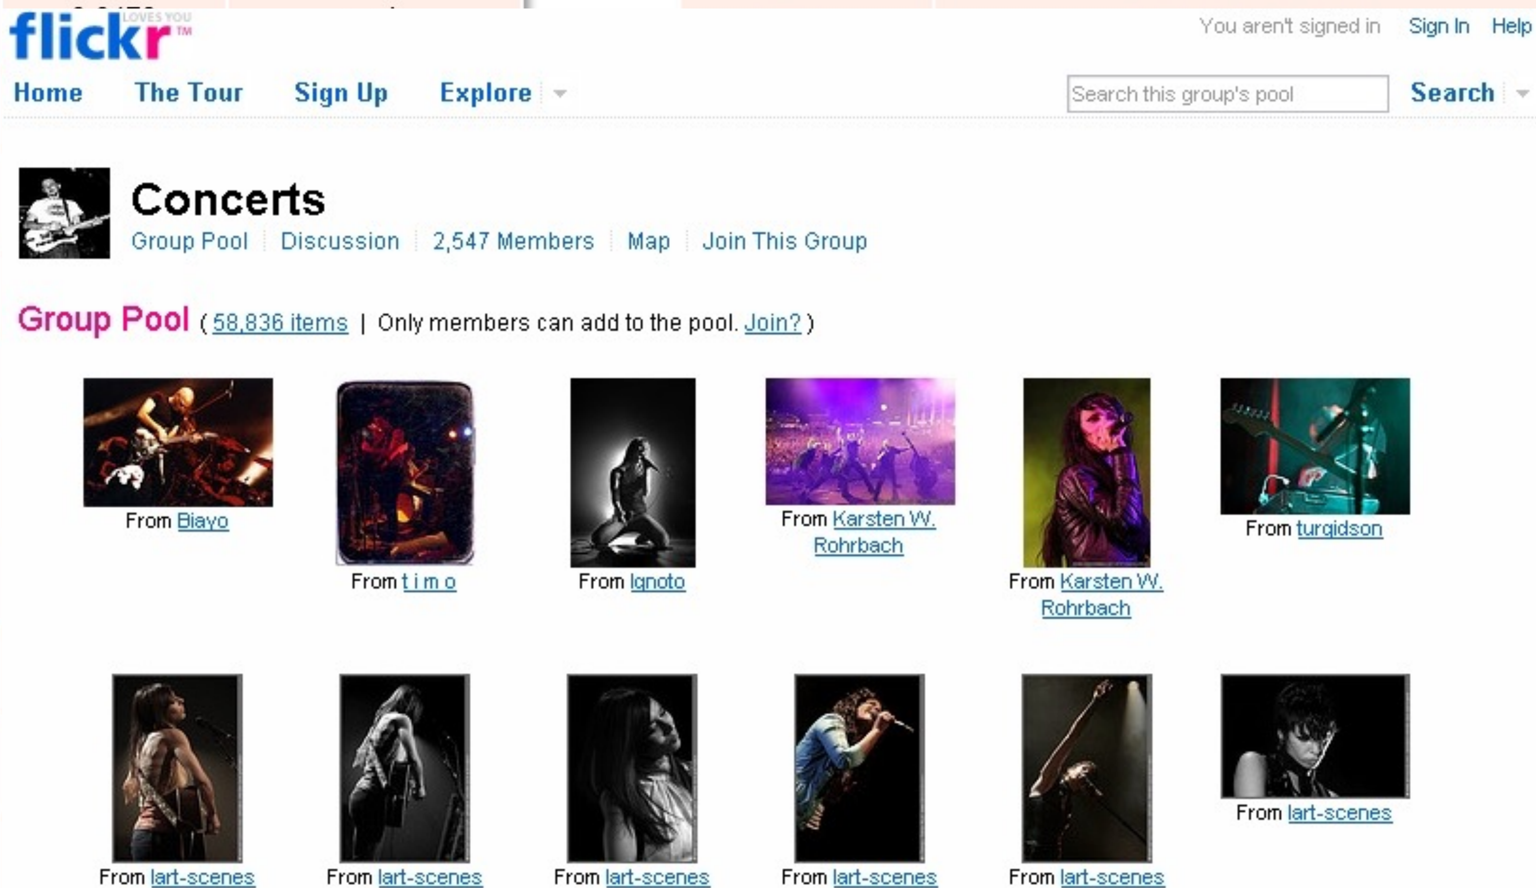

learned topics (3)

Topic 2 top tags	
P(t | z)	Tag
0.0957	canada
0.0397	bc
0.0343	snow
0.0334	vancouver
0.024	britishcolumbia
0.0213	ontario
0.021	winter
0.0129	water
0.0128	mountain
0.0127	ice

Topic 2 top groups	
P(z | G)	Group
0.9978	BC Peaks & Mountains
0.9971	A S C E N T - (how you get to the top)
0.9937	British Columbia Provincial Parks
0.9922	Climbing Photography
0.9809	Rock Climbing
0.9667	Climbing lifestyle
0.965	Climbing
0.9632	Where am I in BC
0.951	ROCKCLIMBING
0.9421	Alpinism

learned topics (4)

Topic 13 top tags	
$P(t \mid z)$	Tag
0.0846	holland
0.0613	netherlands
0.0458	nederland
0.0255	thenetherlands
0.021	amsterdam
0.0182	denhaag
0.0148	bike
0.0141	dutch
0.0136	bw
0.0119	rotterdam

Topic 13 top groups	
$P(z \mid G)$	Group
0.9599	Den Haag (The Hague)
0.9591	Den Haag / The Hague, The Netherlands
0.8626	goingdutch
0.8202	1-2-3 Nederland
0.7831	Nederland/The Netherlands
0.7679	Made in Holland
0.7671	Dutch
0.7665	horses
0.7639	Dutch skylines
0.7566	Amsterdam today

how many topics are flickr groups about?

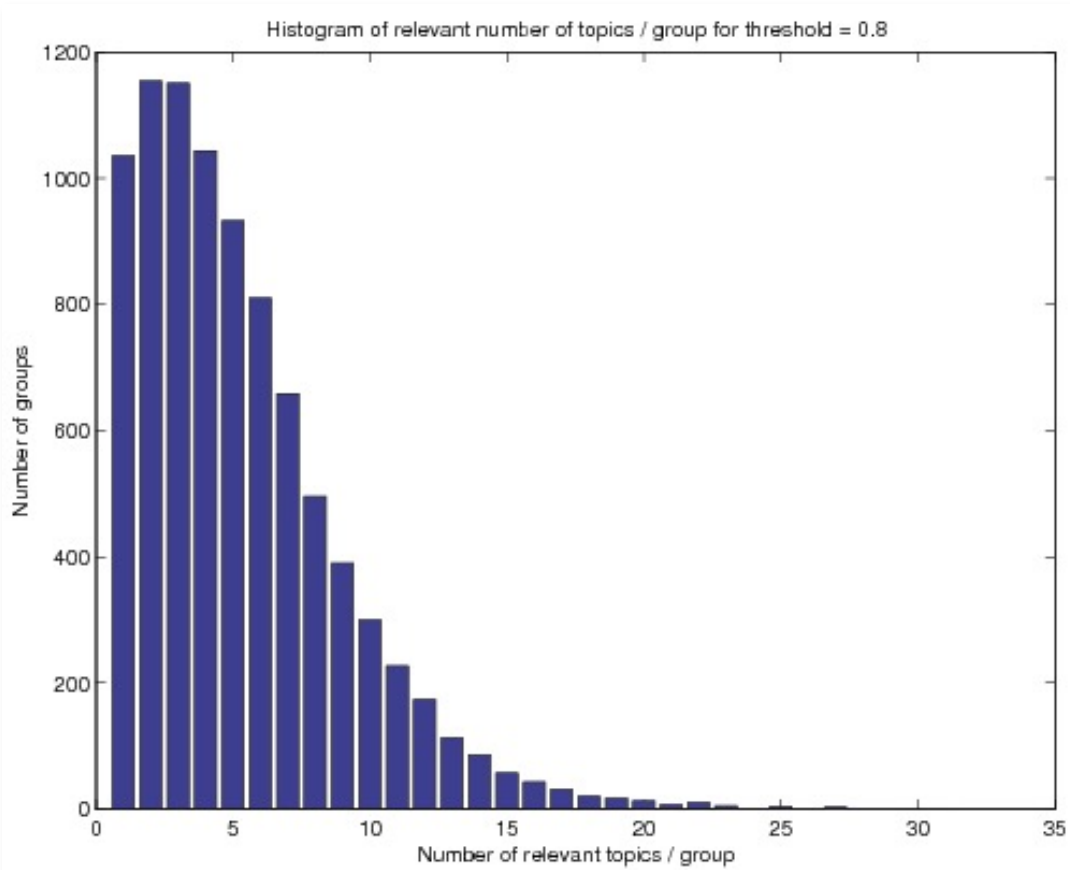

topic-expert groups have spiky topic distributions, with one dominant topic

~35% of groups with 3-5 topics: focused interests (content-oriented groups)

less semantic coherence in large social-oriented groups (20+ topics)

topic decomposition for group *Portrait*

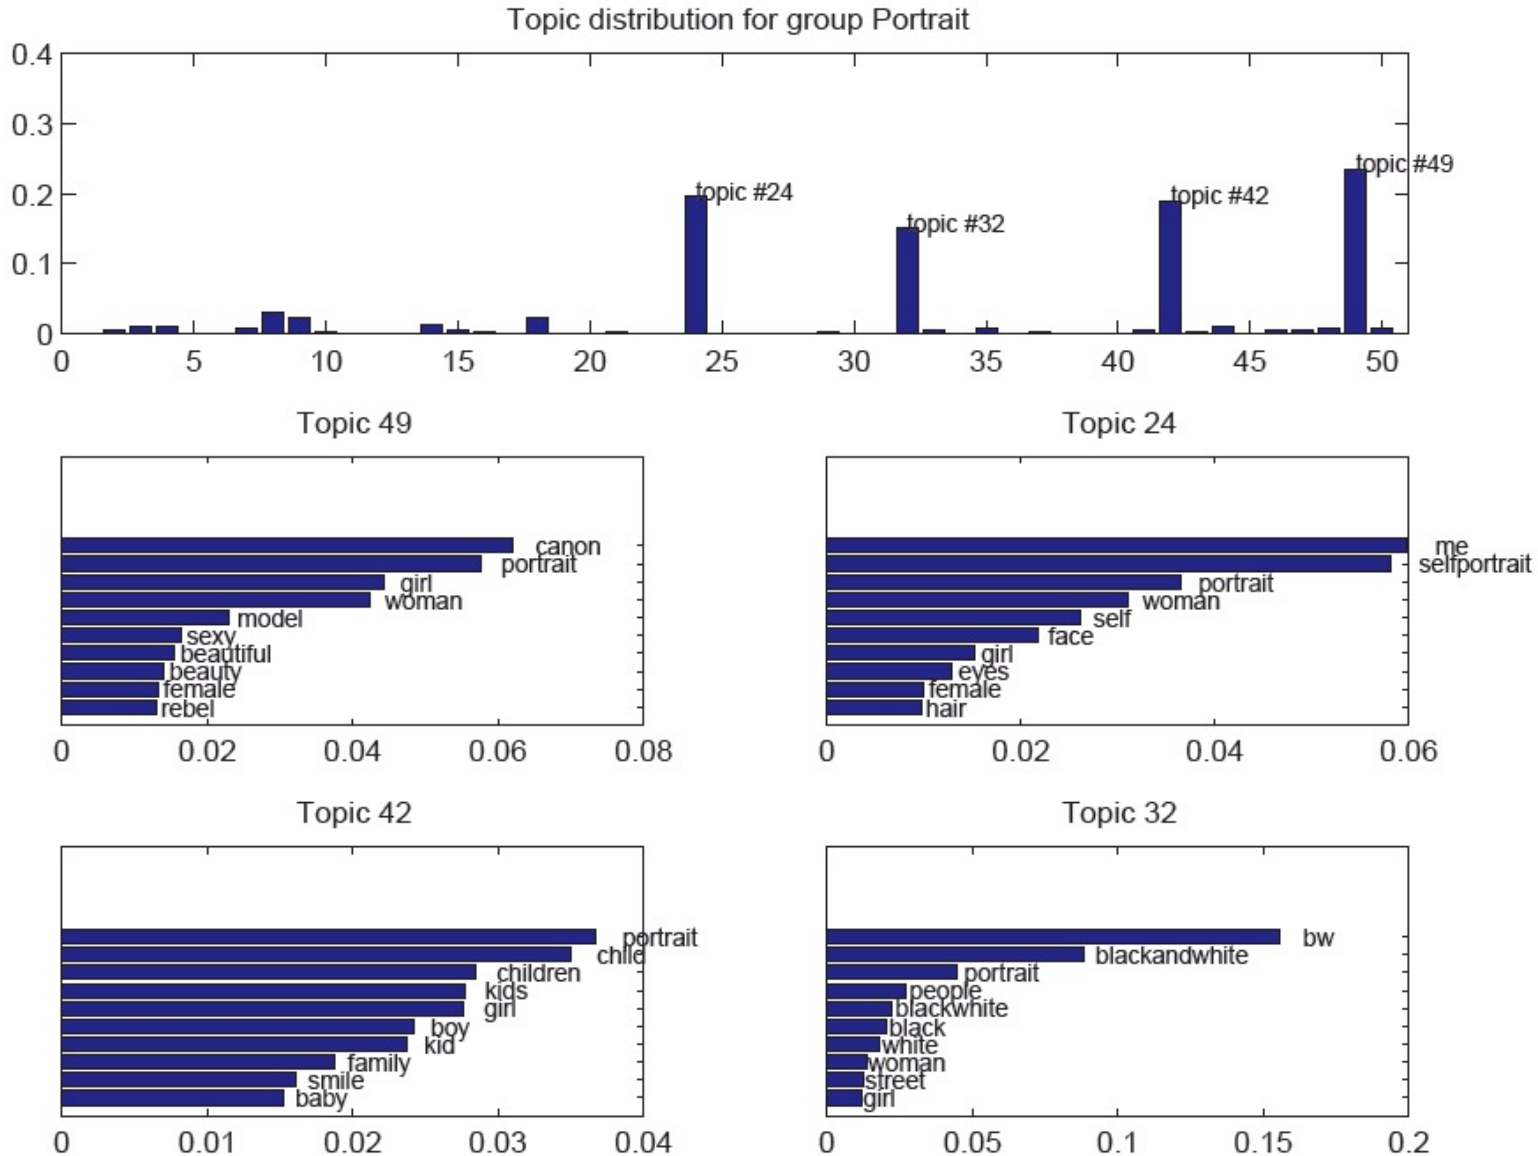

topic decomposition for group *Flickr Central*

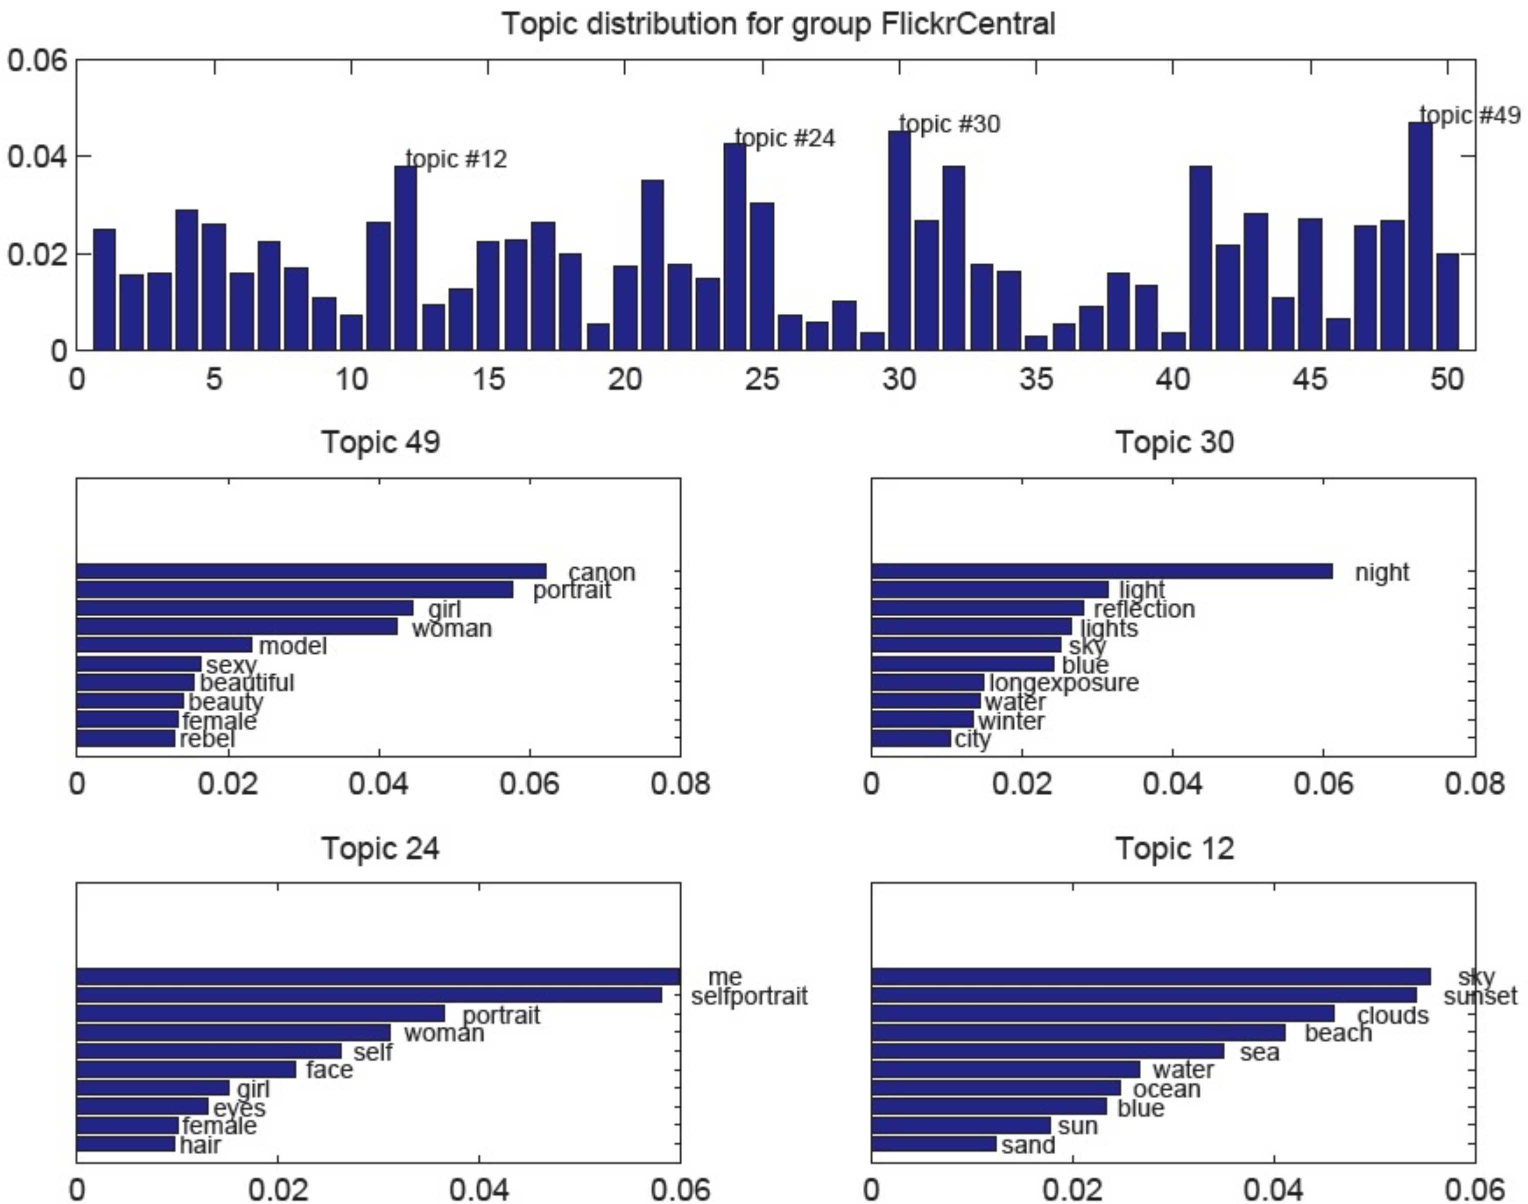

comparing image sharing practices

two image sharing practices

kodak culture: take photos to share with a small existing social group or to archive

snaprs: take photos with the primary objective of sharing with the world

credit (cc): Photo by Philippe Oursel on Unsplash, https://unsplash.com/photos/06y6wukkSKg

A. Miller & W. Edwards, Give and Take: a Study of Consumer Photo-sharing Culture and Practice, in Proc. ACM CHI, 2007.

credit (cc): Klearchos Kapoutsis @flickr:
https://www.flickr.com/photos/klearchos

two datasets

flickr: images with tags
kodak gallery: images with free text

Statistic	Flickr	Kodak
Total photos	4.6M	413,000
Total tag occurrences	13M	900,000
Total users	25,800	5400
Photos / user	157	76
Unique tags / user	81	34

10,000 most popular words (in terms of number of users) for each source

R. Negoescu, A. Loui, & D. Gatica-Perez, Kodak Moments and Flickr Diamonds: How Users Shape Large-Scale Media, in Proc. ACM Multimedia, 2010

comparing kodak & flickr language

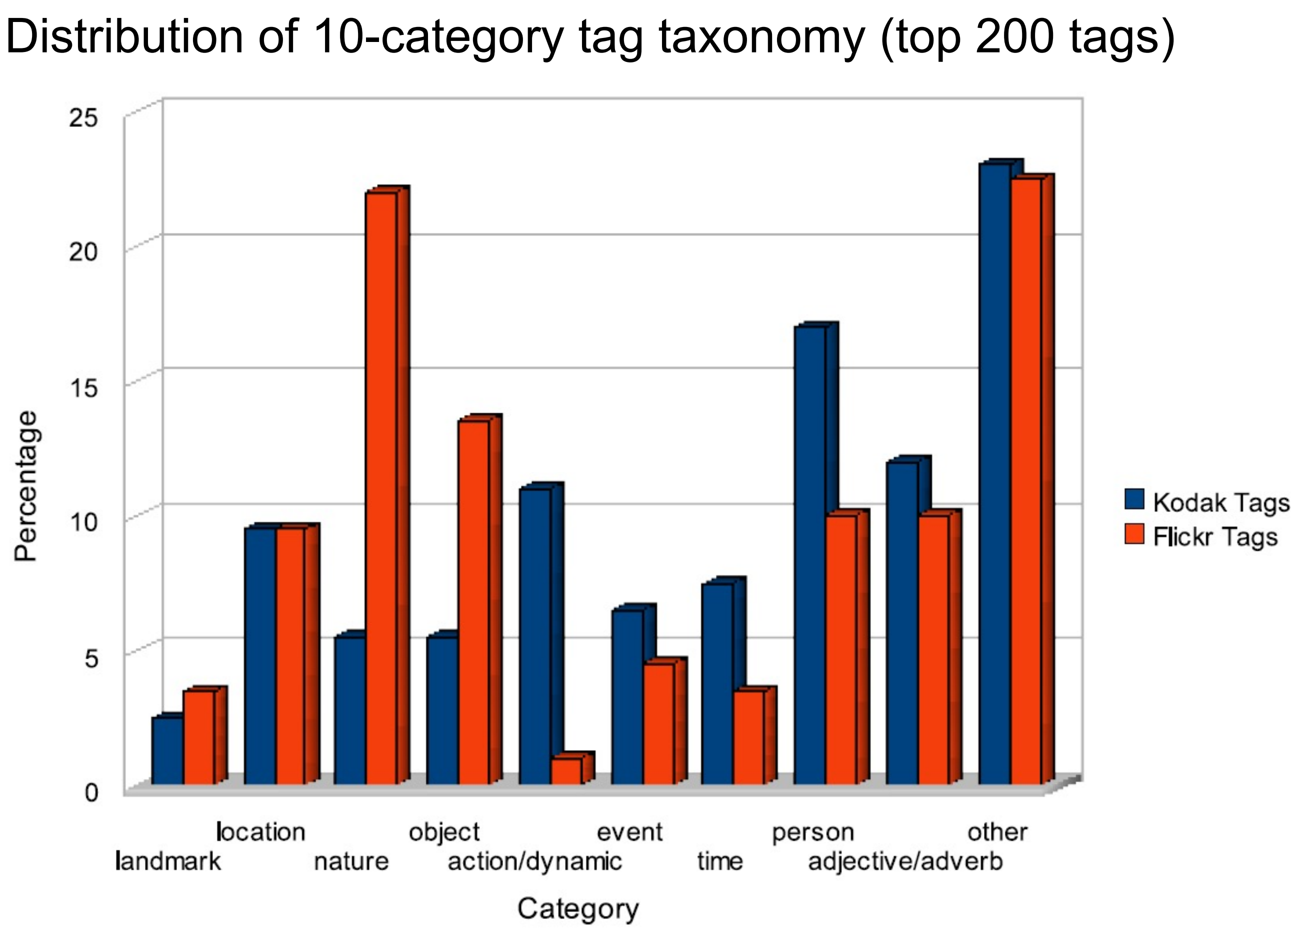

+ **flickr but not kodak**
  macro, selfportrait, blackandwhite, photoshop, flickr, abigfave, geotagged
+ **kodak but not flickr**
  enjoying, showing, giving, checking, loved, visiting, dressed, wearing

using topic models to study topic specificity for kodak & flickr

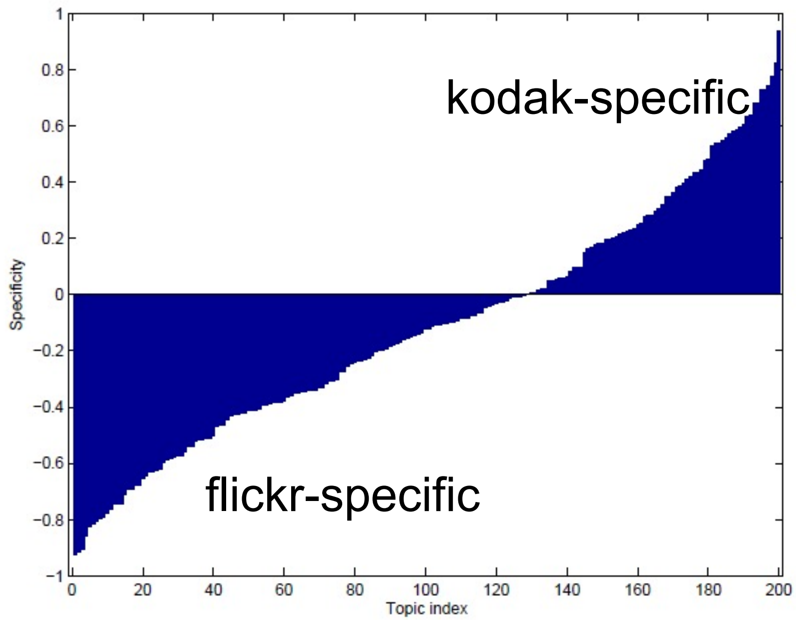

Figure 3: Topic specificity among the two communities. Specificity is computed as the ratio of the difference between Kodak and Flickr users for which that topic is relevant, and the total number of users for which the topic is relevant.

+ 5400 users for each source
+ learn LDA on **joint vocabulary**
+ determine relevant topics for each user based on topic distribution

+ topic specificity

"flickr" topic: abigfave, flickrdiamond,anawesomeshot

"flickr" topic: nature, landscape, flora, ilovenature, plant, animal

"kodak" topic: i, time, daddy, ready, mommy, love, big, happy, playing

"kodak" topic: picture, pictures, edited, png

flickr contributions to open datasets

creative commons

Attribution 2.0 Generic (CC BY 2.0)

This is a human-readable summary of (and not a substitute for) the **license**.

Disclaimer

Free Cultural Works APPROVED FOR

You are free to:

Share — copy and redistribute the material in any medium or format

Adapt — remix, transform, and build upon the material

for any purpose, even commercially.

The licensor cannot revoke these freedoms as long as you follow the license terms.

Under the following terms:

Attribution — You must give **appropriate credit**, provide a link to the license, and **indicate if changes were made**. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use.

No additional restrictions — You may not apply legal terms or **technological measures** that legally restrict others from doing anything the license permits.

https://creativecommons.org/licenses/by/2.0/

DOI:10.1145/2812802

This publicly available curated dataset of almost 100 million photos and videos is free and legal for all.

BY BART THOMEE, DAVID A. SHAMMA, GERALD FRIEDLAND, BENJAMIN ELIZALDE, KARL NI, DOUGLAS POLAND, DAMIAN BORTH, AND LI-JIA LI

YFCC100M: The New Data in Multimedia Research

B. Thomee, D. A. Shamma, G. Friedland, B. Elizalde, K. Ni, D. Poland, D. Borth, and L.-J. Li. YFCC100M: the new data in multimedia research. *Communications of the ACM* 59, 2, 64-73, January 2016

what to remember

sharing images with family & friends is an old practice, transformed by digital, online, social

professional producers & consumers of images also saw their practices transformed

flickr as an early social image system

- convergence of amateur & professional photographers
- photos, tags & communities as key features

probabilistic topic models

- tool to analyze tagged image collections

questions?

daniel.gatica-perez@epfl.ch