Beyond Relevance!

Roelof van Zwol

Yahoo! Research

roelof@yahoo-inc.com

YAHOO!

Beyond Relevance?

Definition: “Closely connected or appropriate to the matter at hand”

Our comfort zone: precision - recall - gain

Beyond Relevance!

The true challenge for multimedia is to find a balance between **relevancy**, **freshness**, **quality**, **interestingness** and **diversity** in order to provide an **engaging** rich **media experience** to the user.

What to expect?

- Understanding users
- Diversity
- Reforming image search
- Image fingerprinting
- Faceted search
- Smarter thumbnails
- Explore / exploit of images
- Personalization in social media
- Predicting Flickr favorites
- Automated slideshow generation

Understanding Users

Based on image search query log analysis

Understanding users

- 3 months in oct - dec 2009
- Query volume: 1.3 B
- Unique queries 75.2 M
- Terms per queries 2.39

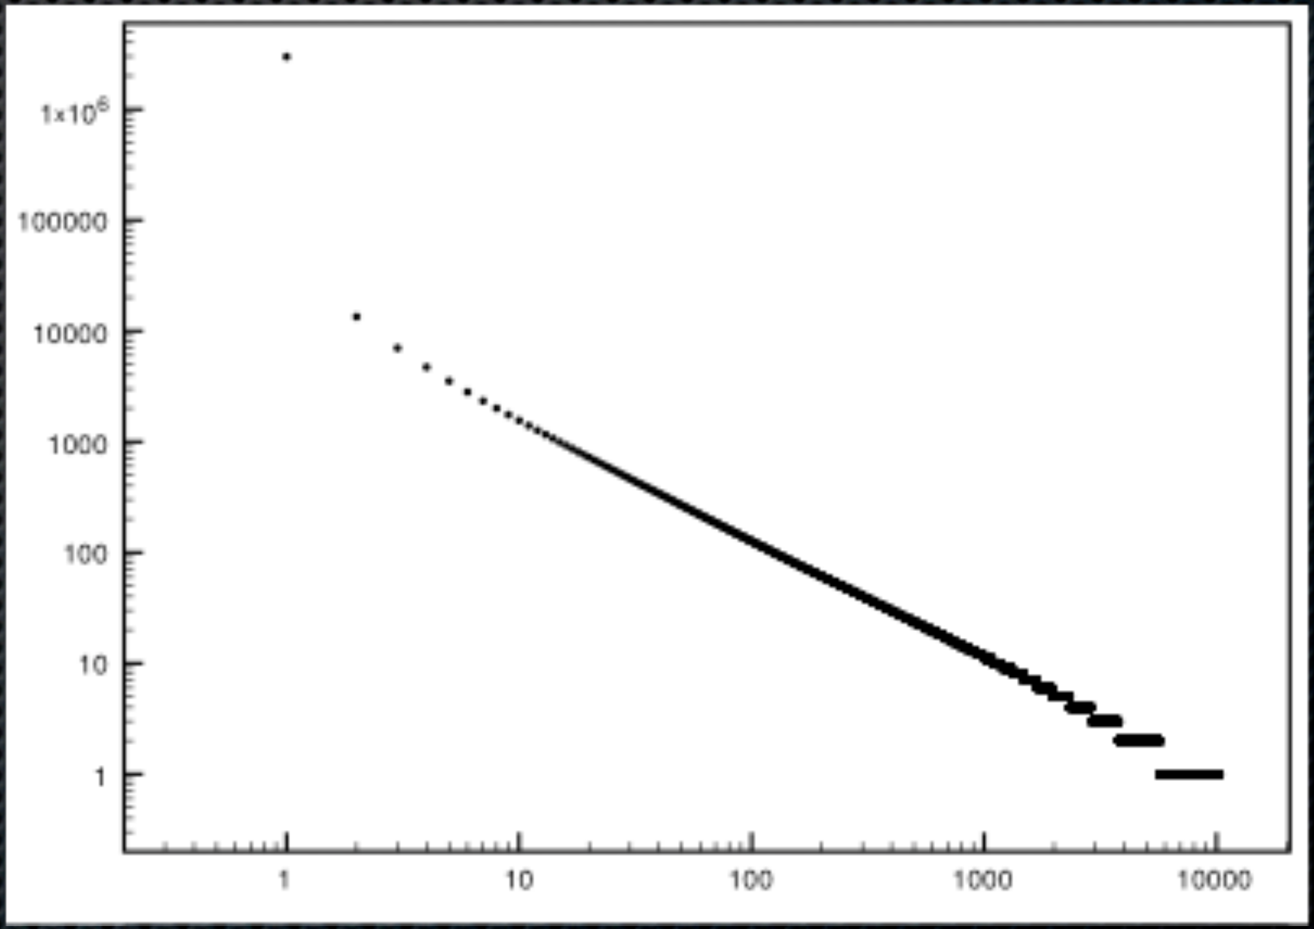

What users search for...

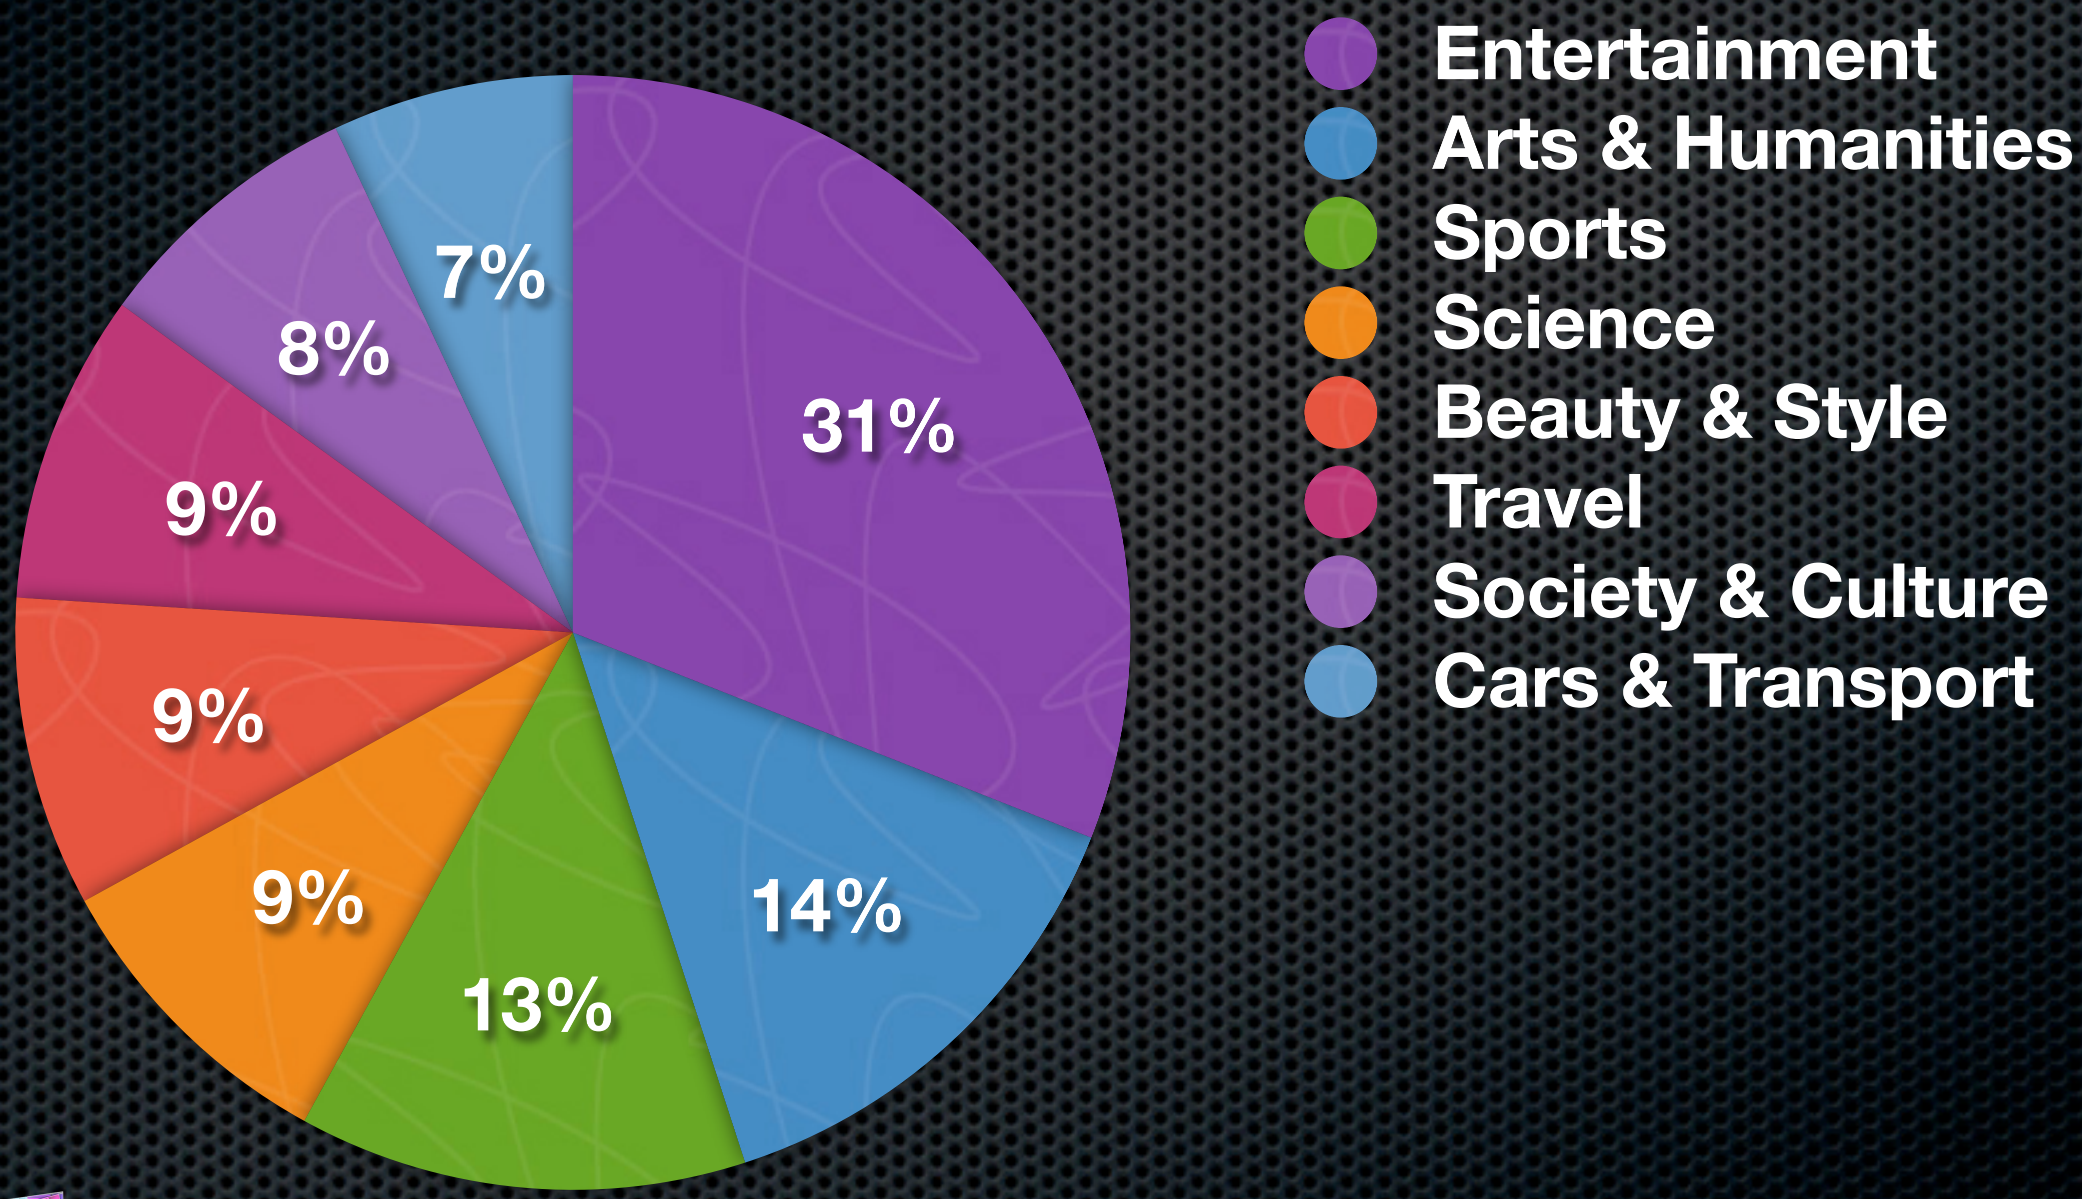

Breakdown of entertainment

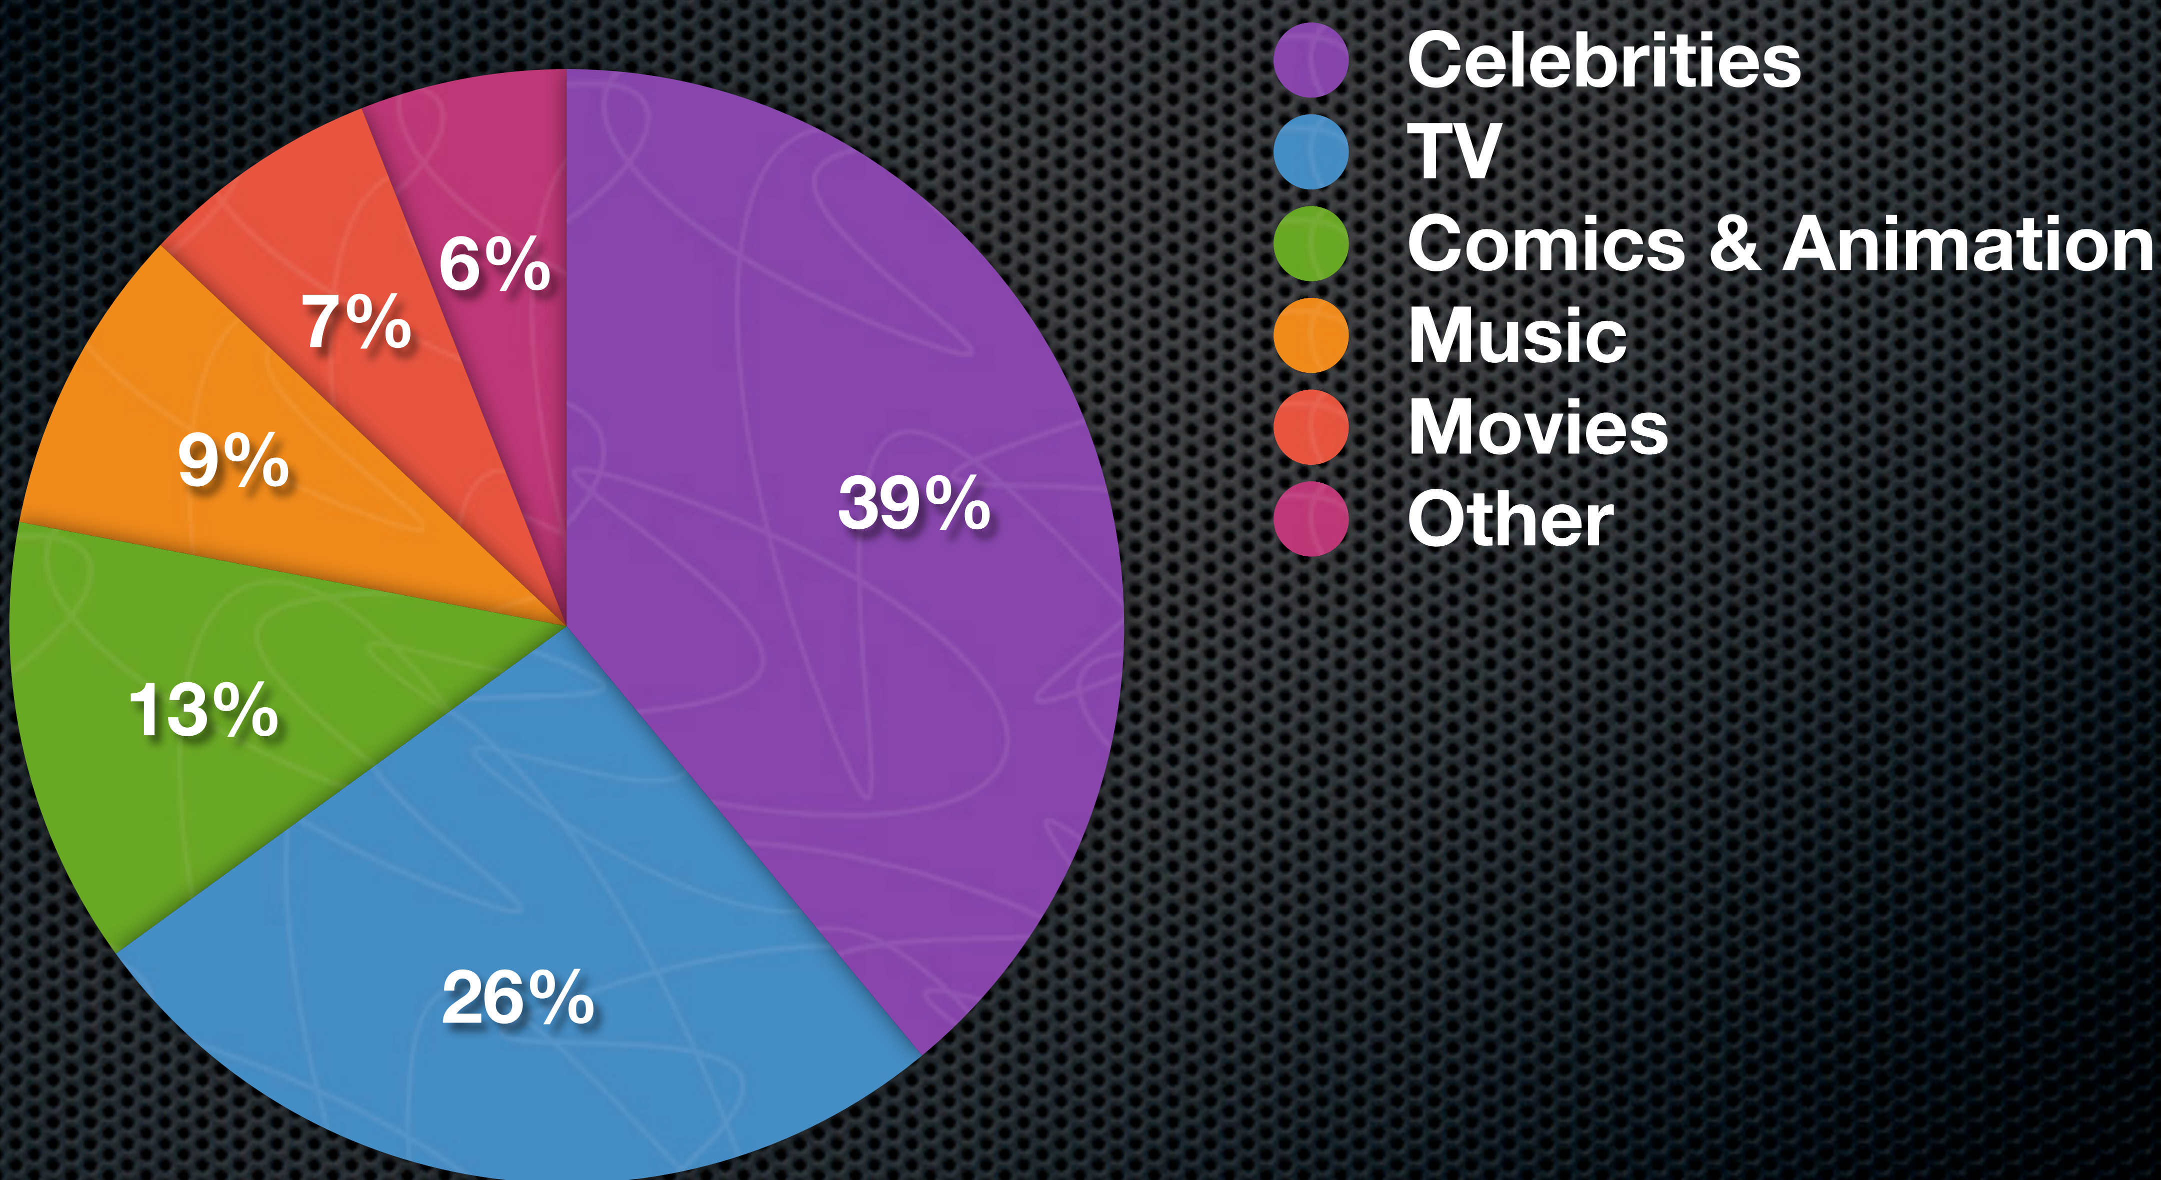

User session analysis

	Image Search	Web Search
Number of Sessions	187,422,028	?
Queries per Session	4.52	1.25

Classification of user actions

Action	Description
Initial query	The first query issued by a user.
Next page	The user goes to the next (or previous) search page.
More specific	The user refined the search query.
More generic	The user made the query more generic.
Minor rewrite	The user made minor edits to the query, e.g. spelling corrections (Levenshtein distance <5).
Major rewrite	The user made significant changes with respect to the previous search query.

Session 1	Action Type
cursive letters	Initial
cursive letters	Next page
tattoo fonts	Rewrite
tattoo fonts	Next page (4x)
tattoo fonts lettering	Specific
tattoo fonts lettering	Next page (2x)
fonts	Generic
fonts	Next page (4x)
tattoo fonts	Specific
tattoo fonts	Next page
cursive fonts	Rewrite
cursive fonts	Next page (2x)
cursive tattoos	Rewrite
graffiti fonts	Rewrite

Session 2	Action Type
nokia n95	Initial
nokia n series	Rewrite
nokia n97 cell	Rewrite
nokia n97 cell	Next page (2x)
lg	Rewrite

User session trends

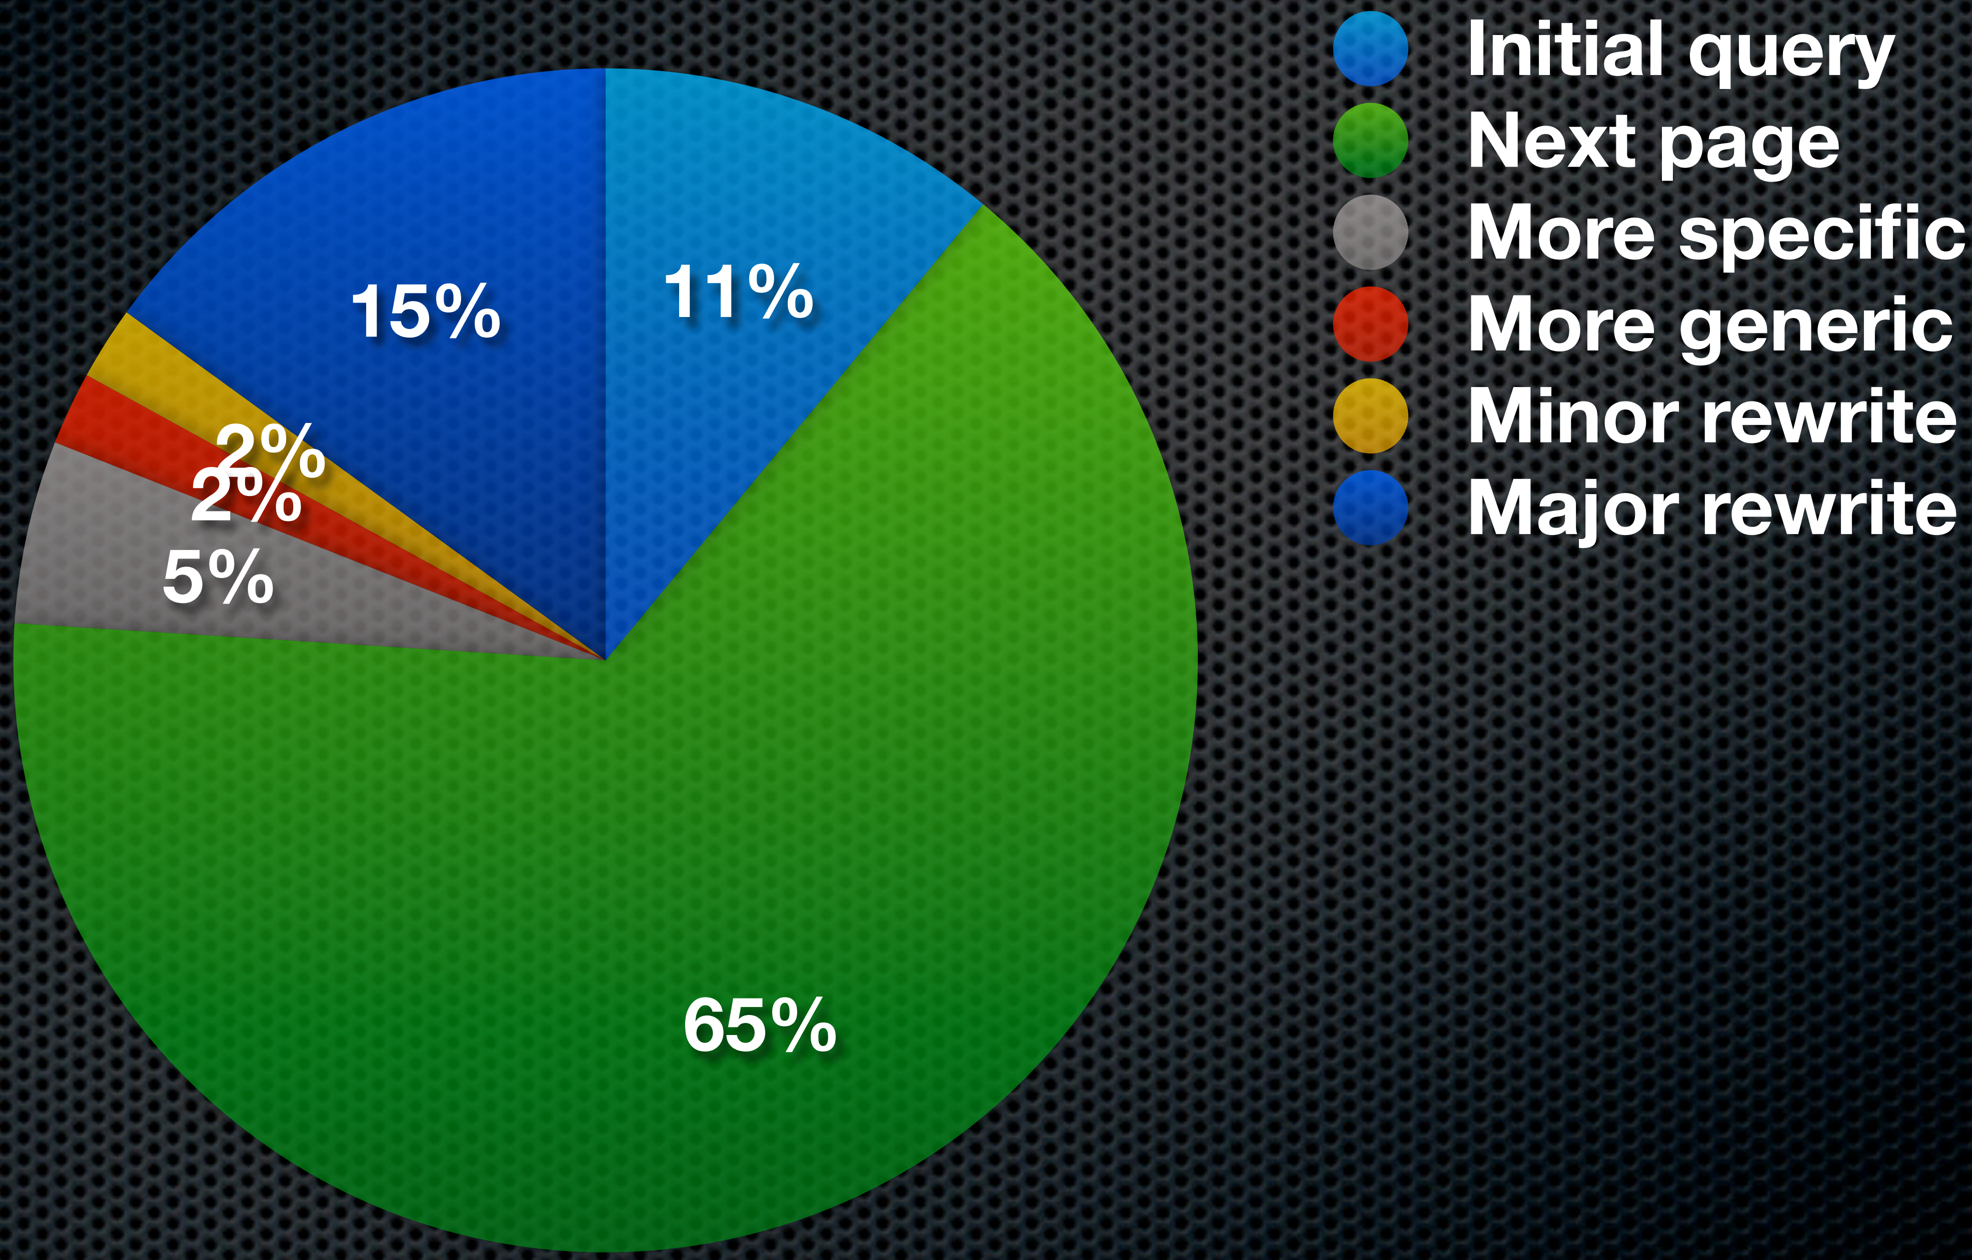

YAHOO!

What do they click on?

- Canonical image?
- Consensus of what the topic is about?

Cond. CTR @ position

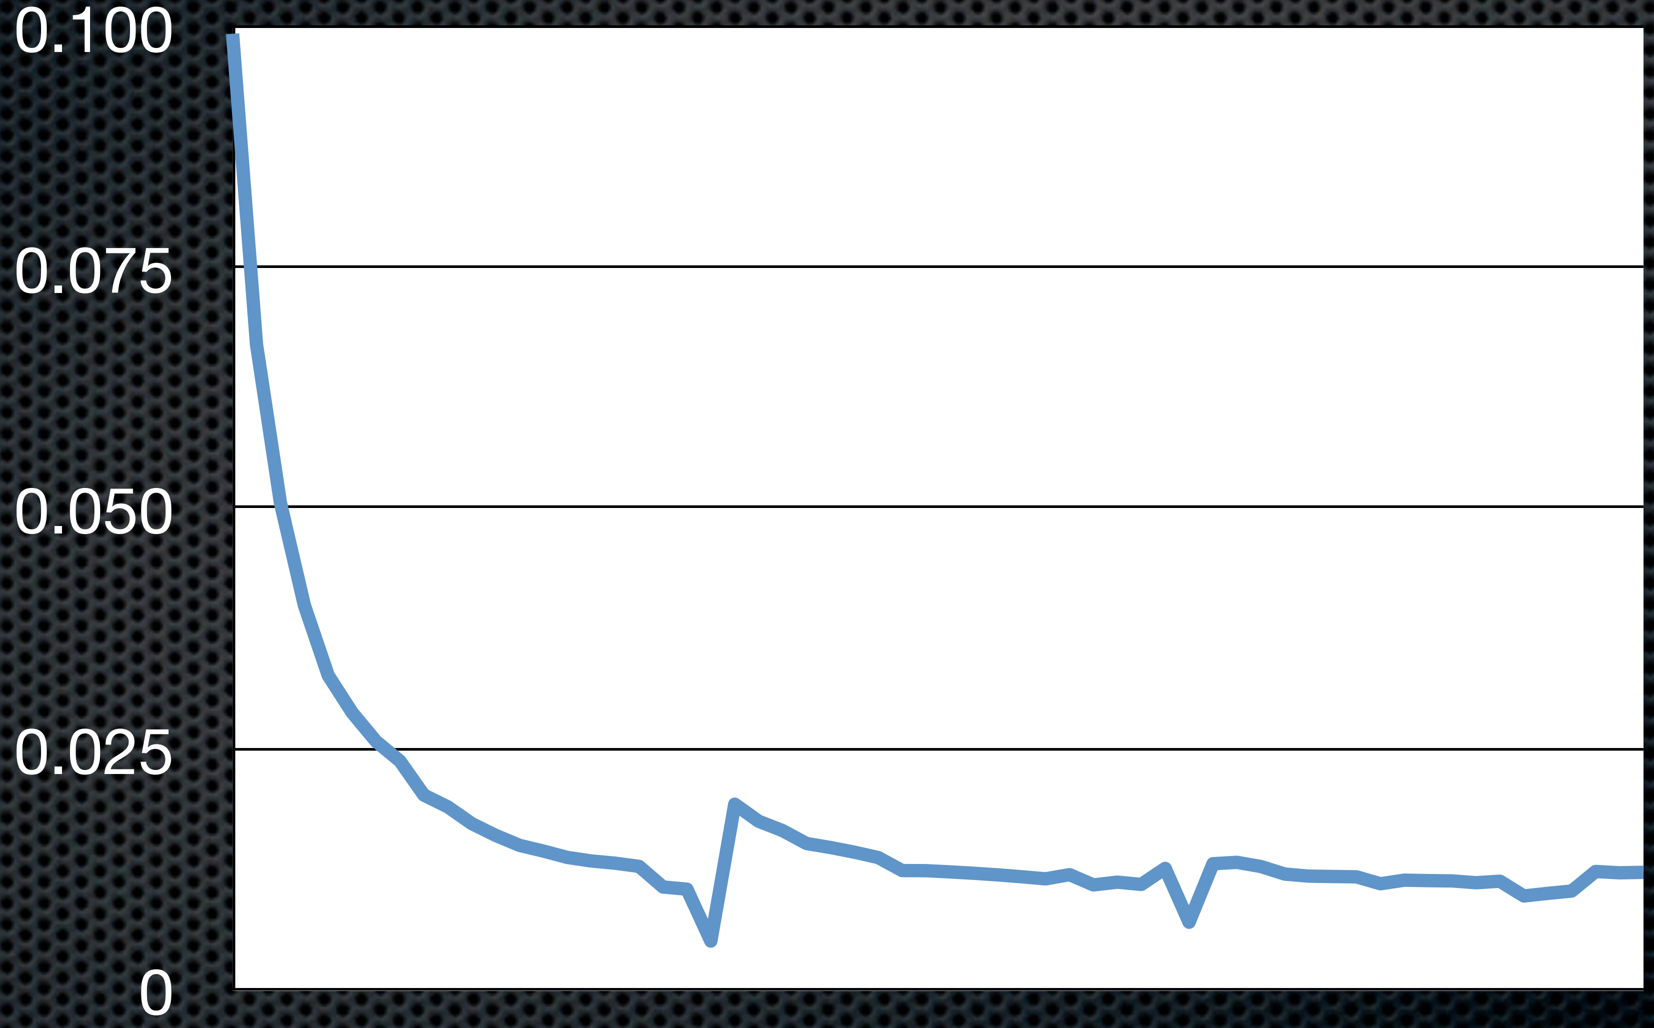

Conditional CTR

Query	Max	Avg	Stddev
Jaguar	0.097	0.020	0.019
Valentine	0.133	0.011	0.014
Johnny Depp	0.126	0.024	0.033
Pumpkin Carving	0.279	0.054	0.016

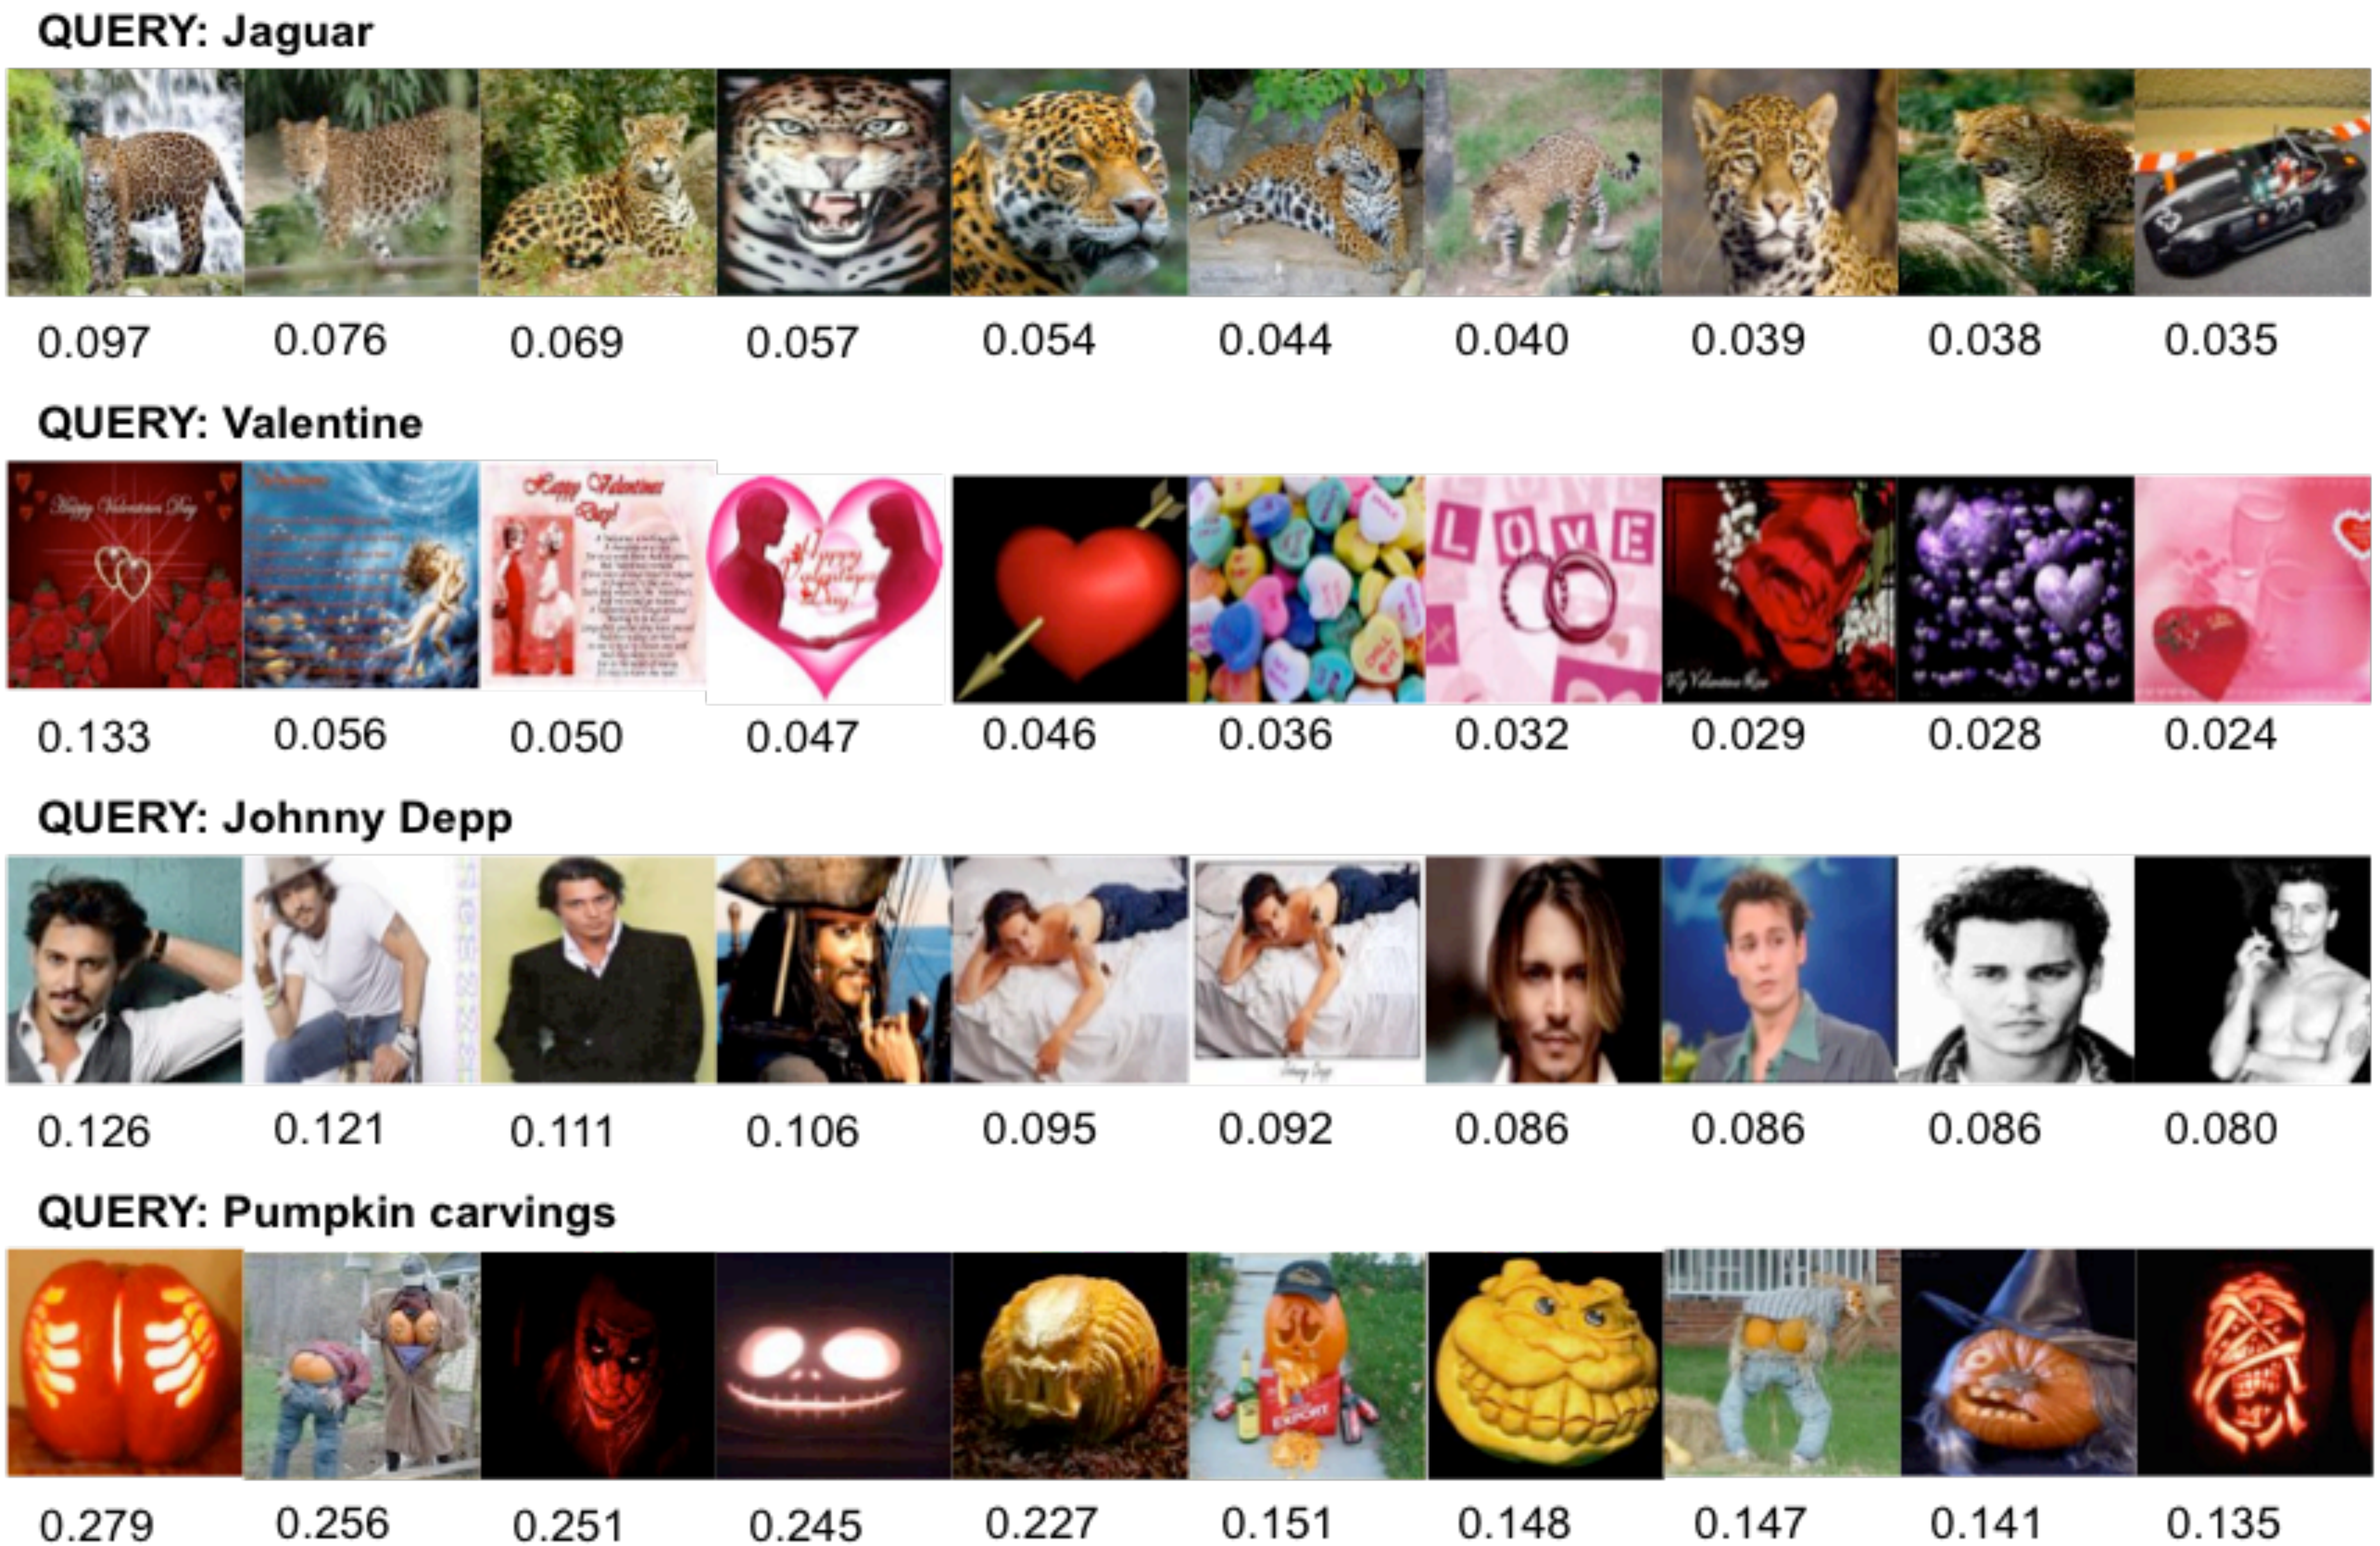
QUERY: Jaguar
0.097 0.076 0.069 0.057 0.054 0.044 0.040 0.039 0.038 0.035
QUERY: Valentine
0.133 0.056 0.050 0.047 0.046 0.036 0.032 0.029 0.028 0.024
QUERY: Johnny Depp
0.126 0.121 0.111 0.106 0.095 0.092 0.086 0.086 0.086 0.080
QUERY: Pumpkin carvings
0.279 0.256 0.251 0.245 0.227 0.151 0.148 0.147 0.141 0.135

Understanding users?

- Searching or Browsing?
- How to reform media search?

Diversity

Blending search results, composing interesting slideshows

Joint work with:

Vanessa Murdock, Lluis Garcia,

Reinier van Leuken, Malcolm Slaney,

Killian Weinberger, Ximena Olivares

YAHOO!

Dimensions of diversity

- Topical

- Visual

- Temporal, Spatial, Social

Detecting and resolving tag ambiguity

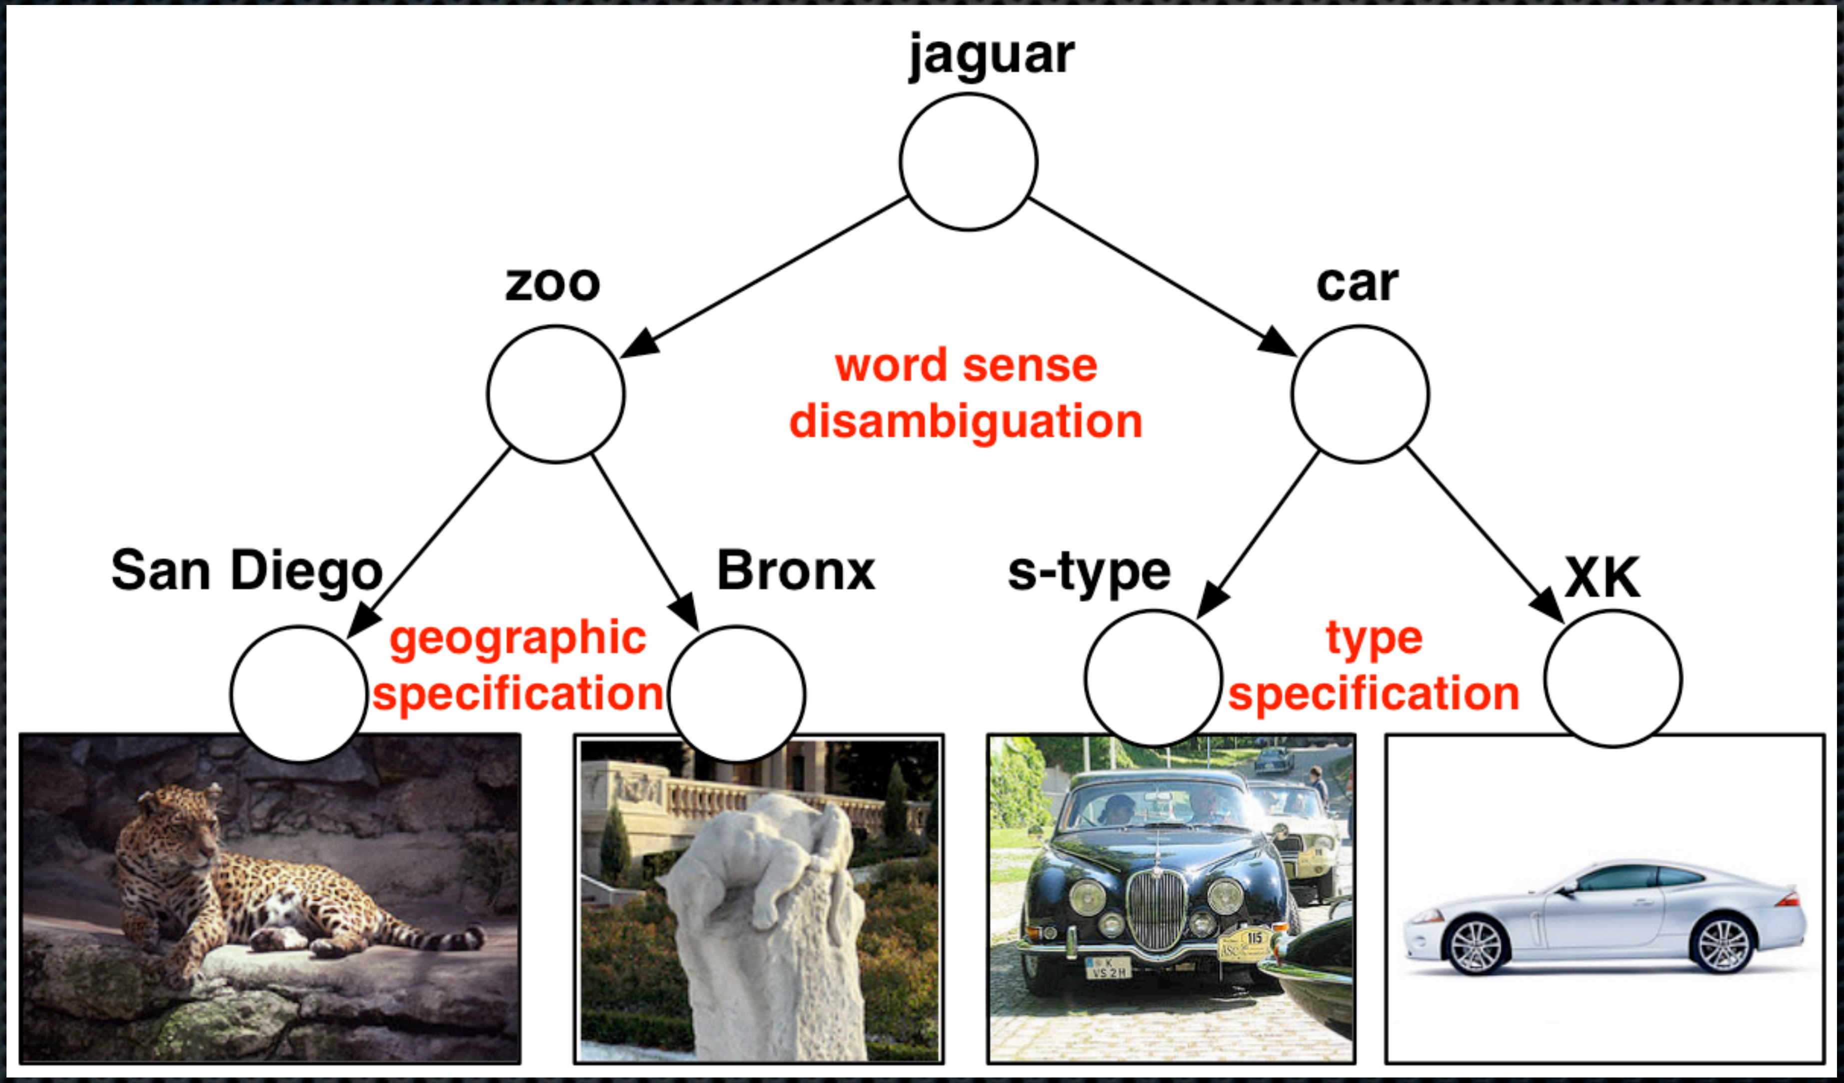

Resolving tag ambiguity

Two contributions:

1. A statistical approach is proposed to measure the ambiguity of a tag set, and the user is only interrupted, when the ambiguity score is above a certain threshold
2. The method introduces ***pairwise*** disambiguation to recommends two tags that would reduce the ambiguity of the existing tag set the most

Nature of ambiguity

- A tag set is ambiguous if it can appear in two different tag contexts
  - Geographic locations, time-based events, languages, topical, social, or any combination of the mentioned contexts (“Java”: location, programming language, coffee, etc.)
- Example: “Cambridge”
  - Considered ambiguous, based on spatial context
  - Tag suggestions: “Massachusetts” or “United Kingdom”
  - Alternative tag suggestion “university” is highly relevant, but will not resolve the ambiguity.

Flickr tag co-occurrence

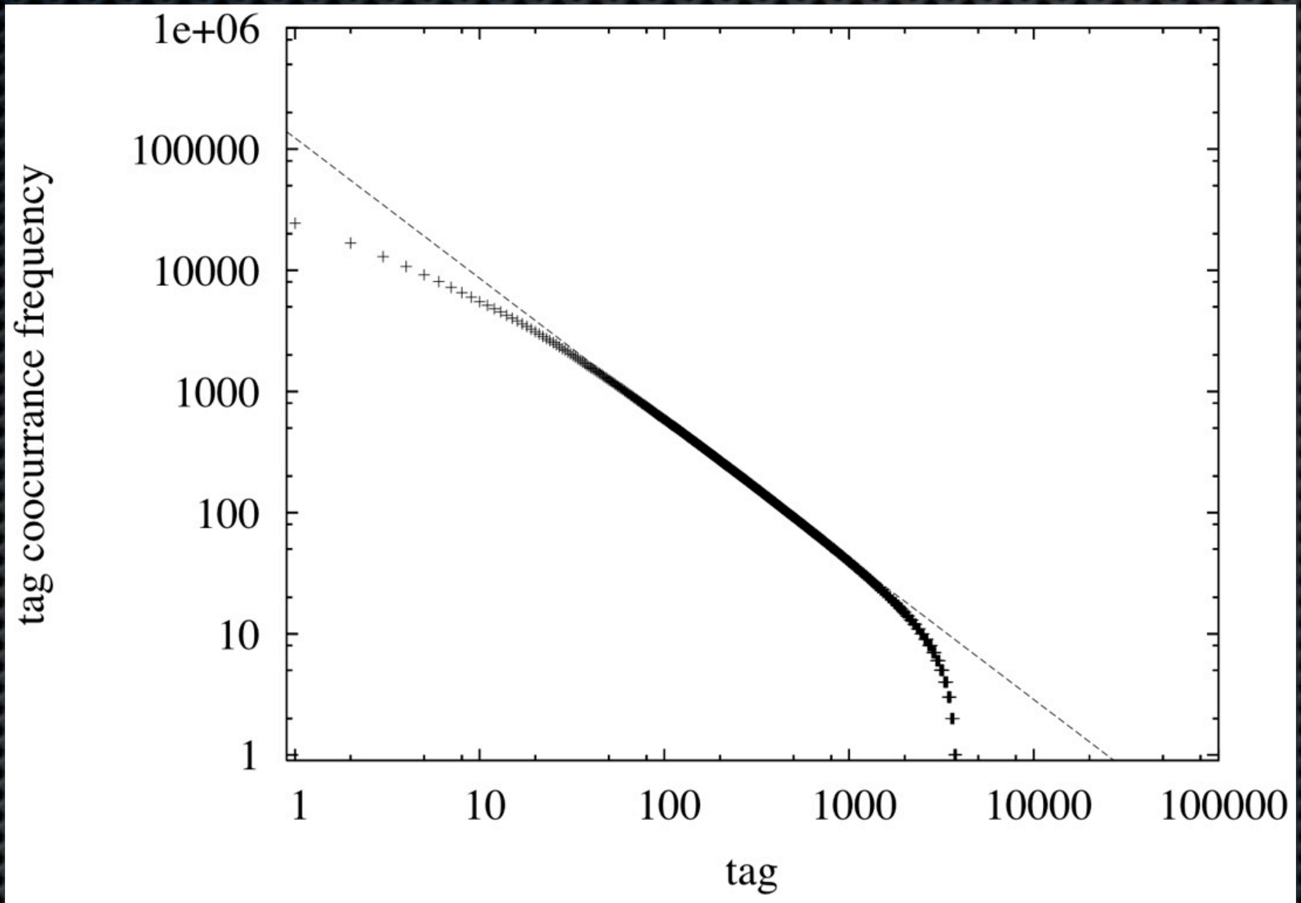

- Use Flickr tag co-occurrence stats
- Asked users if a tag-set would require additional tags to reduce ambiguity (4-point scale)

Use a weighted KL divergence for detecting pairs of tags that have the largest impact on reducing the ambiguity

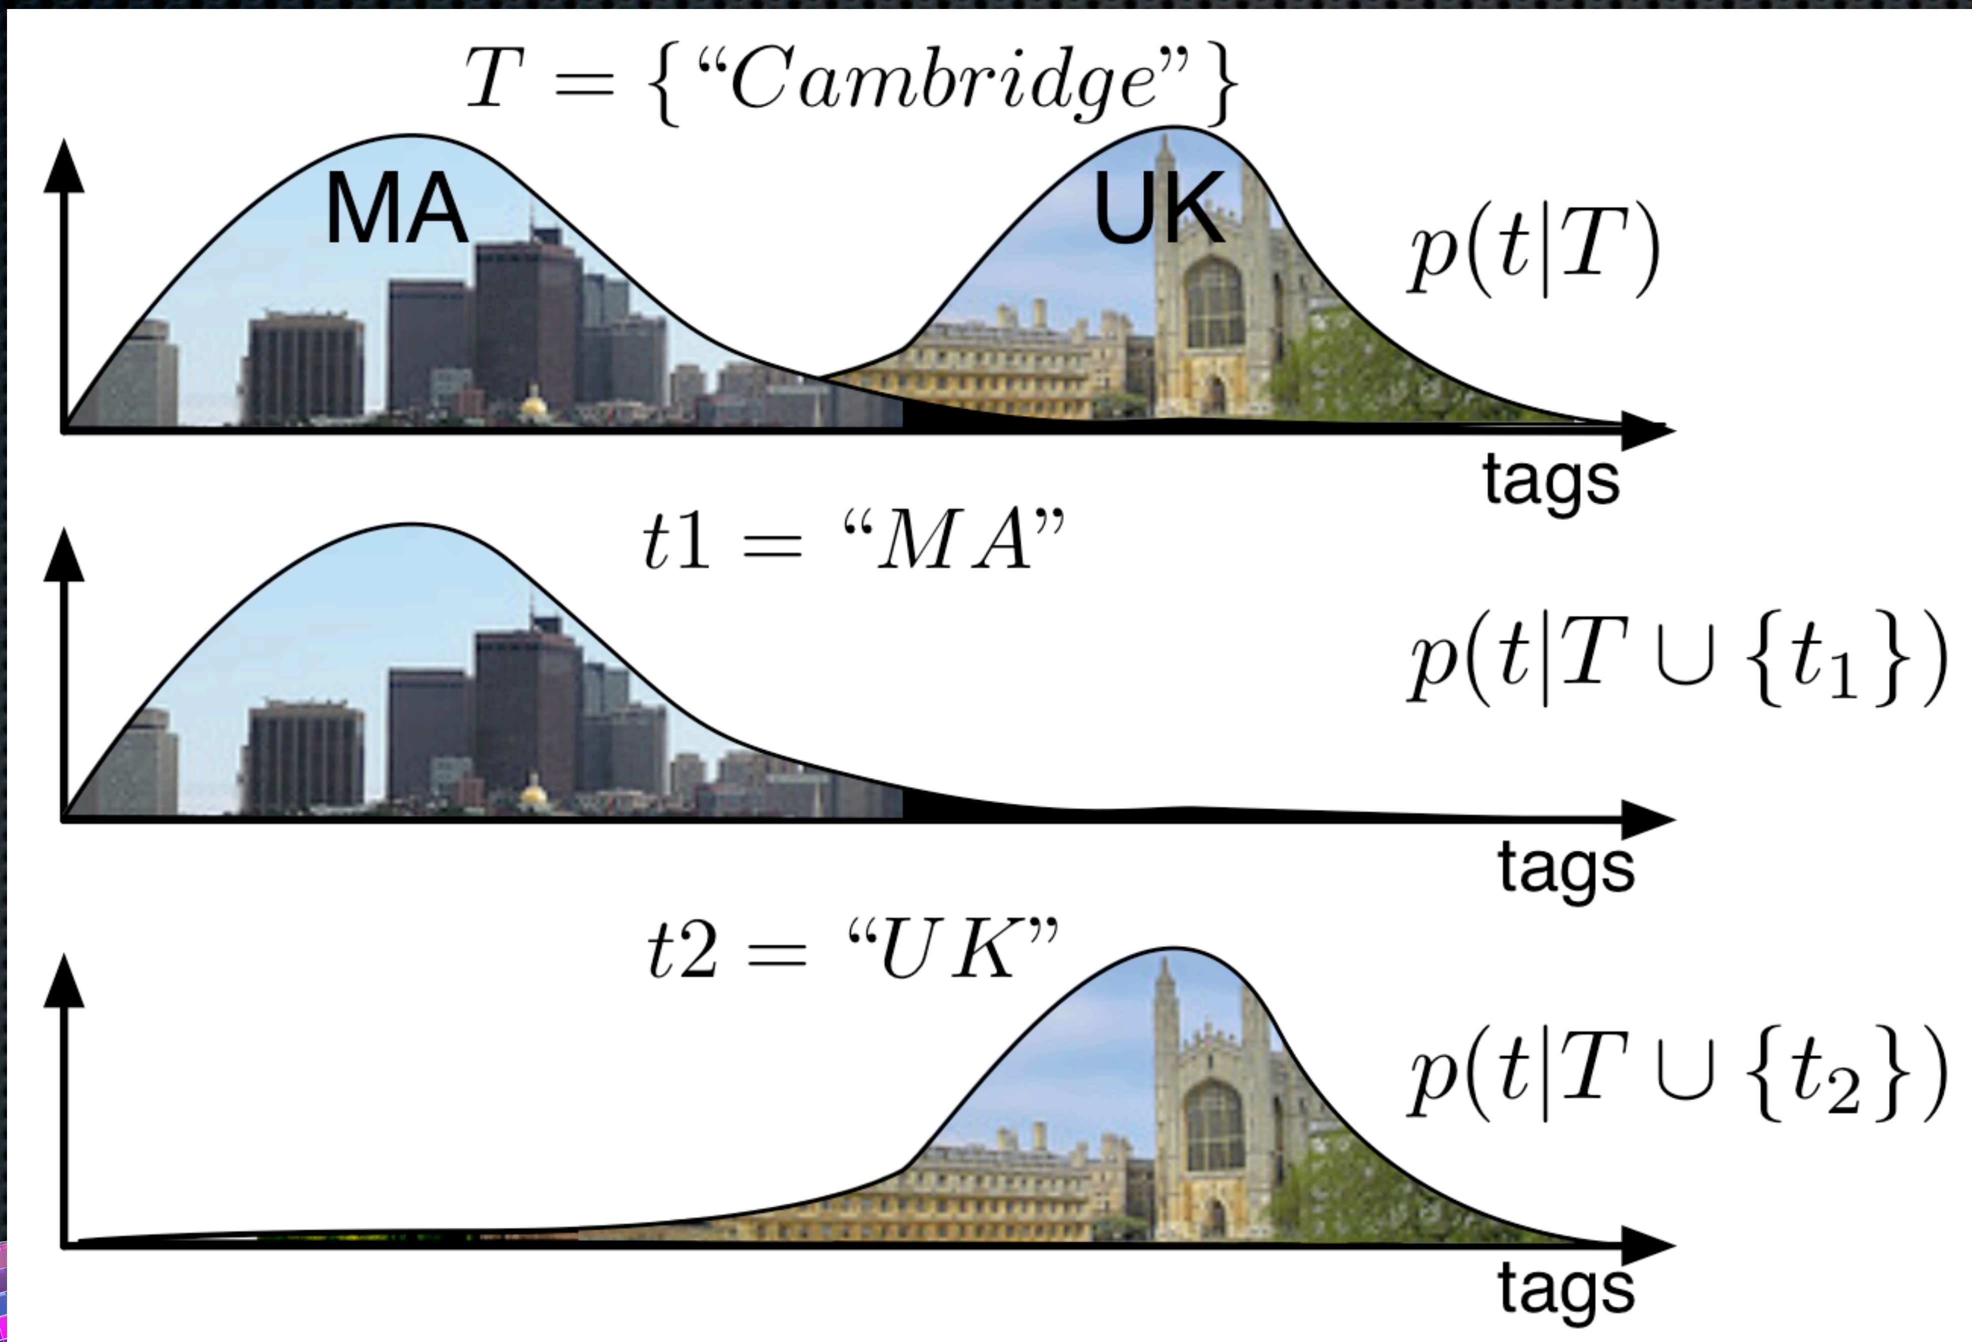

Probabilistic framework

$$p(t_i|T) = \frac{p(T|t_i)p(t_i)}{p(T)} = \frac{p(t_i)\prod_{t \in T} p(t|t_i)}{\sum_j p(t_j)\prod_{t \in T} p(t|t_j)}$$

$$KL(t_i||t_j) = \sum_t p(t|T \cup \{t_i\}) \log\left(\frac{p(t|T \cup \{t_i\})}{p(t|T \cup \{t_j\})}\right)$$

$$\bar{KL}(t_i, t_j) = KL(t_i||t_j) + KL(t_j||t_i)$$

$$div(t_i, t_j) = p(t_i|T)p(t_j|T)g(\bar{KL}(t_i, t_j)),$$

Optimization

- Optimize trade-off between tag probabilities and the KL divergence by finding the function, g, that gives the best fit to human judgments of ambiguity. We restrict ourselves to functions of the form, g(x) = x^p
  - max. Pearson correlation: P=2,3,4
- Reduction in computation cost is possible by only considering top 30 co-occurring tags:

Ambiguity explained

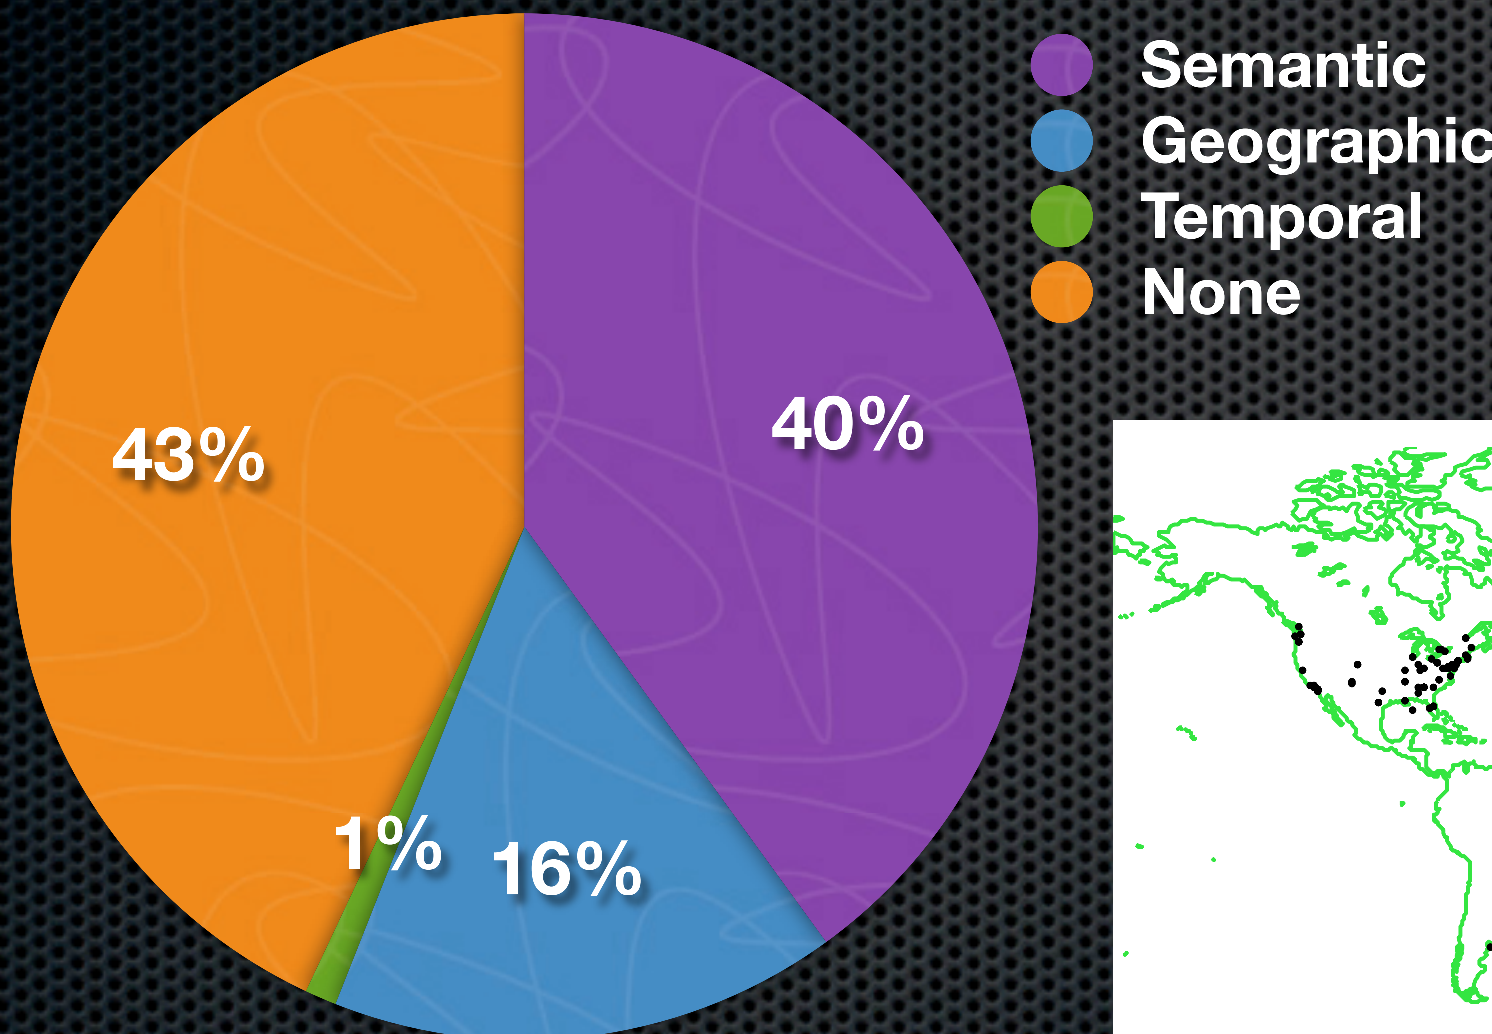

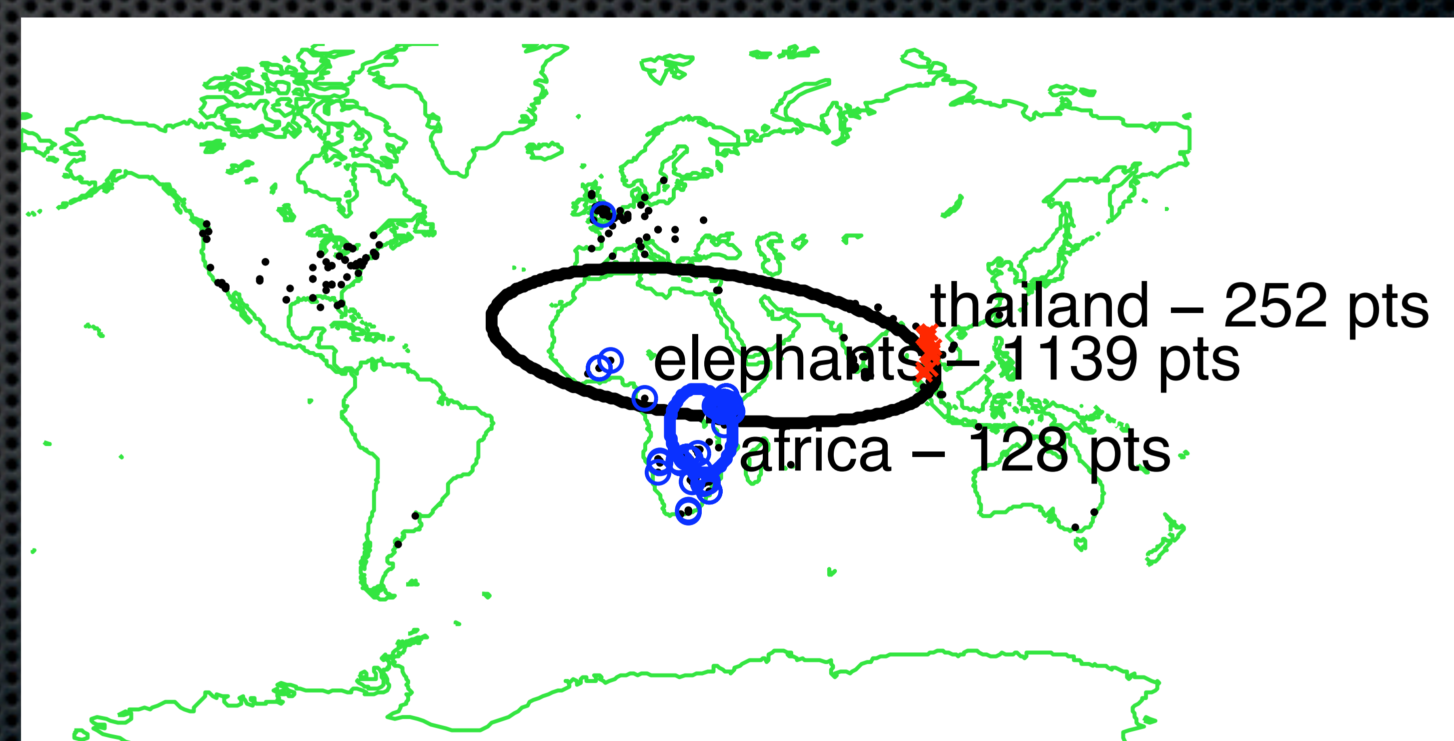

Visual Diversity

Need for visual diversity

- **Applications: SERP, Slideshows**

- Both rely on textual information associated with an image
- Textual information lacks discriminative power to deliver visually diverse search results
- “Limited” query formulation power

Towards visual diversification

- Dynamic weighting of visual features
  - To capture the discriminative aspects of a set of images

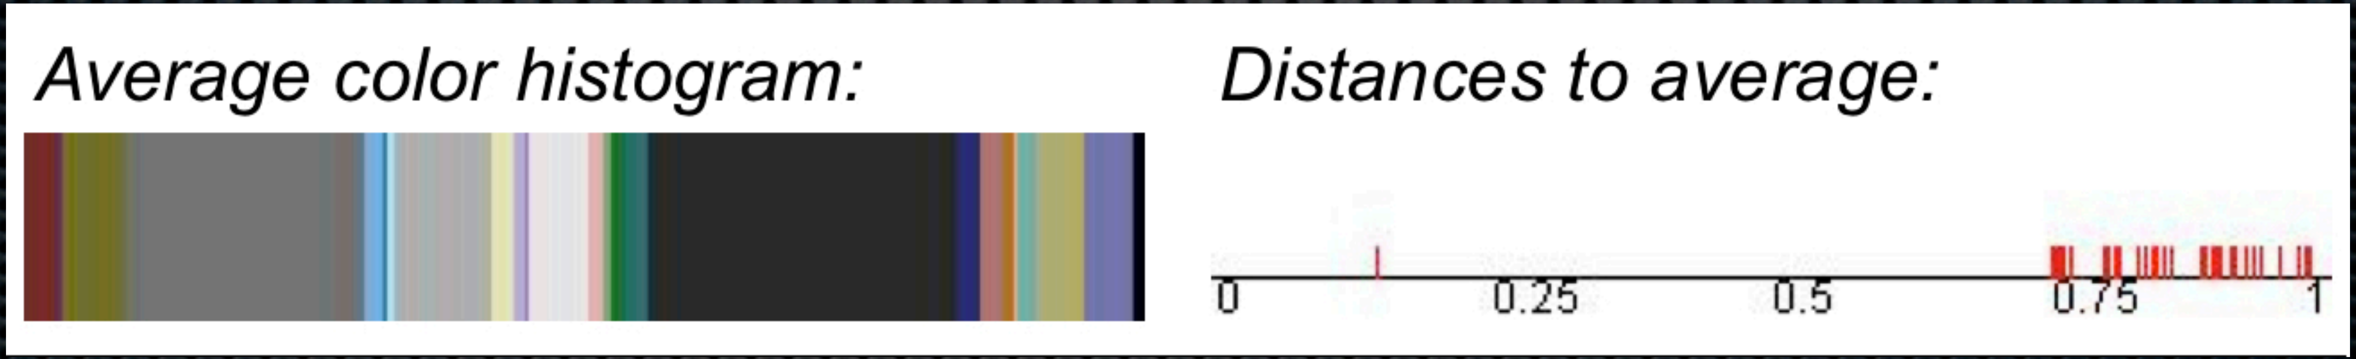

- Methods for visual diversification of image search results
  - Post-retrieval step
  - We assume relevance of images retrieved is good.

Towards visual diversification

- Deploy lightweight clustering methods
  - Folding -- obey original clustering
  - Reciprocal election -- images cast votes for other images to be its representative.

Visual characteristics

- In context of a set of images, the relative importance of the different visual features is a-priori undefined
- Depends on the characteristics of the images in the set
- Jaguar:

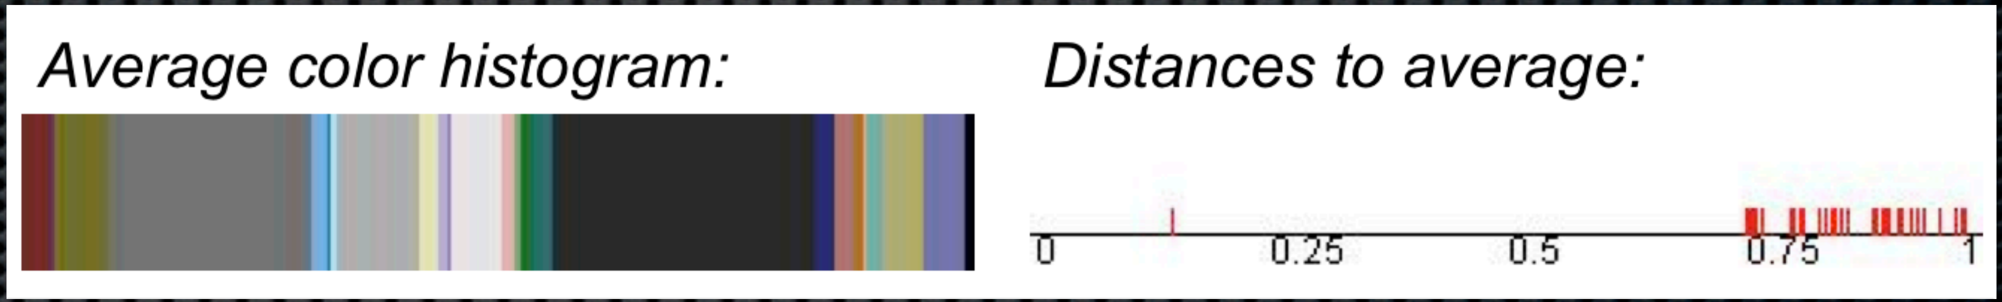

- Fireworks:

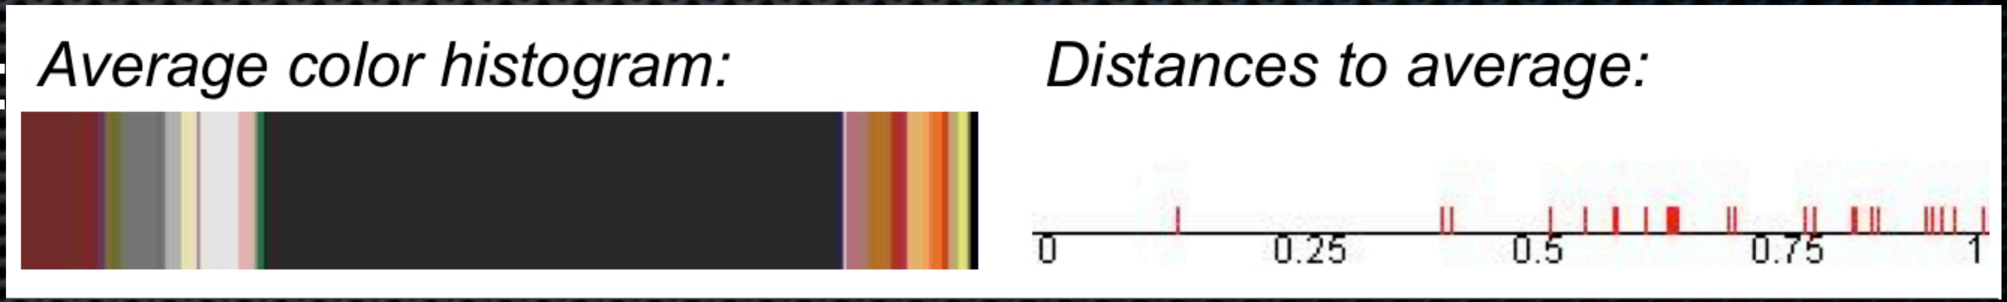

Visual characteristics

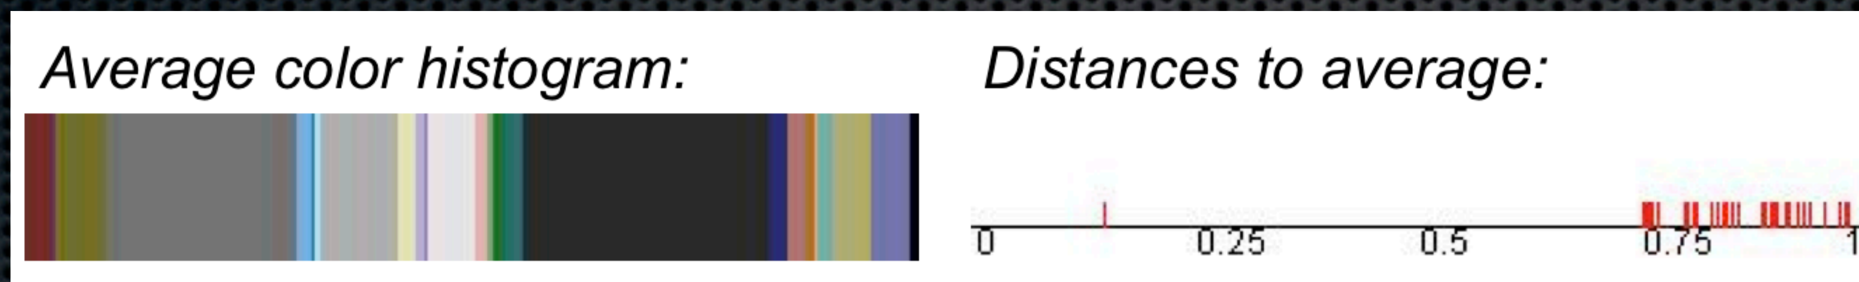

Distance between a and b according to i^{th} feature

$$d(a,b) = \frac{1}{f}\sum_{i=0}^{f}\frac{1}{{\sigma_i}^2}\,d_i(a,b)$$

Total number of features

Variance of distances according to i^{th} feature

Notation (image clustering)

- A set of images search results I contains n images
- I can be stored in:
  1. A ranked list $L=L_1,L_2,\ldots,L_n$, with decreasing relevance
  2. An unorder set $S=S_1,S_2,\ldots,S_n$
- Methods
  1. Input: L or S
  2. Output: a clustering C (partitioning of I)
  3. Images divided over K clusters: $C_1, C_2,\ldots,C_k$, with:

  4. One image is declared cluster representative R_k
  5. All representatives together form the set R

  7. Parameter free -- threshold is set dynamically

YAHOO!

Folding

Algorithm 1 Folding

Input: Ranked list L of I

Output: Clustering C

```
Let the image L_1 be the first representative R_1
for Each image L_i do
    if d(L_i, R_j) > ε(*) for all representatives R_j then
        add L_i to the set of representatives R
for Each image L_i ∉ R do
    Find representative R_j that is closest to L_i
    Assign L_i to the cluster of R_j
```

(*)ϵ is defined as the mean distance all images have to the *average image* in I

Reciprocal election

Algorithm 3 Reciprocal election

Input: Set S containing I, parameter m

Output: Clustering C

```
Initialize Votes map V[0, ..., k] = 0, ..., 0
for Each image i in S do
    Rank S into L_i based on visual similarity to i
    for Each image j in L_i do
        V[j] += 1/r, where r is the rank of j in L_i
while V is not empty do
    Let R_i be the item with the highest score in V
    Remove R_i from V
    Initialize new cluster C with representative R_i
    for All items s in V do
        if R_i is in top-m of L_s then
            add s to cluster C
            remove s from V
```

Evaluation

Human assessments:

- 8 independent, unbiased assessors
- Task: “cluster images for a given topic into clusters, based on visual characteristics”.
- Assessment tool:
  1. Select topic, and inspect the top 50 results during >1 minute
  2. Assign each image to a cluster (max. 20 clusters, undo last action)
  3. Label each cluster, and select cluster representative
- 200 human clusterings collected.
- Inter-assessor variability provides baseline for algos

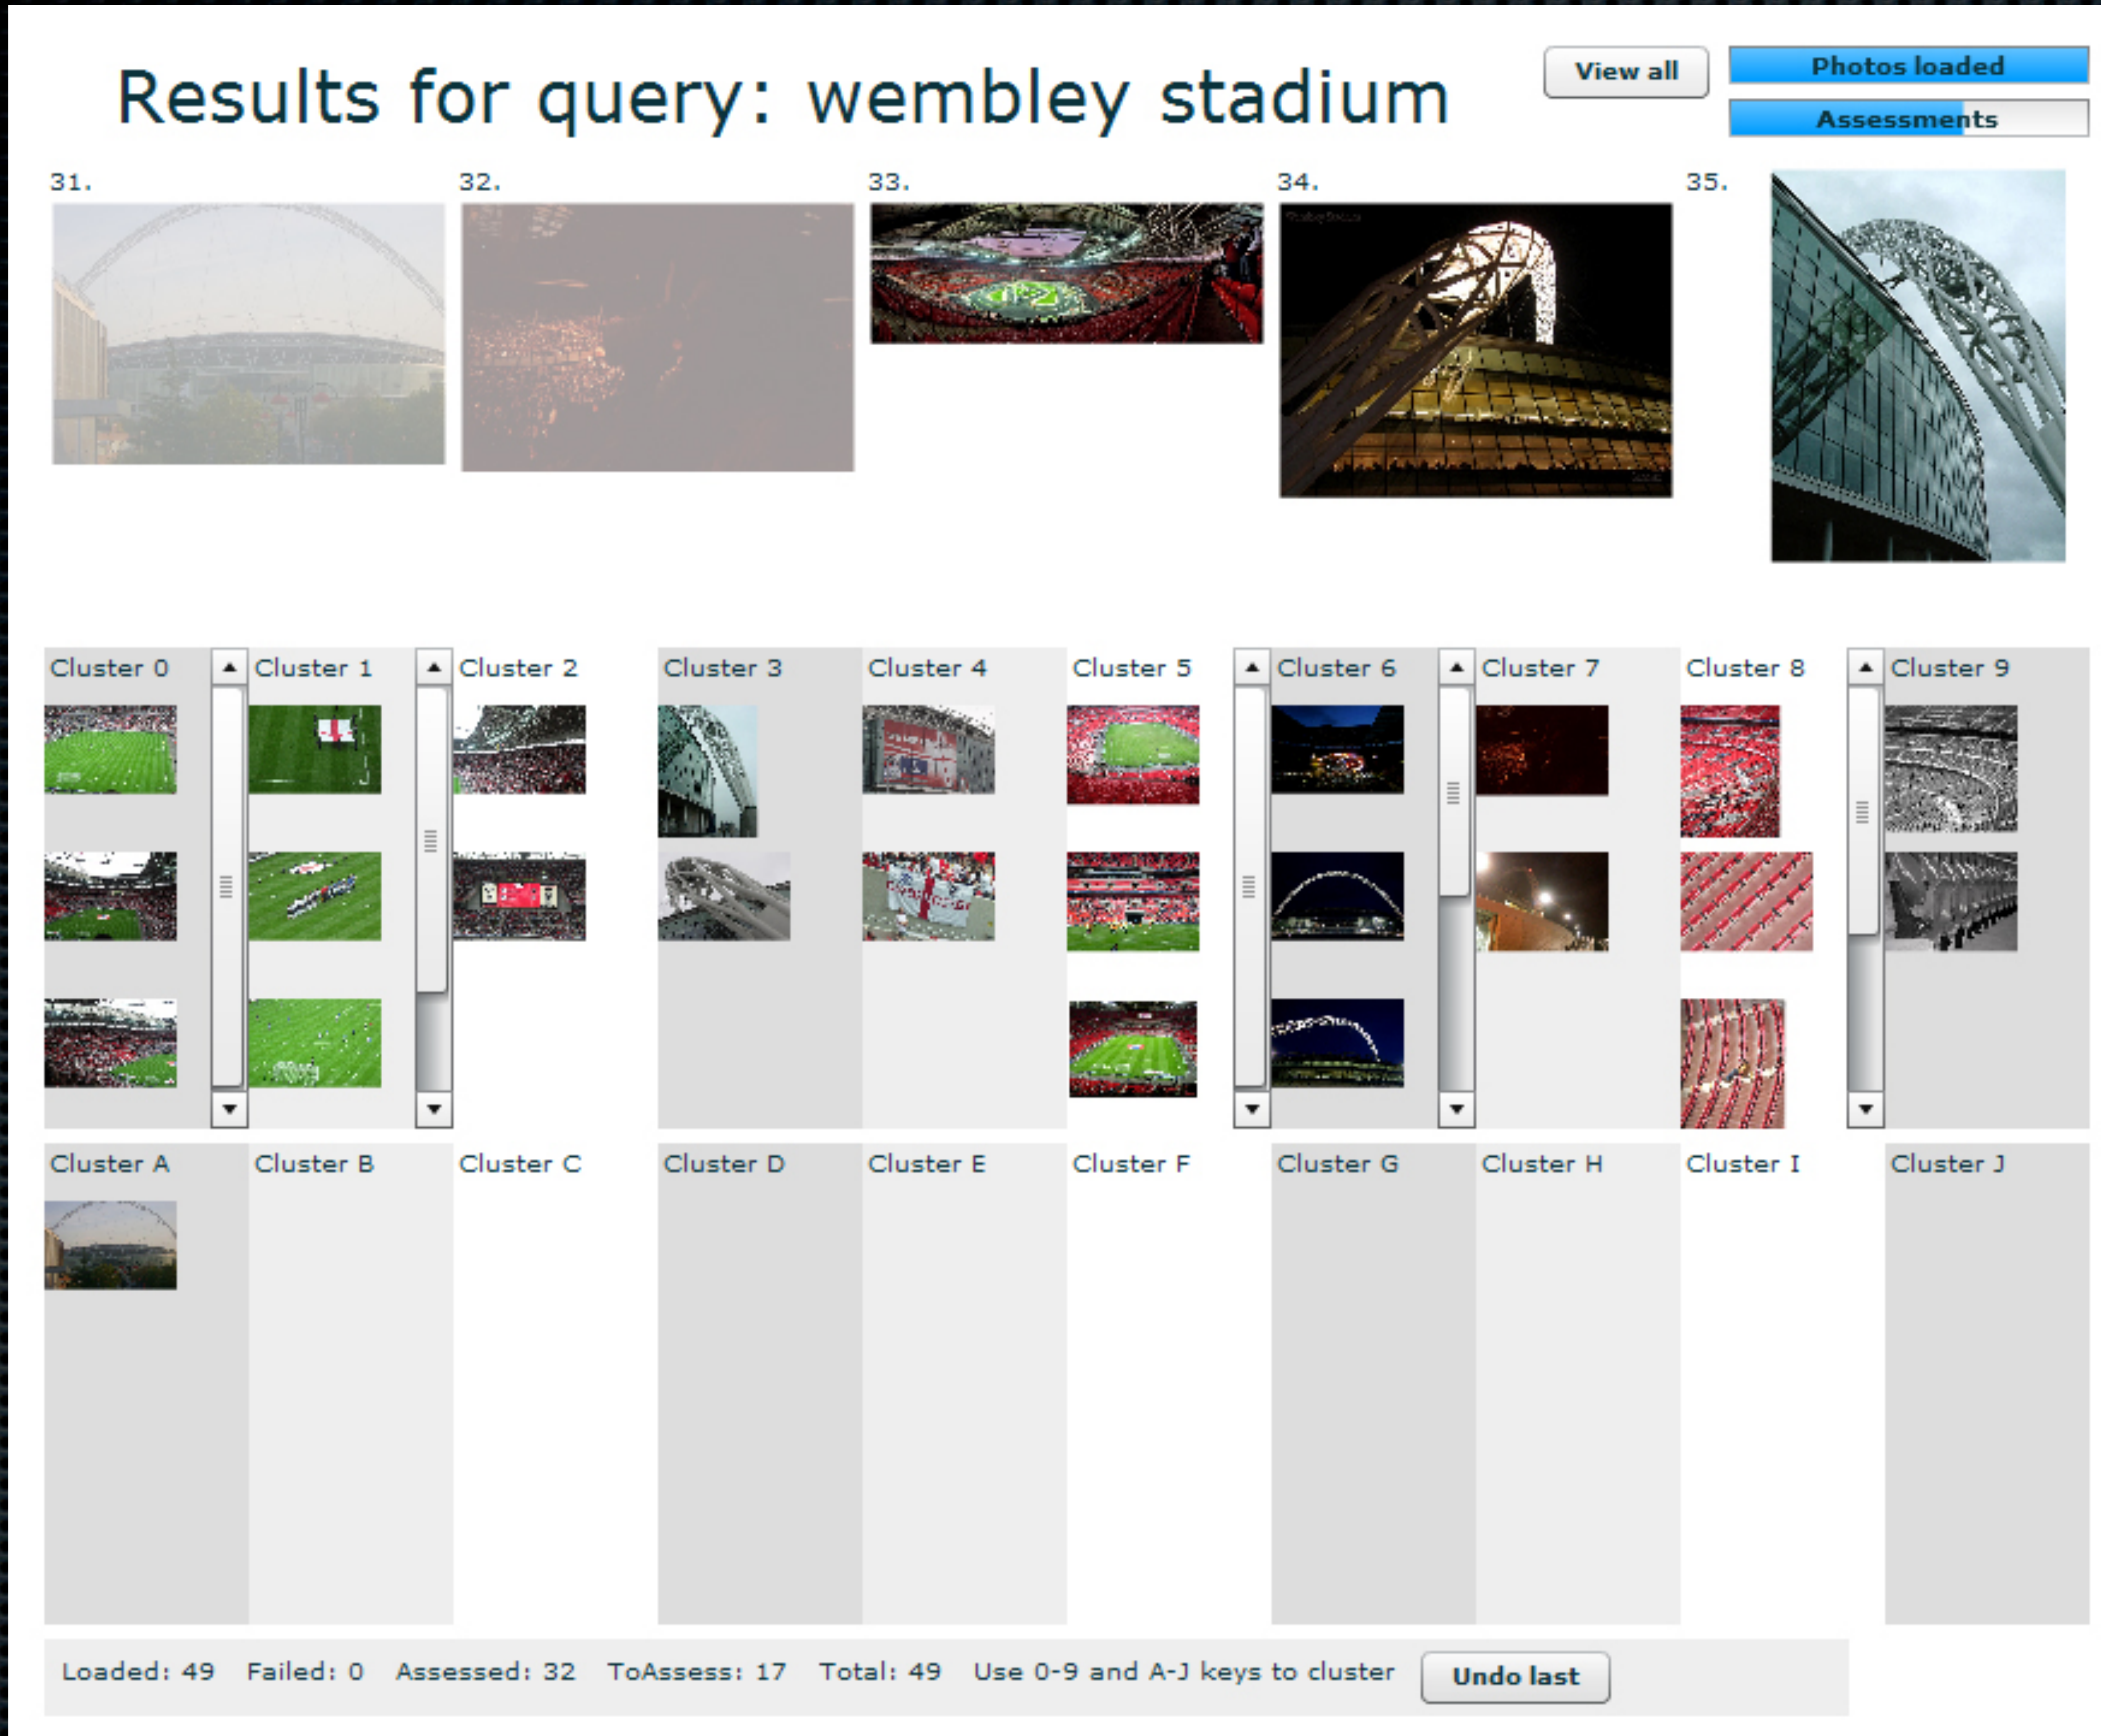
Results for query: wembley stadium
View all
Photos loaded
Assessments
31.
32.
33.
34.
35.
Cluster 0
Cluster 1
Cluster 2
Cluster 3
Cluster 4
Cluster 5
Cluster 6
Cluster 7
Cluster 8
Cluster 9
Cluster A
Cluster B
Cluster C
Cluster D
Cluster E
Cluster F
Cluster G
Cluster H
Cluster I
Cluster J
Loaded: 49 Failed: 0 Assessed: 32 ToAssess: 17 Total: 49 Use 0-9 and A-J keys to cluster
Undo last

Evaluation criteria

- Objective: Compare the quality of (two) clusterings
- Given a set of images I and two clusterings C and C’:
  1. N11: image pairs in same cluster under both C and C’
  2. N00: image pairs in different cluster under both under C and C’
  3. N10: image pairs in same cluster in C but not in C’
  4. N01: image pairs in same cluster in C’ but not in C
- Fowlkes-Mallows Index:
  1. Clustering equivalent of precision/recall
  2. High score on FM index indicates cluster similarity
- Variation of Information:
  1. Measures the difference in the relationship between a point and a cluster over two clusterings
  2. Low score on VI indicates cluster similarity

YAHOO!

Results

- Performance Bounds:
  - Upper-bound: Inter assessor agreement
  - Lower-bound: Random clustering
- Overall performance (over all topics):

	Inter assessor	Random	Folding	Reciprocal election
FM index	0.419	0.139	**0.282**	0.25
VI	1.463	2.513	2.081	**1.975**

Ambiguous vs un-ambiguous - FM index

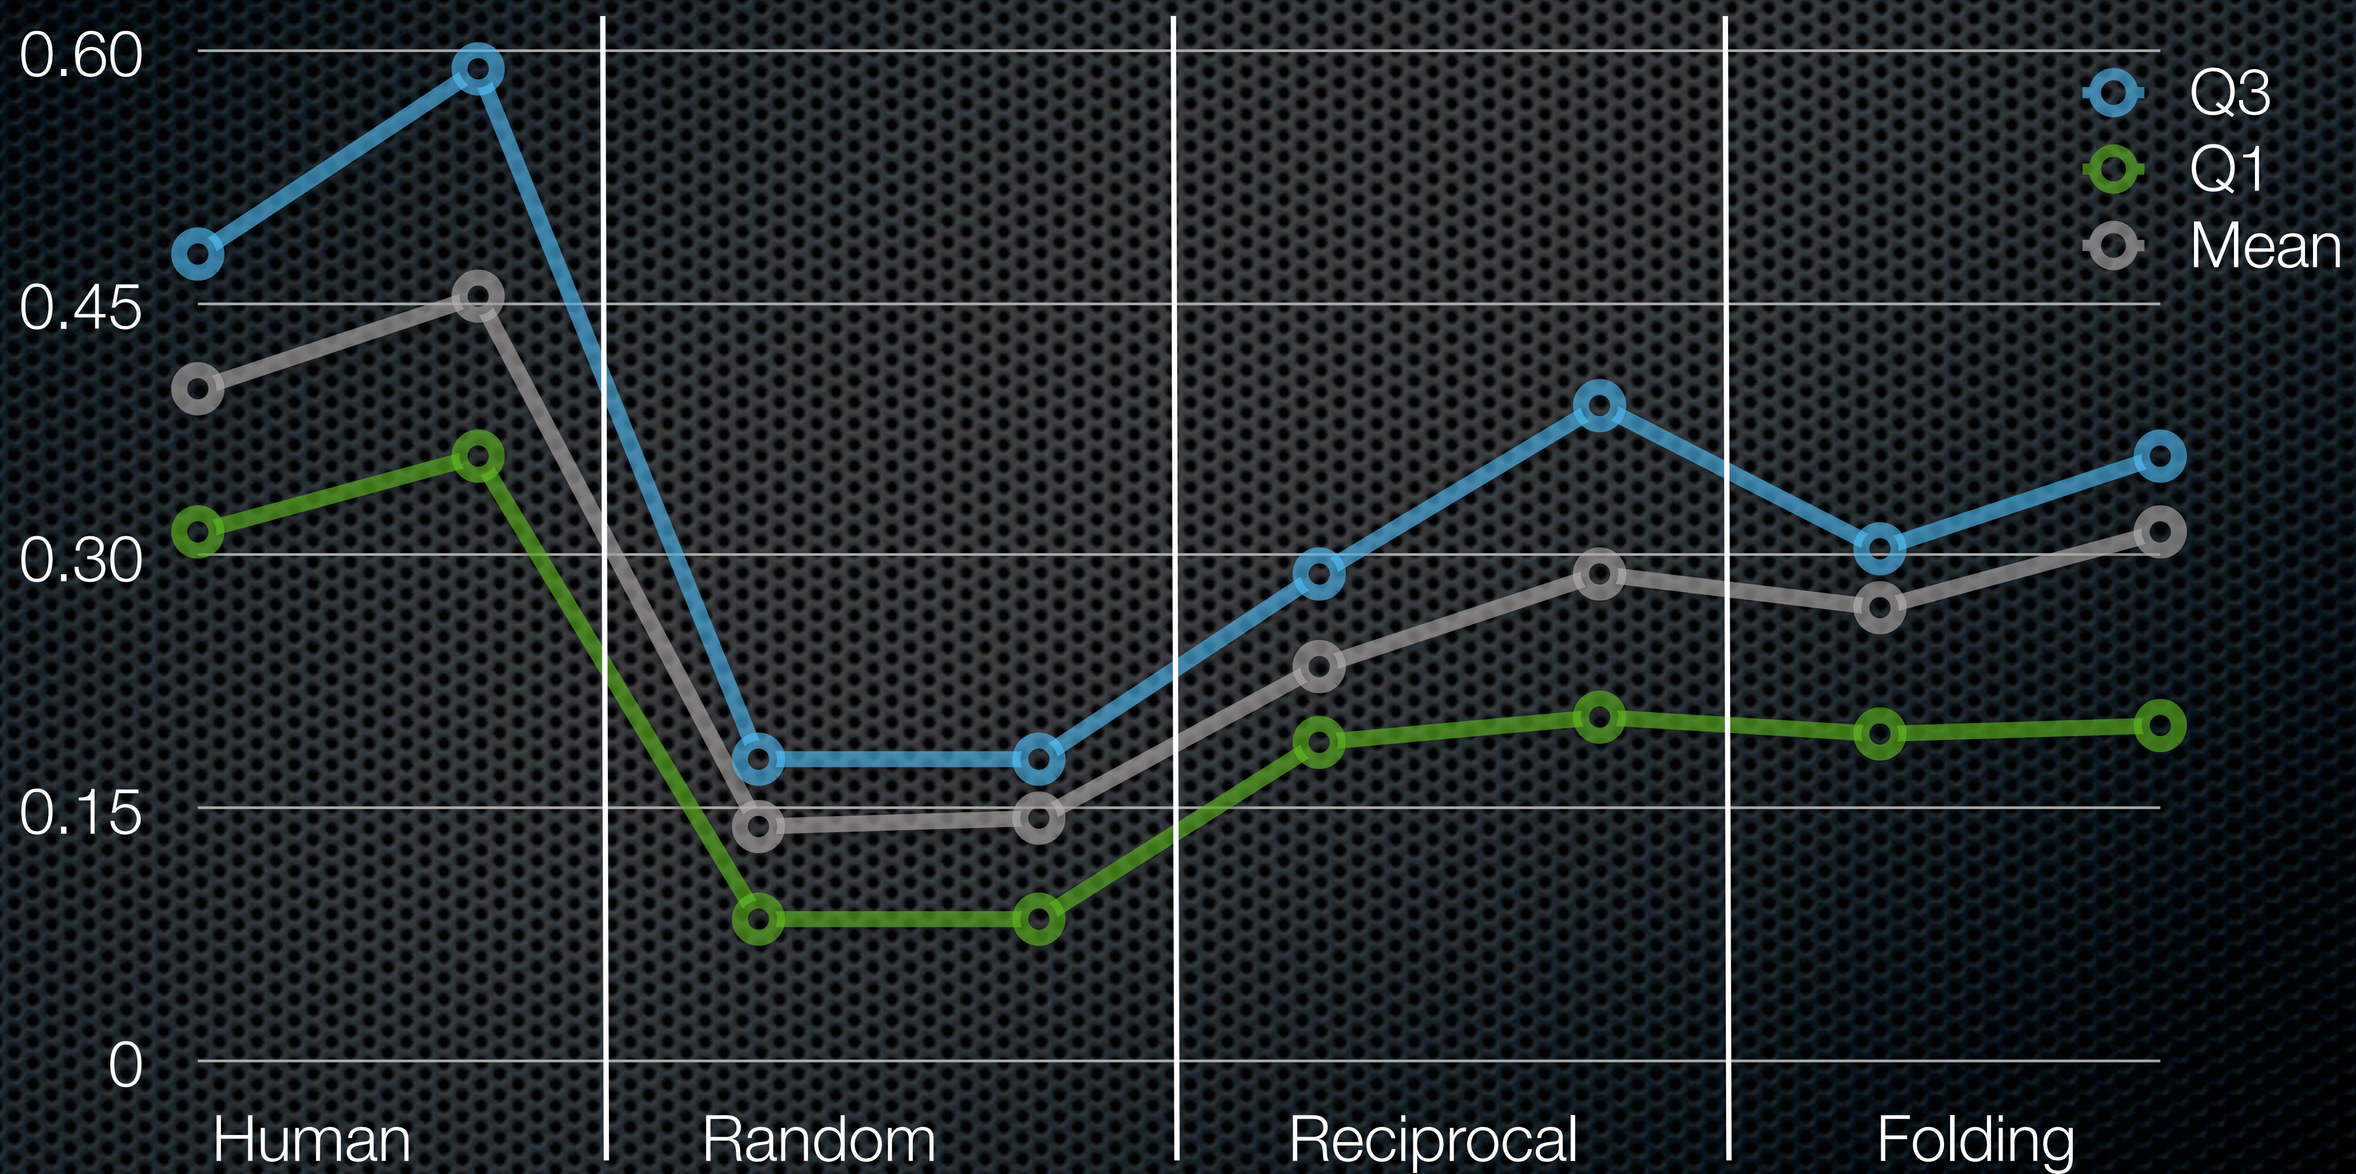

Freshness, Trending, Facets, Slideshows, and Explore

http://images.search.yahoo.com/

Image Fingerprinting

Finding near-duplicates on the Web, and in image slideshows

DCT fingerprint

- Examples of duplicate images on the Web

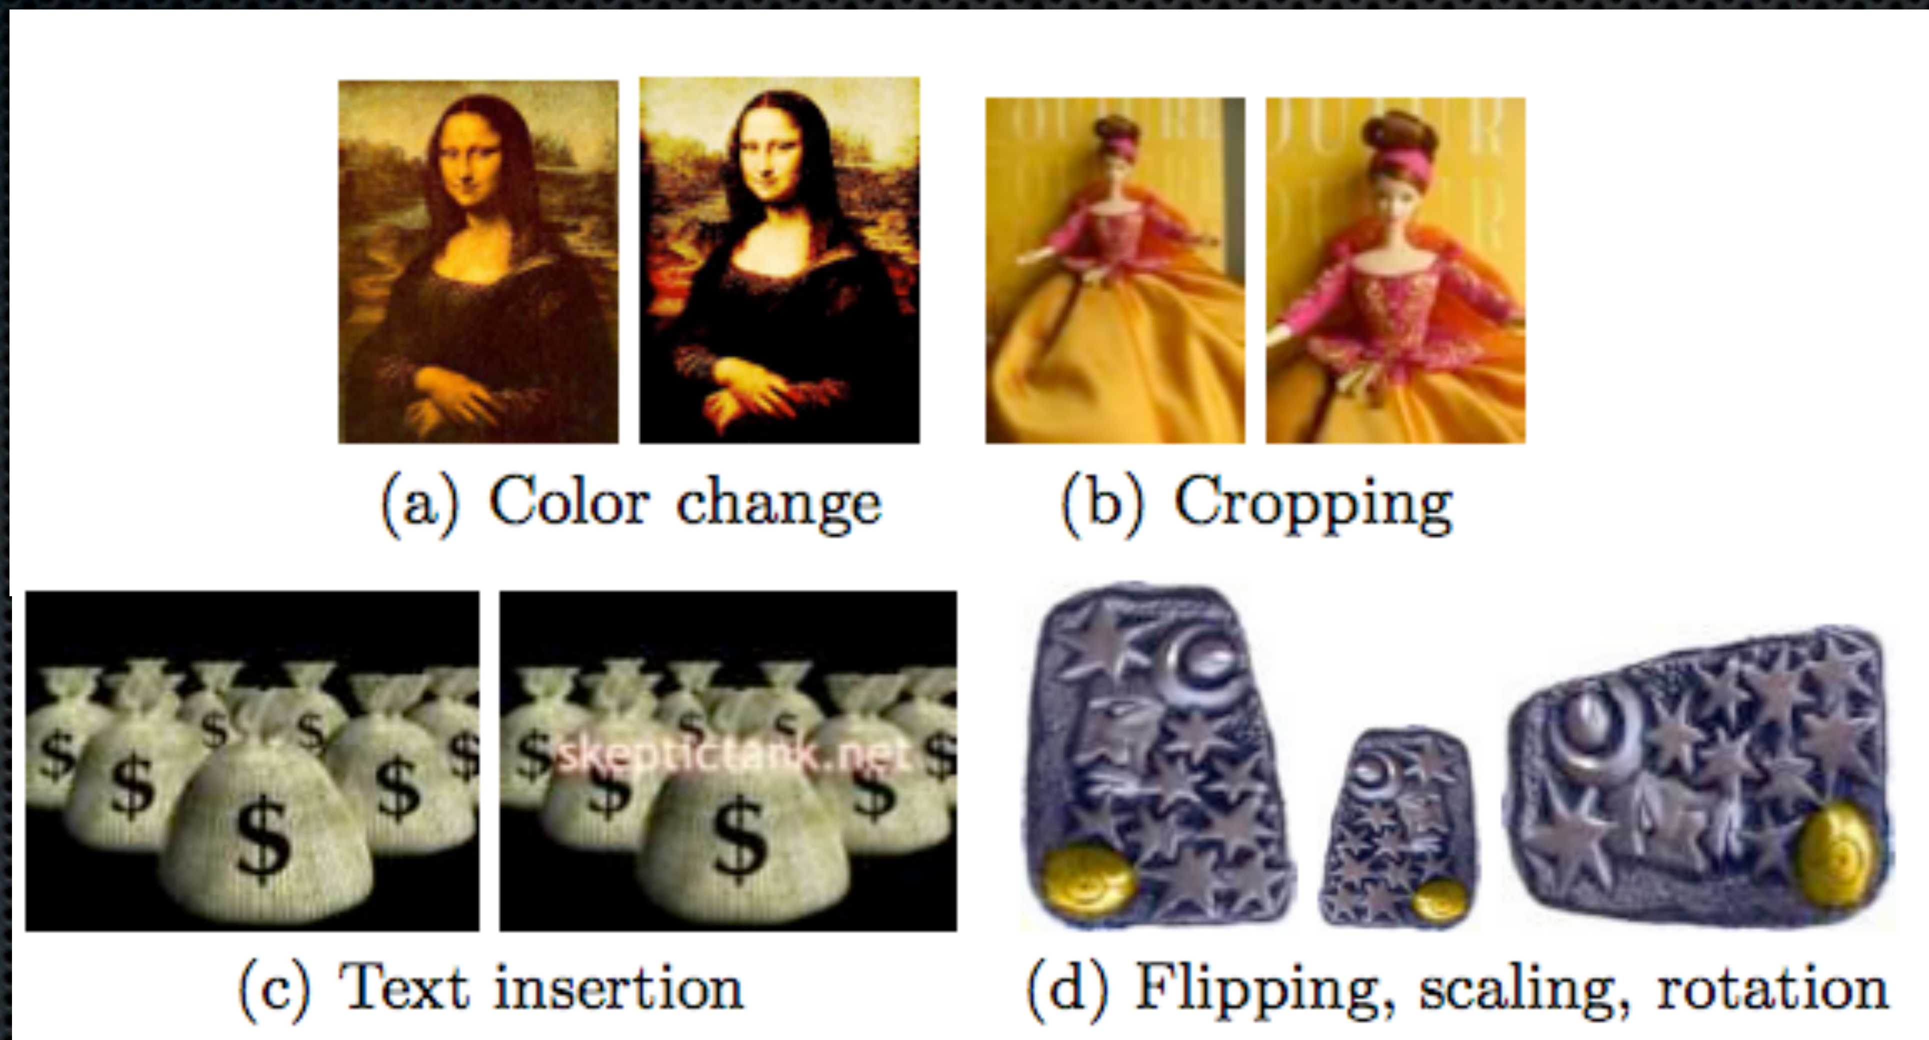

tineye.com

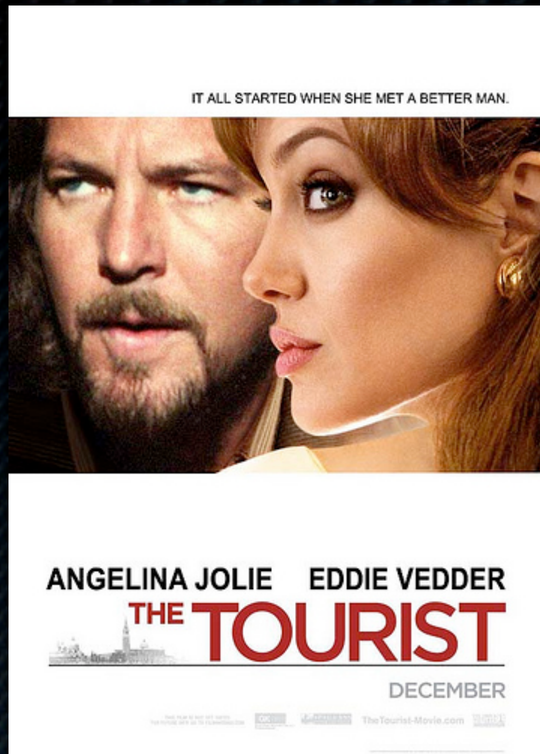
IT ALL STARTED WHEN SHE MET A BETTER MAN.
ANGELINA JOLIE EDDIE VEDDER
THE TOURIST
DECEMBER

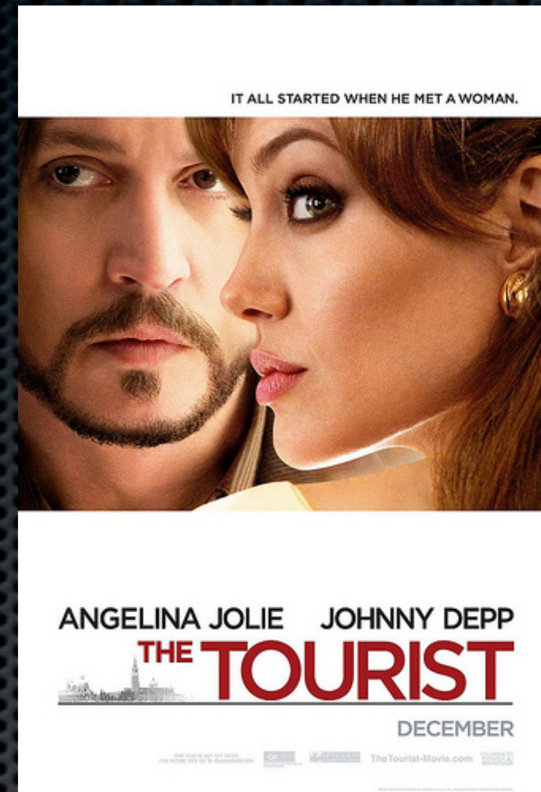
IT ALL STARTED WHEN HE MET A WOMAN.
ANGELINA JOLIE JOHNNY DEPP
THE TOURIST
DECEMBER

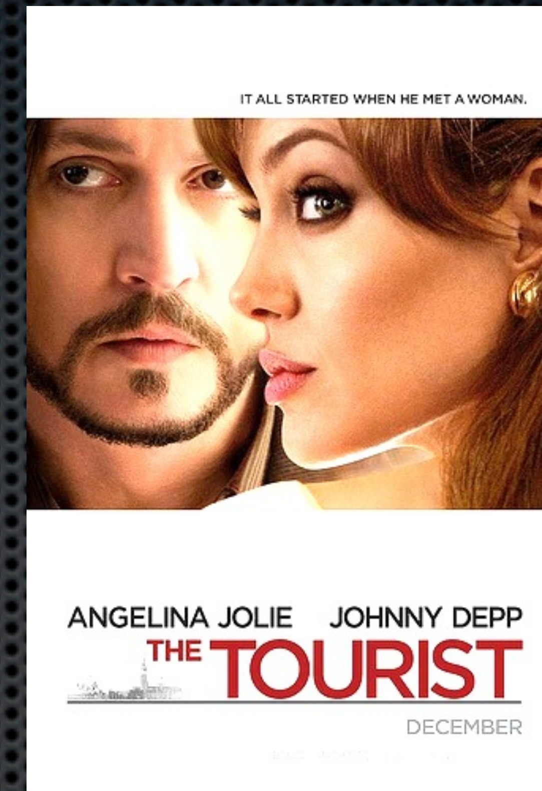
IT ALL STARTED WHEN HE MET A WOMAN.
ANGELINA JOLIE JOHNNY DEPP
THE TOURIST
DECEMBER

YAHOO!

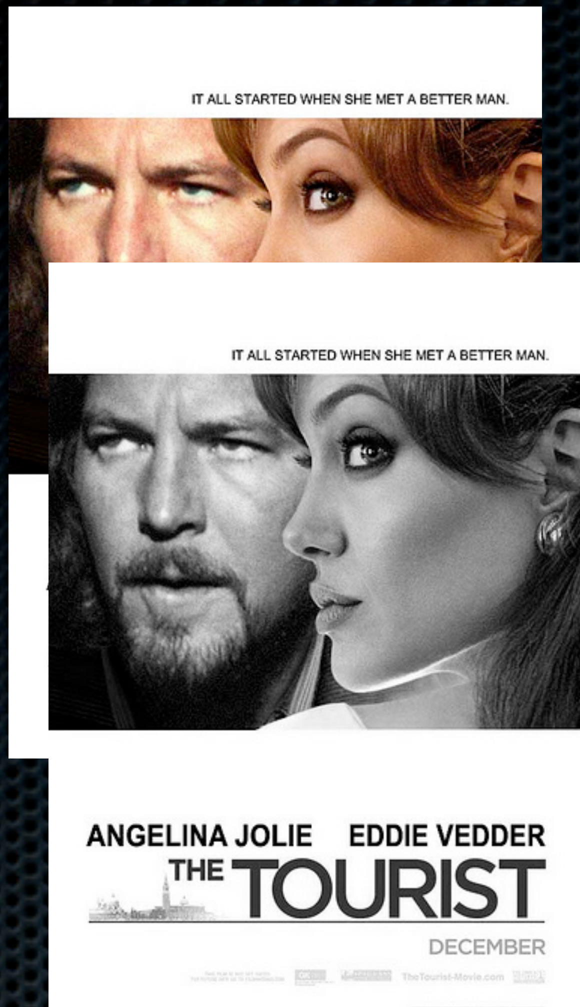
IT ALL STARTED WHEN SHE MET A BETTER MAN.
IT ALL STARTED WHEN SHE MET A BETTER MAN.
ANGELINA JOLIE EDDIE VEDDER
THE TOURIST
DECEMBER

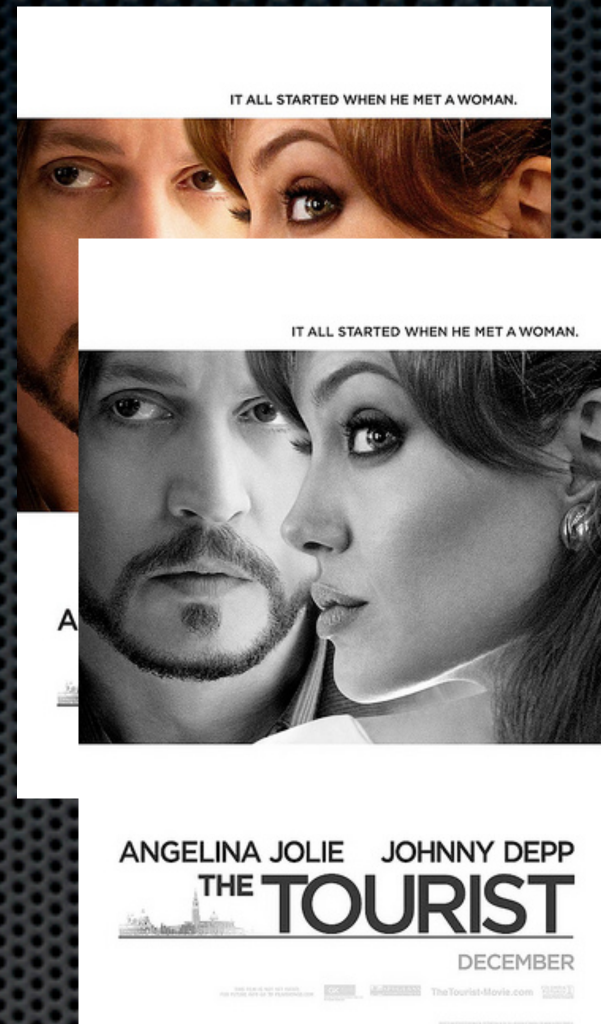
IT ALL STARTED WHEN HE MET A WOMAN.
IT ALL STARTED WHEN HE MET A WOMAN.
ANGELINA JOLIE JOHNNY DEPP
THE TOURIST
DECEMBER


IT ALL STARTED WHEN HE MET A WOMAN.
IT ALL STARTED WHEN HE MET A WOMAN.
ANGELINA JOLIE JOHNNY DEPP
THE TOURIST
DECEMBER

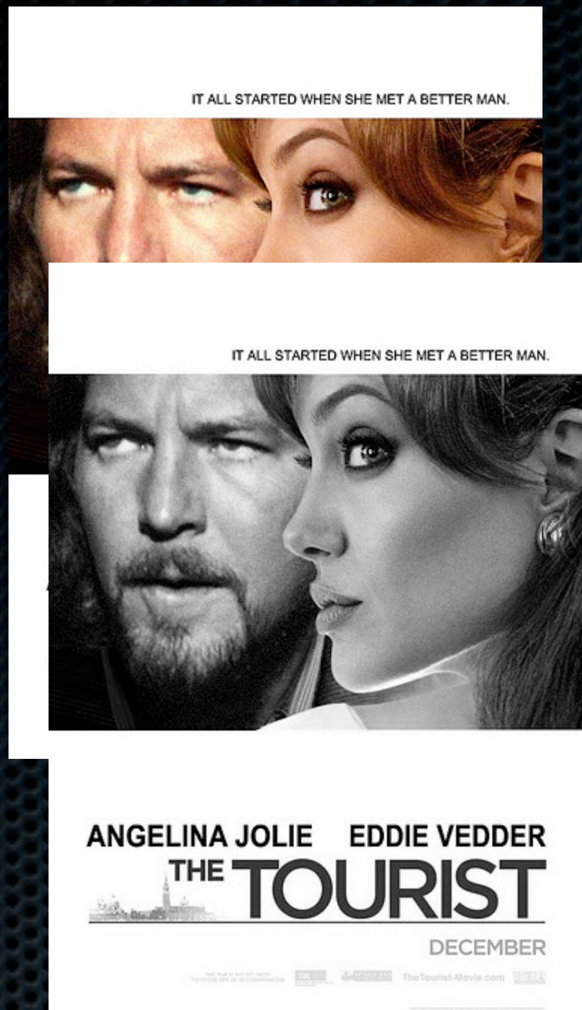
IT ALL STARTED WHEN SHE MET A BETTER MAN.
IT ALL STARTED WHEN SHE MET A BETTER MAN.
ANGELINA JOLIE
EDDIE VEDDER
THE TOURIST
DECEMBER

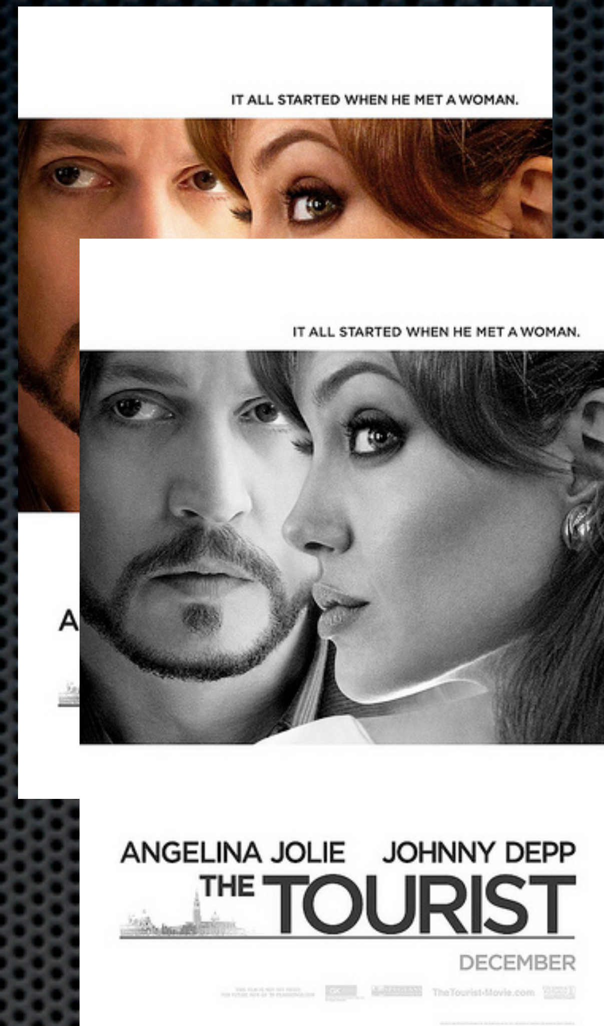
IT ALL STARTED WHEN HE MET A WOMAN.
IT ALL STARTED WHEN HE MET A WOMAN.
ANGELINA JOLIE
JOHNNY DEPP
THE TOURIST
DECEMBER

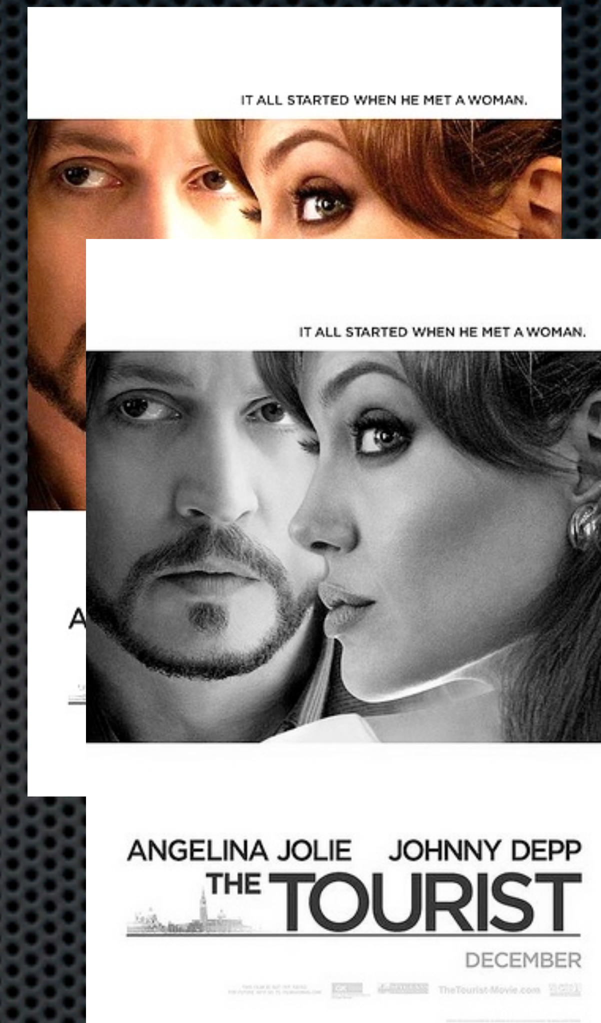
IT ALL STARTED WHEN HE MET A WOMAN.
IT ALL STARTED WHEN HE MET A WOMAN.
ANGELINA JOLIE
JOHNNY DEPP
THE TOURIST
DECEMBER

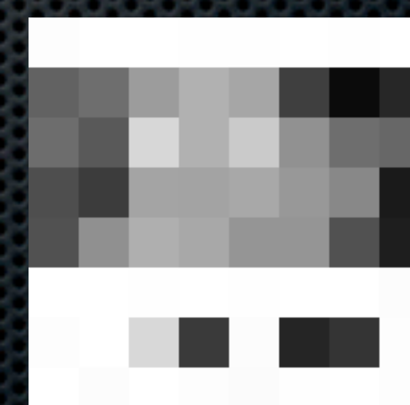

Multi-dimensional Discrete Cosine Transform

$$X_{k_1,k_2} = \sum_{n_1=0}^{N_1-1} \left(\sum_{n_2=0}^{N_2-1} x_{n_1,n_2} \cos\left[\frac{\pi}{N_2}\left(n_2+\frac{1}{2}\right)k_2\right]\right) \cos\left[\frac{\pi}{N_1}\left(n_1+\frac{1}{2}\right)k_1\right]$$

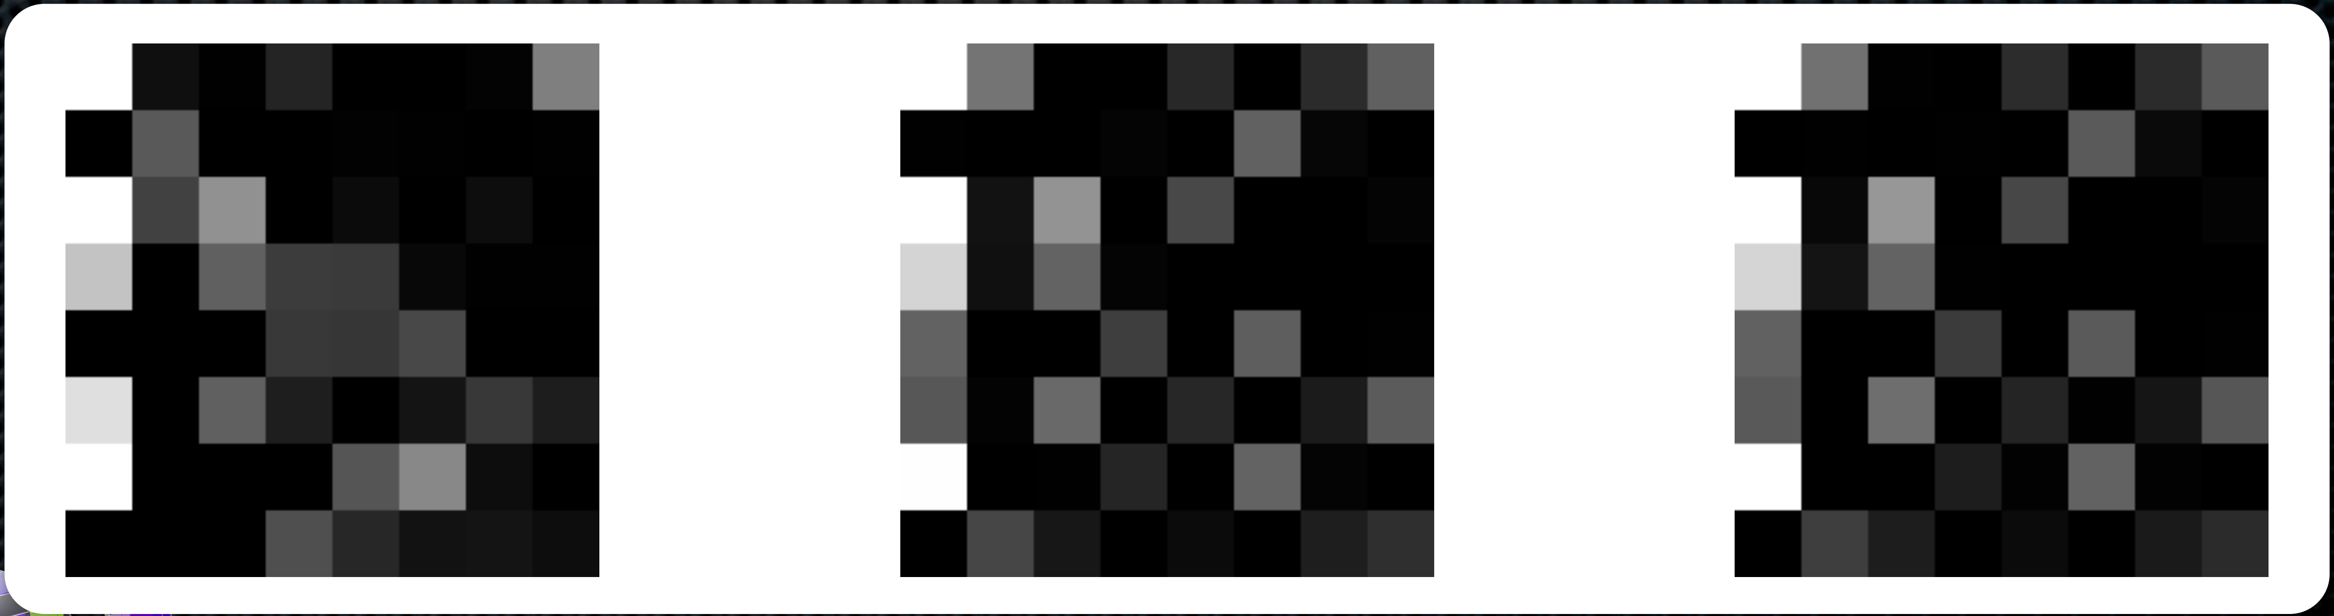

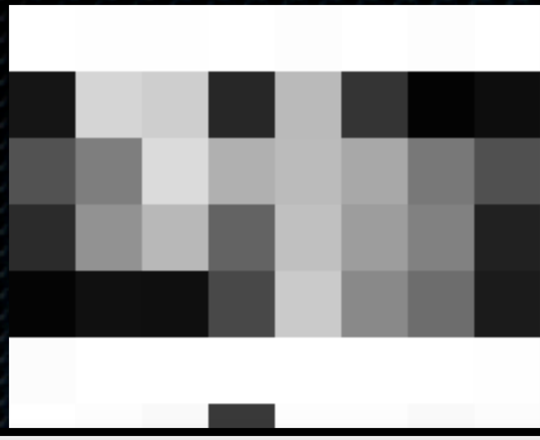
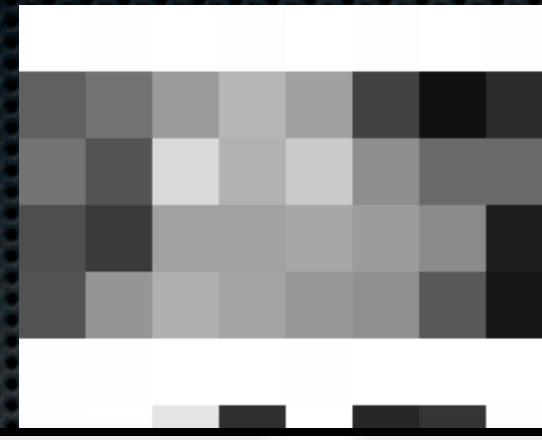


$$\begin{bmatrix}
6.1917 & -0.3411 & 1.2418 & 0.1492 & 0.1583 & 0.2742 & -0.0724 & 0.0561 \\
0.2205 & 0.0214 & 0.4503 & 0.3947 & -0.7846 & -0.4391 & 0.1001 & -0.2554 \\
1.0423 & 0.2214 & -1.0017 & -0.2720 & 0.0789 & -0.1952 & 0.2801 & 0.4713 \\
-0.2340 & -0.0392 & -0.2617 & -0.2866 & 0.6351 & 0.3501 & -0.1433 & 0.3550 \\
0.2750 & 0.0226 & 0.1229 & 0.2183 & -0.2583 & -0.0742 & -0.2042 & -0.5906 \\
0.0653 & 0.0428 & -0.4721 & -0.2905 & 0.4745 & 0.2875 & -0.0284 & -0.1311 \\
0.3169 & 0.0541 & -0.1033 & -0.0225 & -0.0056 & 0.1017 & -0.1650 & -0.1500 \\
-0.2970 & -0.0627 & 0.1960 & 0.0644 & -0.1136 & -0.1031 & 0.1887 & 0.1444
\end{bmatrix}$$

BDF2AB0DF0AAA433



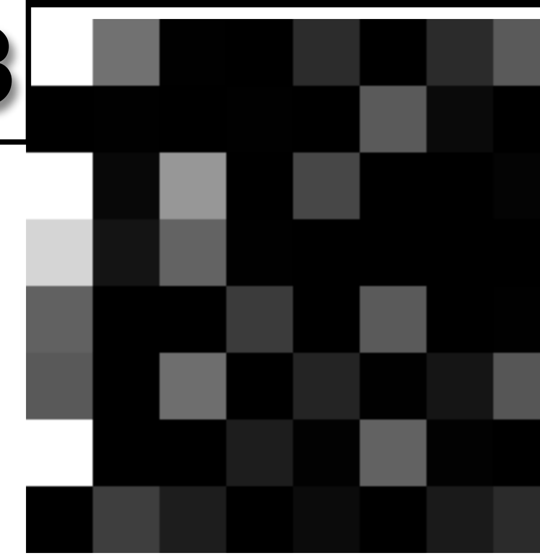

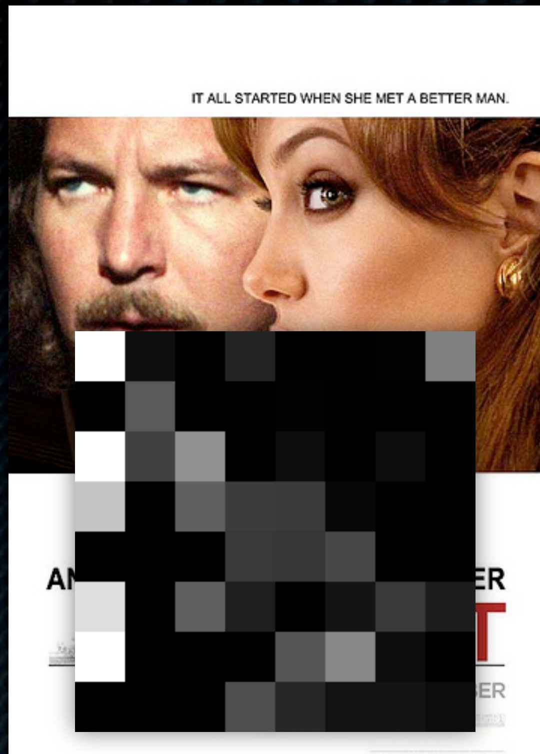

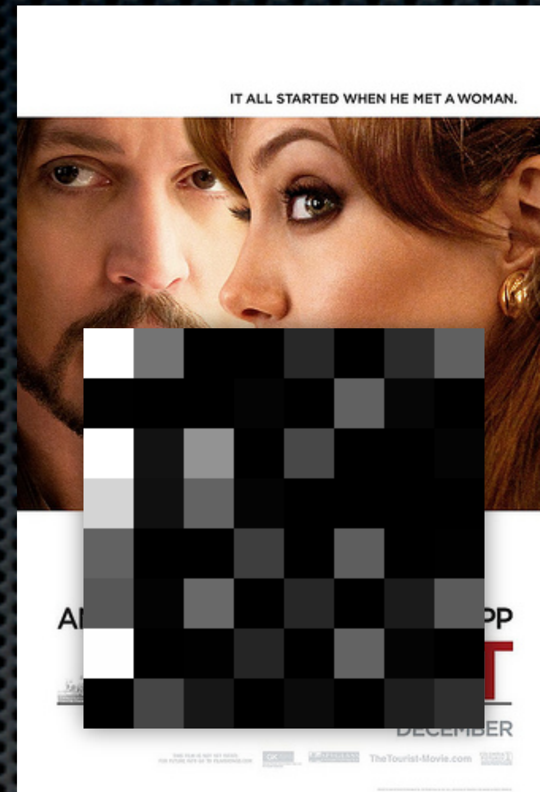

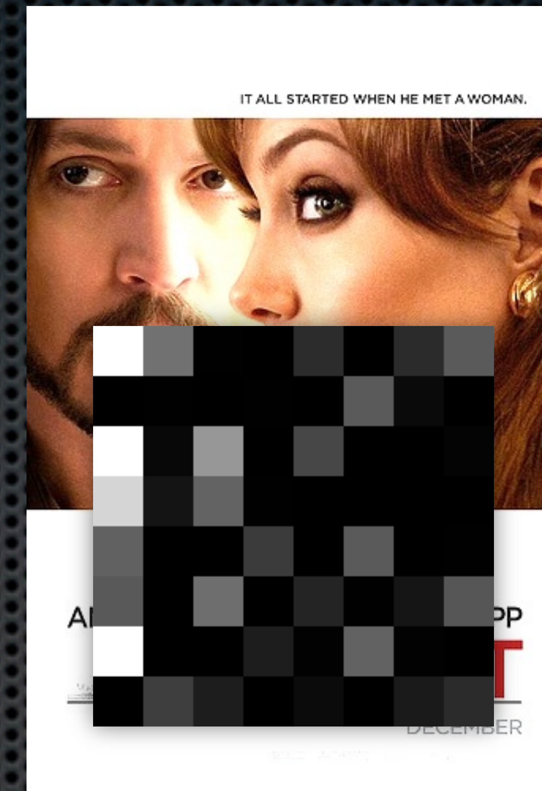

d341eabc1cb78e1f cb16e9f094ab966b cb16e9f094ab966b

hd = 23

hd = 2

hd = 23

Alternative methods

- Fourier-Mellin transform: Finding near-duplicate images on the web using fingerprints, S.H. Srinivasan et al.
- Perceptual hashing: www.phash.org/docs/pubs/thesis_zauner.pdf
- Beyond “Near-Duplicates”: Learning Hash Codes for Efficient ...: research.google.com/pubs/archive/36579.pdf
- ...

Faceted Search

Joint work with:

Lluis Garcia, Mridul Muralidharan, Borkur Sigurbjornsson and many others!

WWW'10

YAHOO!

Overview

- Serving facets for image search
- Extracting entity facets
- Extracting ranking features
- Ranking candidate entity facets
- Evaluation
- Conclusions

People

- Product & design
  - Kaushal Kurapati
  - Anuj Sahai
  - Polly Ng
- Engineering
  - Anand Ramani
  - Sriram ‘Thiru’ Sathish
  - Ramu Adapala
  - Abhinav Katiyar
  - Murali Krishna
  - Balaji Kanan
- Research
  - Roelof van Zwol
  - Borkur Sigurbjornsson
  - Lluis Garcia
  - Mridul Muralidharan
- Sciences
  - Nicolas Torzec

Facets in Image Search

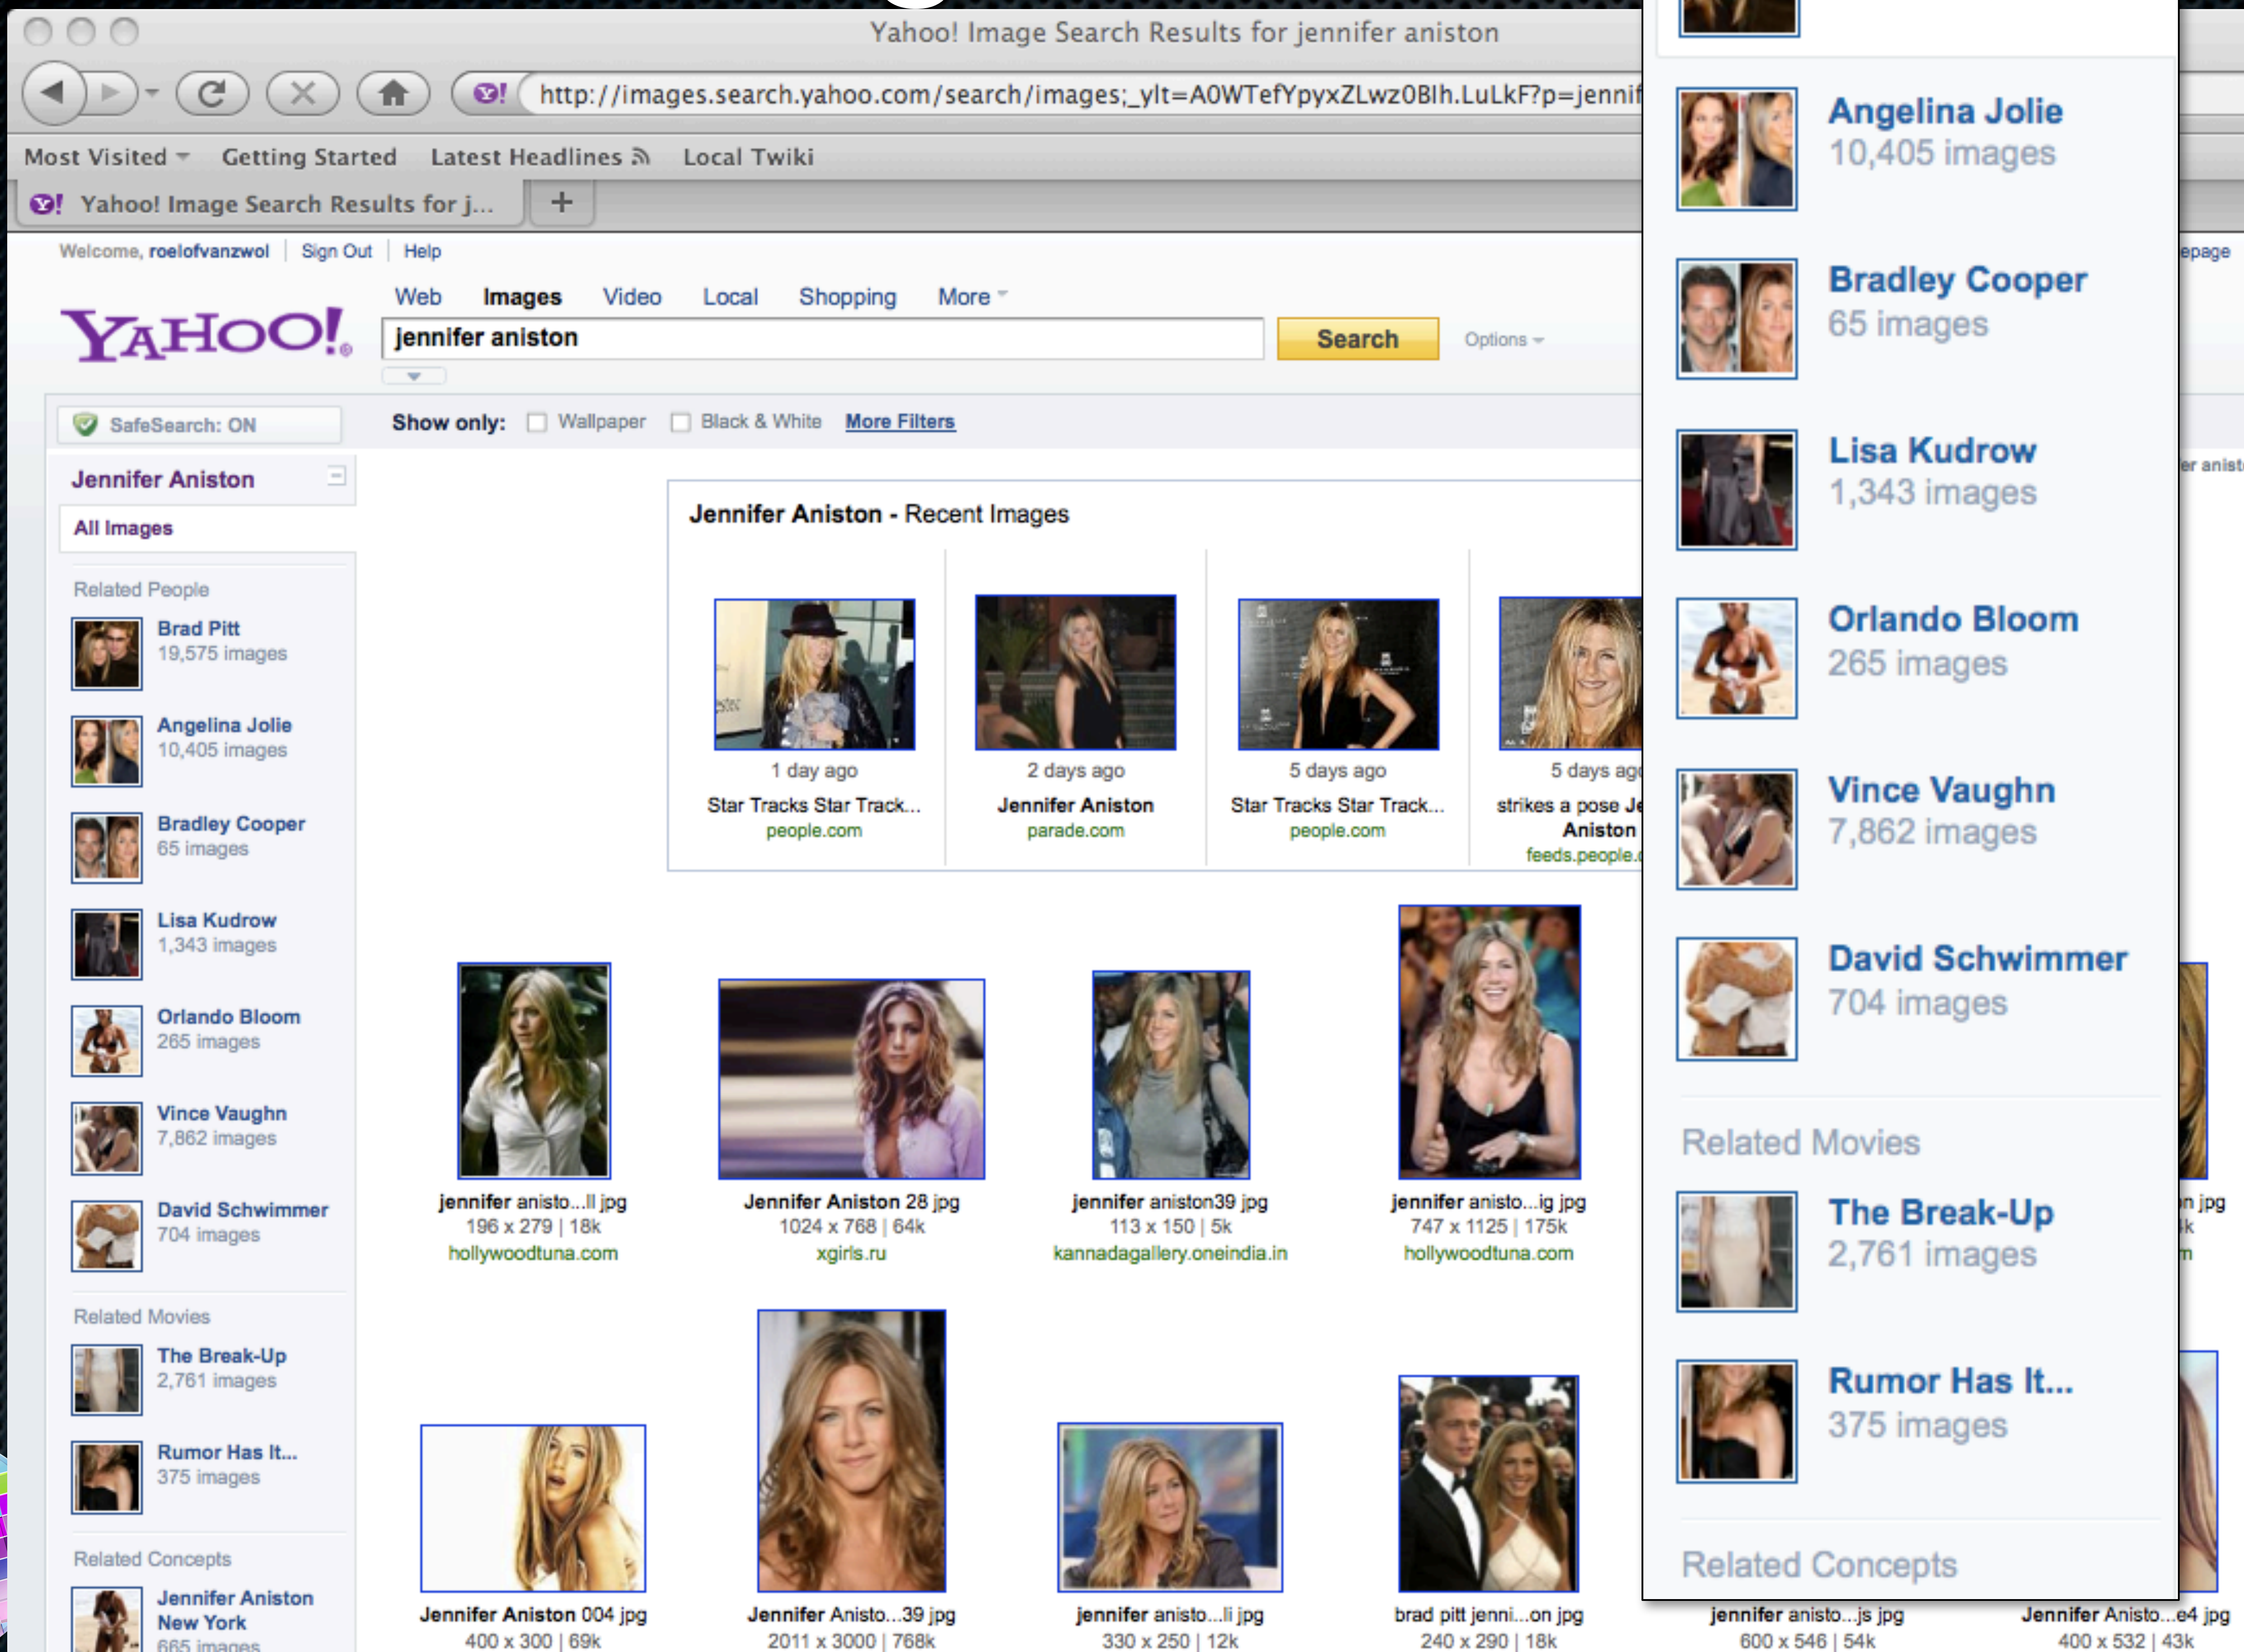

Facets in Image Search

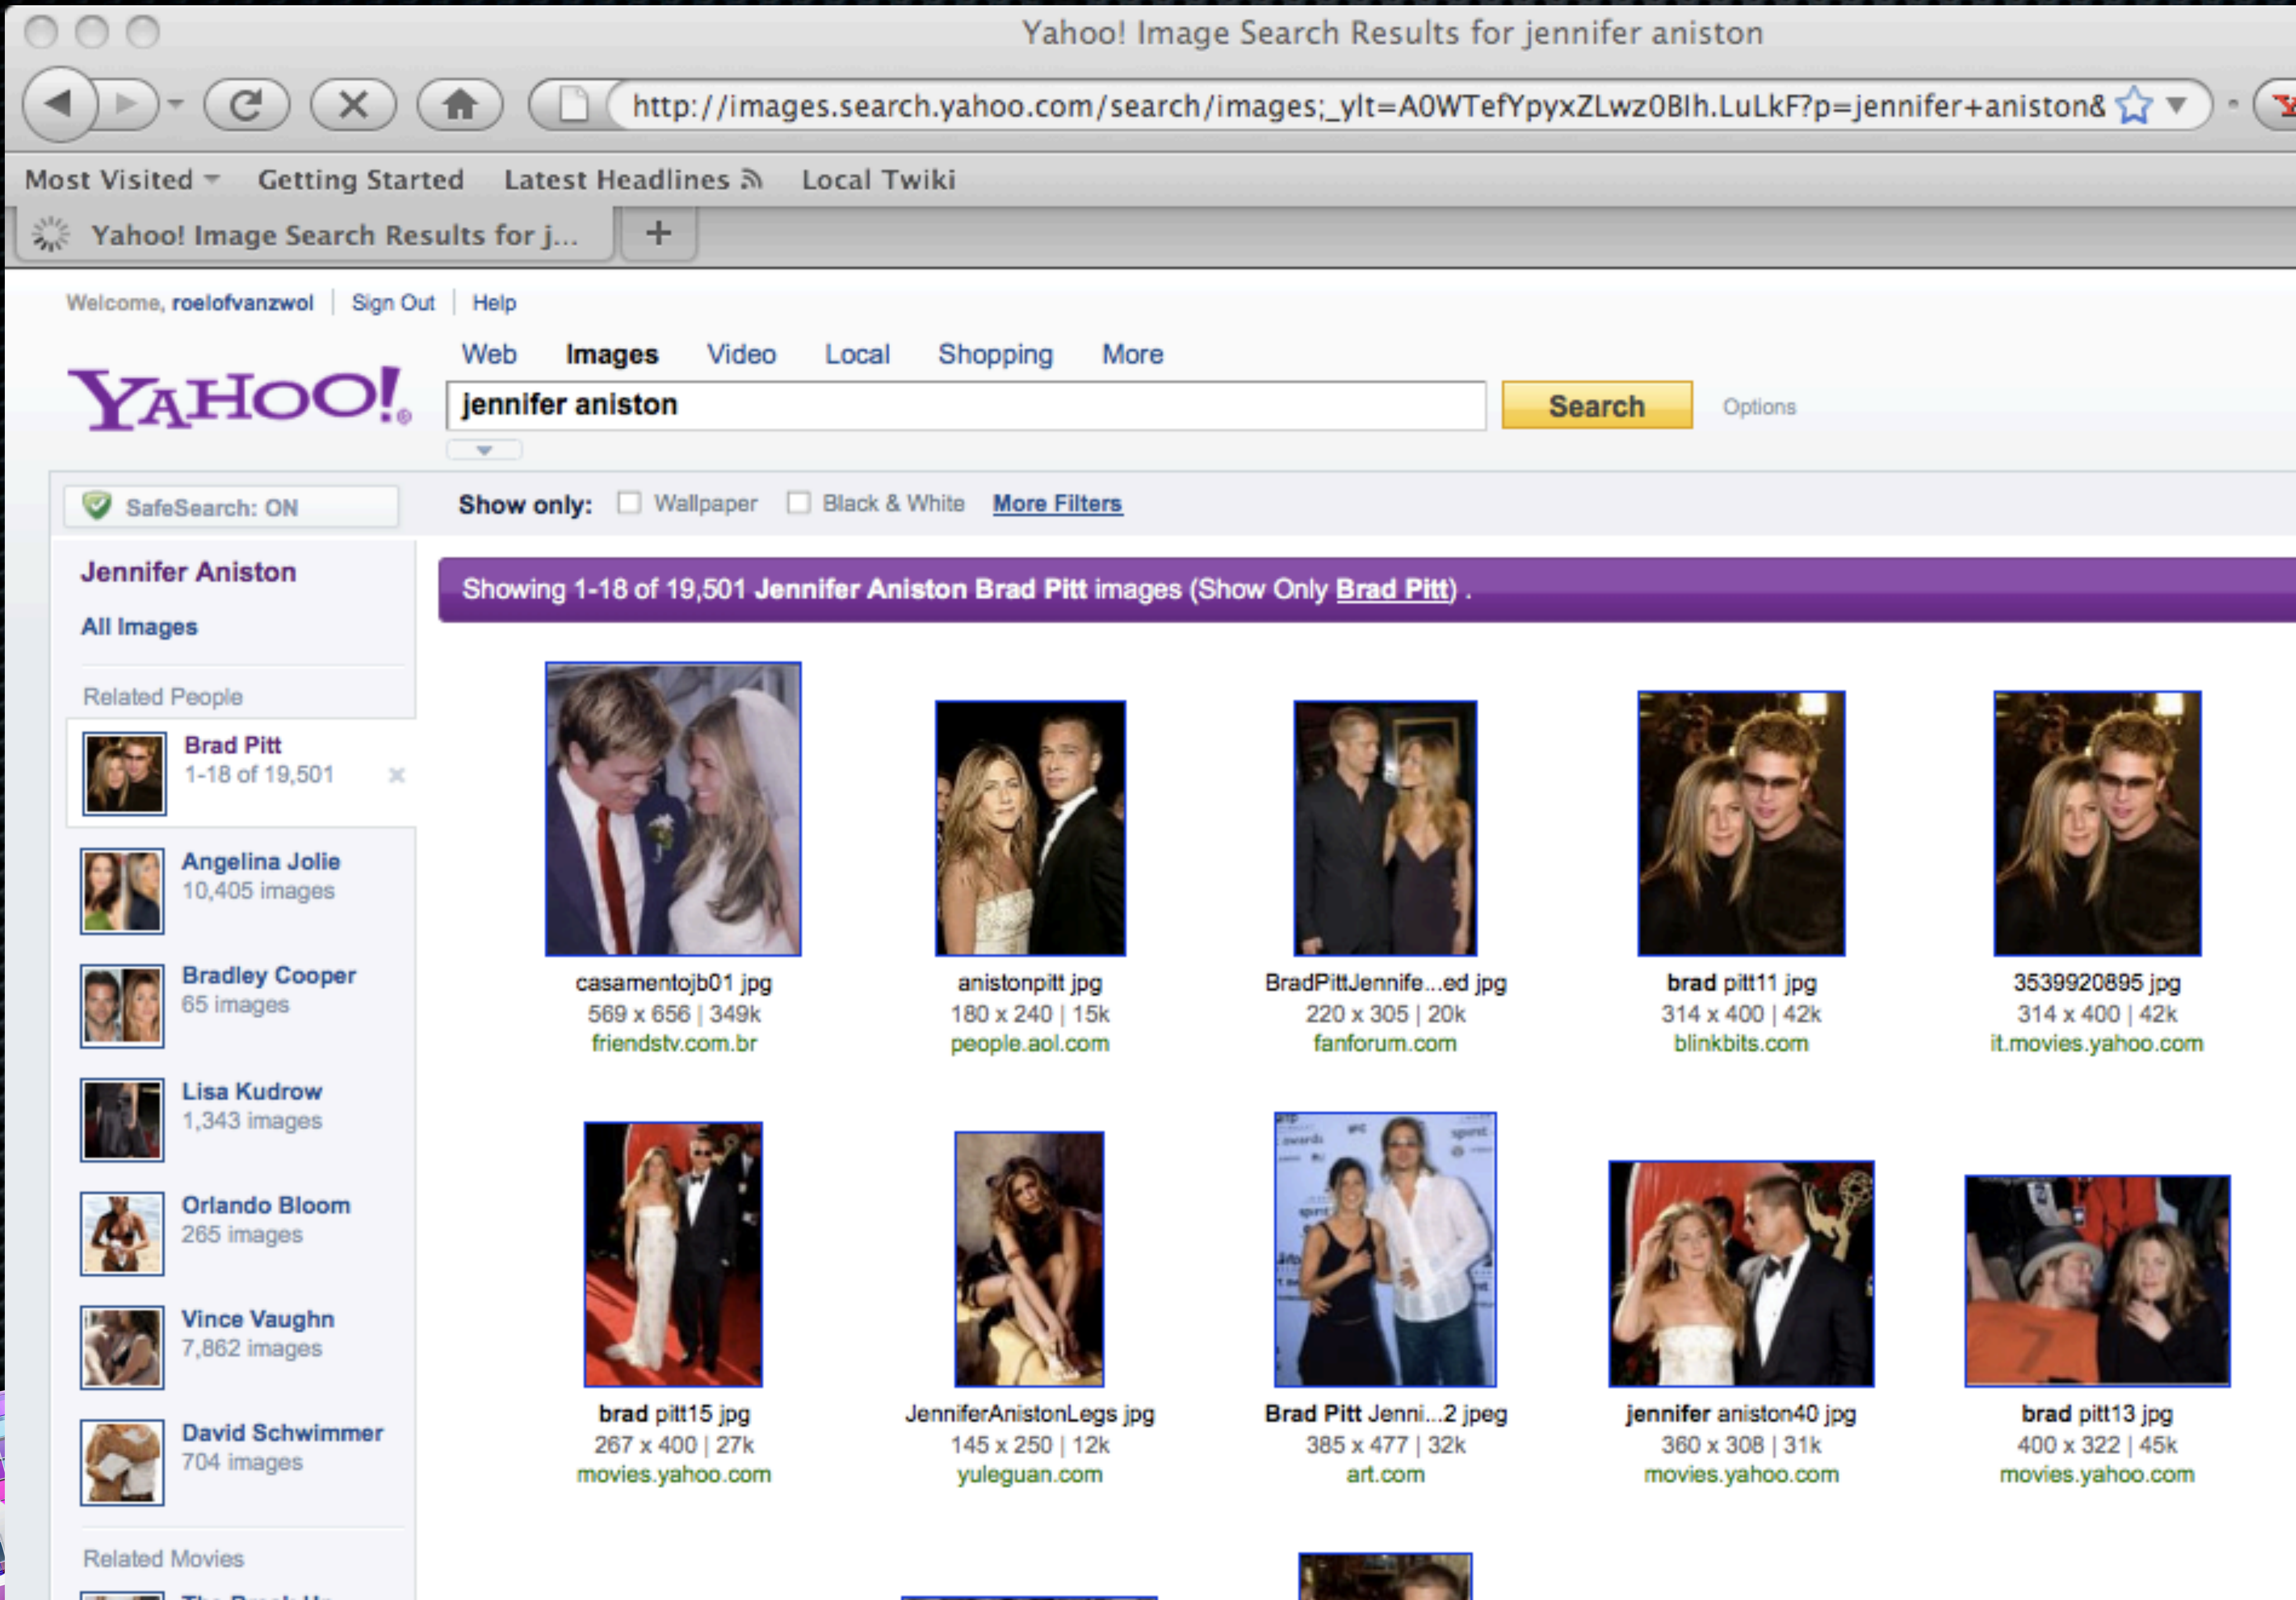

GridFaces

- **Goals**
  - Power the faceted search experience of image search
  - Promote the "Web-of-Objects" paradigm through the introduction of facets in the SERP.
- **Main milestones**
  - May 15th: Travel facets in bucket test
  - July 23th: Travel facets launched
  - August 6th: Celebrity facets in bucket test
  - September 21st: Celebrity facets launched
  - December 2010: video search, mobile adopts facets
- **Now: smart thumbnails, explore/exploit, provide context**

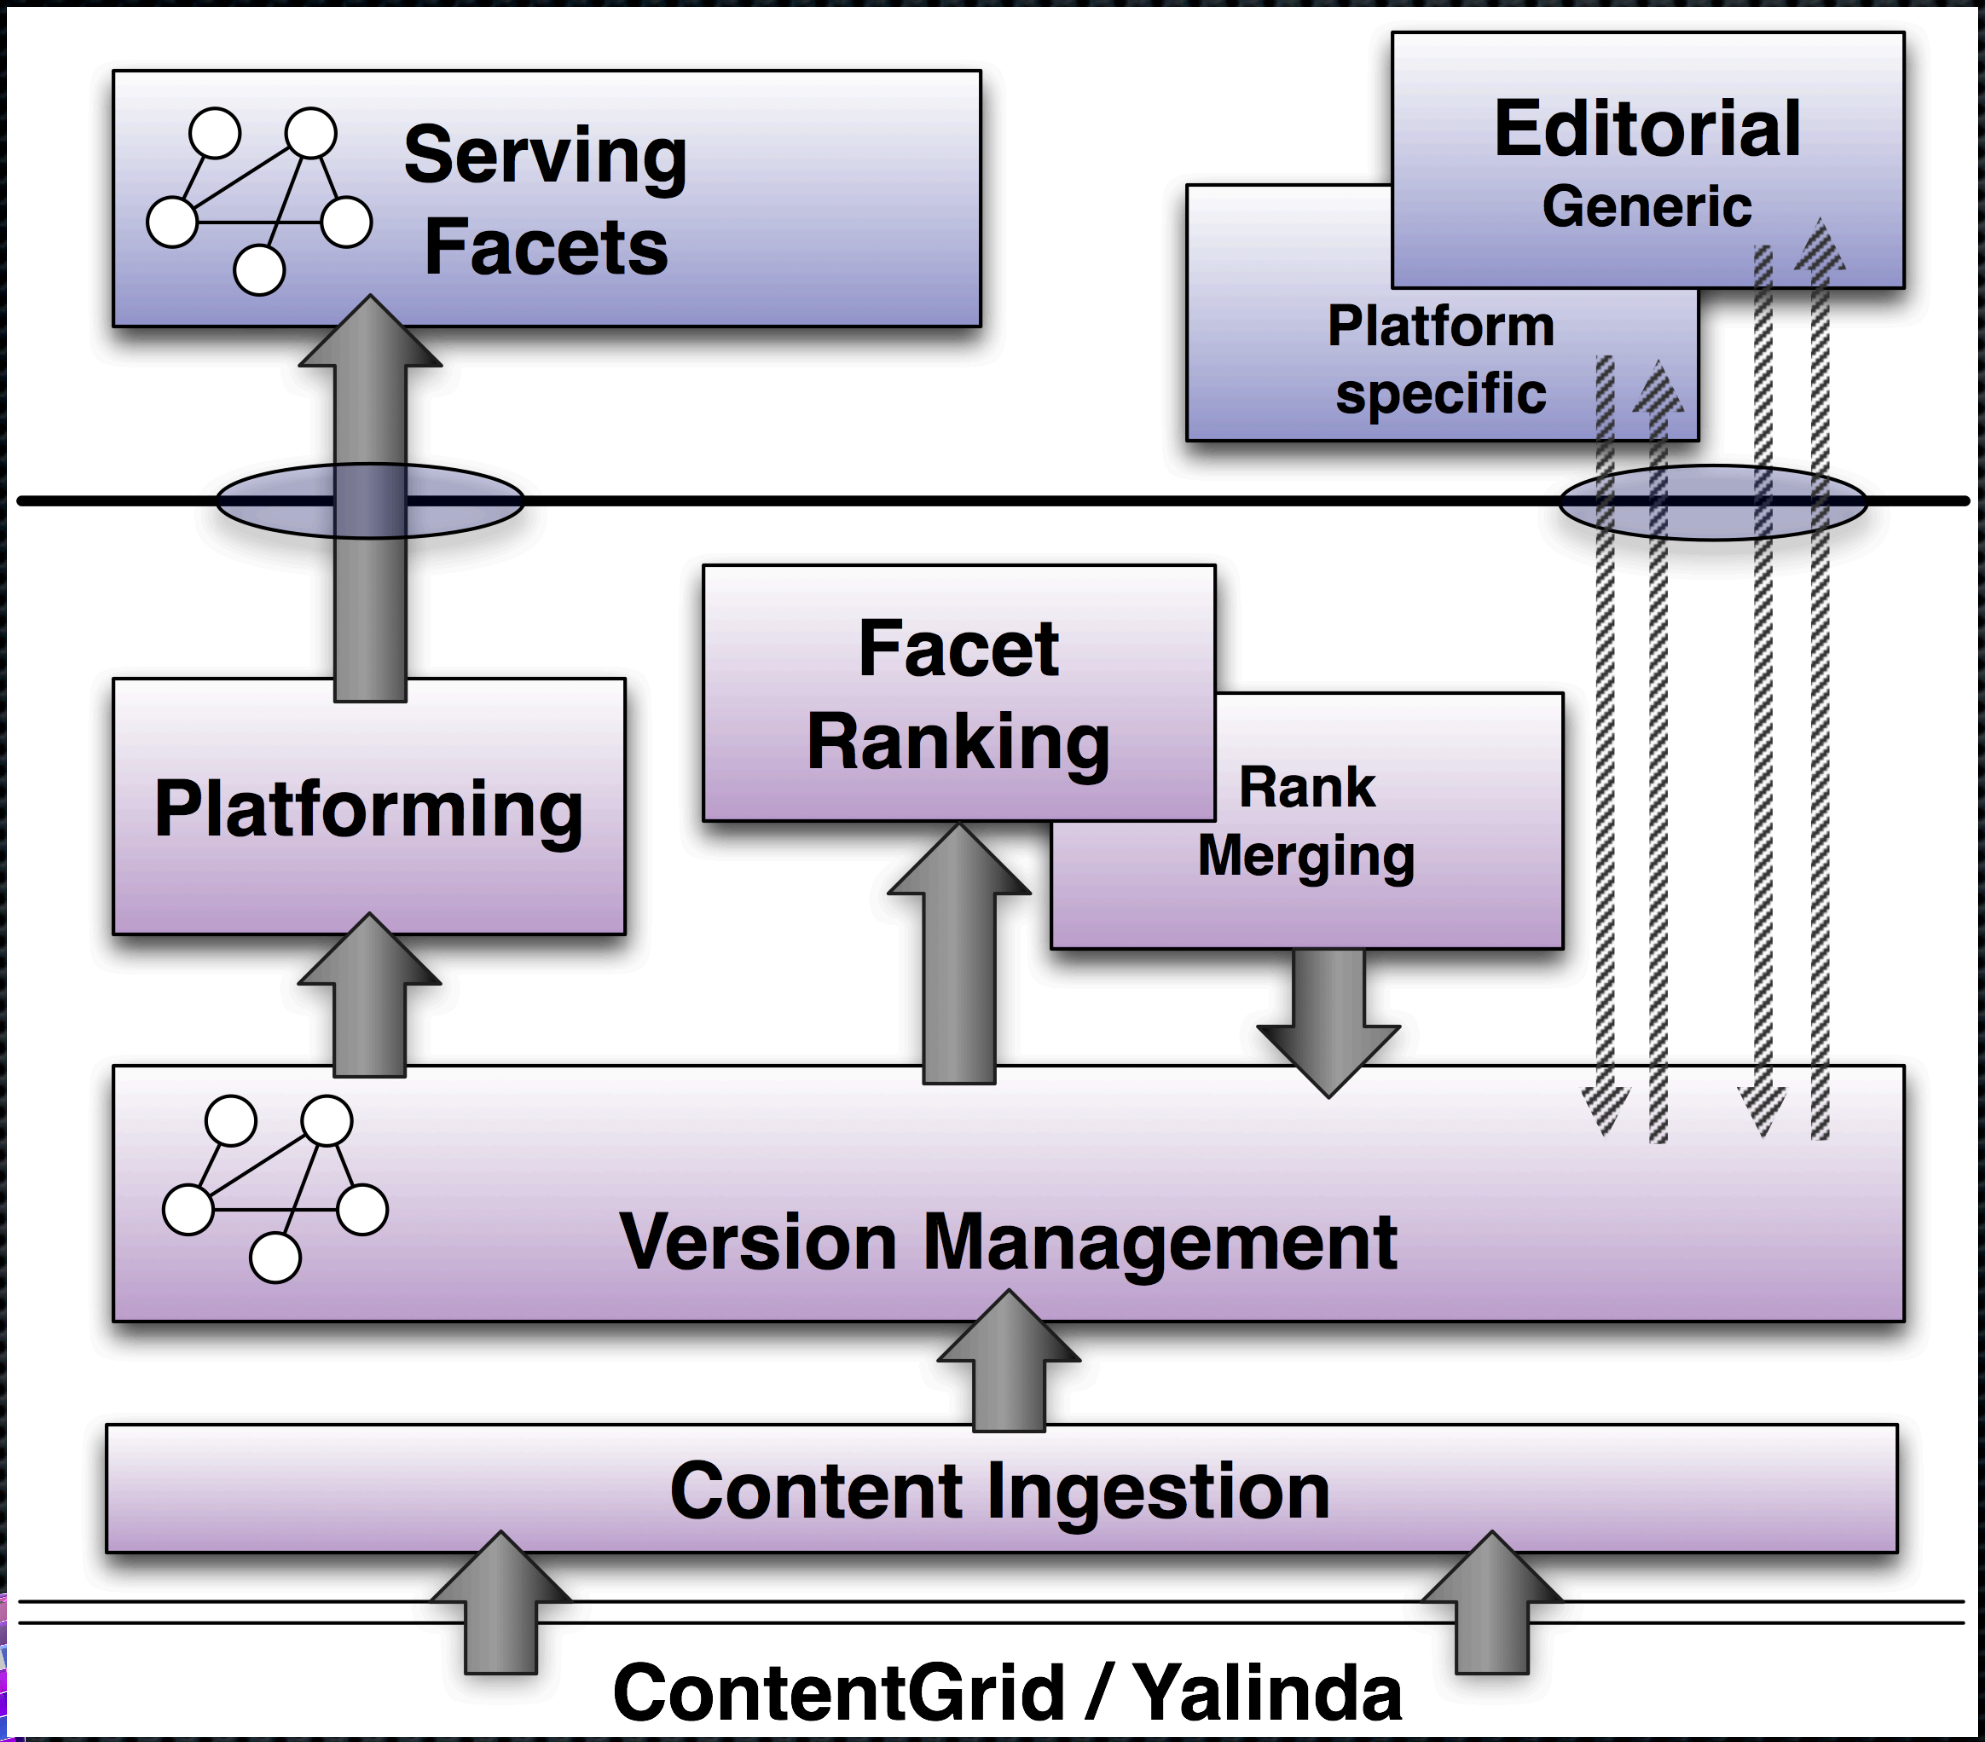

Serving Facets
Editorial
Generic
Platform specific
Facet Ranking
Platforming
Rank Merging
Version Management
Content Ingestion
ContentGrid / Yalinda

Extracting Facets

- A facet is defined as the directed relationship between two entities (e,f).
- For a given entity e, a set of candidate facets F is collected. We refer to entity $f \in F$ as the facet of entity e.
- Entities and facets are extracted from structured sources, such as: Y! Movies, GeoPlanet, Y! Travel, Wikipedia, etc.
- Pool of 6M+ entities and 80M+ candidate facets

name	**Justin Timberlake**
aliases	JT, Justin Randall Timberlake, J. Timberlake
type	Person (musician, actor)

entity *e*	Justin Timberlake
entity *f*	Jessica Biel
type	Romantic relationship

Extracting Features

- Extract a set of features from different ranking sources:
  - Image search query terms
  - Image search user sessions
  - Annotated photos in Flickr
  - Favorites in Y!music
- Pre-process sources into a common format
- Extract statistical features:
  - Atomic, Symmetric, A-symmetric, and Combined features

Extracting Features

- Query term analysis:
  - For every query entered by a user, we extract co-occurring entity pairs:

User query:	**Cubbon park in Bangalore, India**
Tokenization:	Cubbon+park+in+Bangalore+India
Normalization:	cubbon+park+in+bangalore+india
Segmentation:	cubbon+park+in+bangalore+india
Entity detection:	cubbon park; bangalore; india; bangalore india
Cooc pairs:	(cubbon park, bangalore india), (cubbon park, india), (cubbon park, bangalore), (bangalore,india)

- Per event collect (common format):
  - eventId, userId, timestamp, (e1,e2)+

Extracting Features

- Independent from the source, the following set of features is extracted:

Atomic features
$P(e)$, $P(f)$
$E(e)$, $E(f)$
Symmetric features
$P(e,f)$, $P_u(e,f)$, $SI(e,f)$, $CS(e,f)$
A-symmetric features
$P(e\|f)$, $P(f\|e)$, $P_u(f\|e)$, $KL(e\|\|f)$, …
Combined features
$P_u(e\|f)xP(f)$, $P_u(f\|e)xP(f)$ …

- * Pu(f|e) is a variant of P(f|e), where each entity e and entity pair (e,f) is counted once per user. To make the feature less prune to the impact of a single user.

User Click Feedback

Adopted two click-feedback models:

$$ctr_{e,f} = \frac{clicks_{e,f}}{views_{e,f}}$$

$$coec_{e,f} = \frac{clicks_{e,f}}{\sum_{p=1}^{P} views_{e,f_p} \times ctr_p}$$

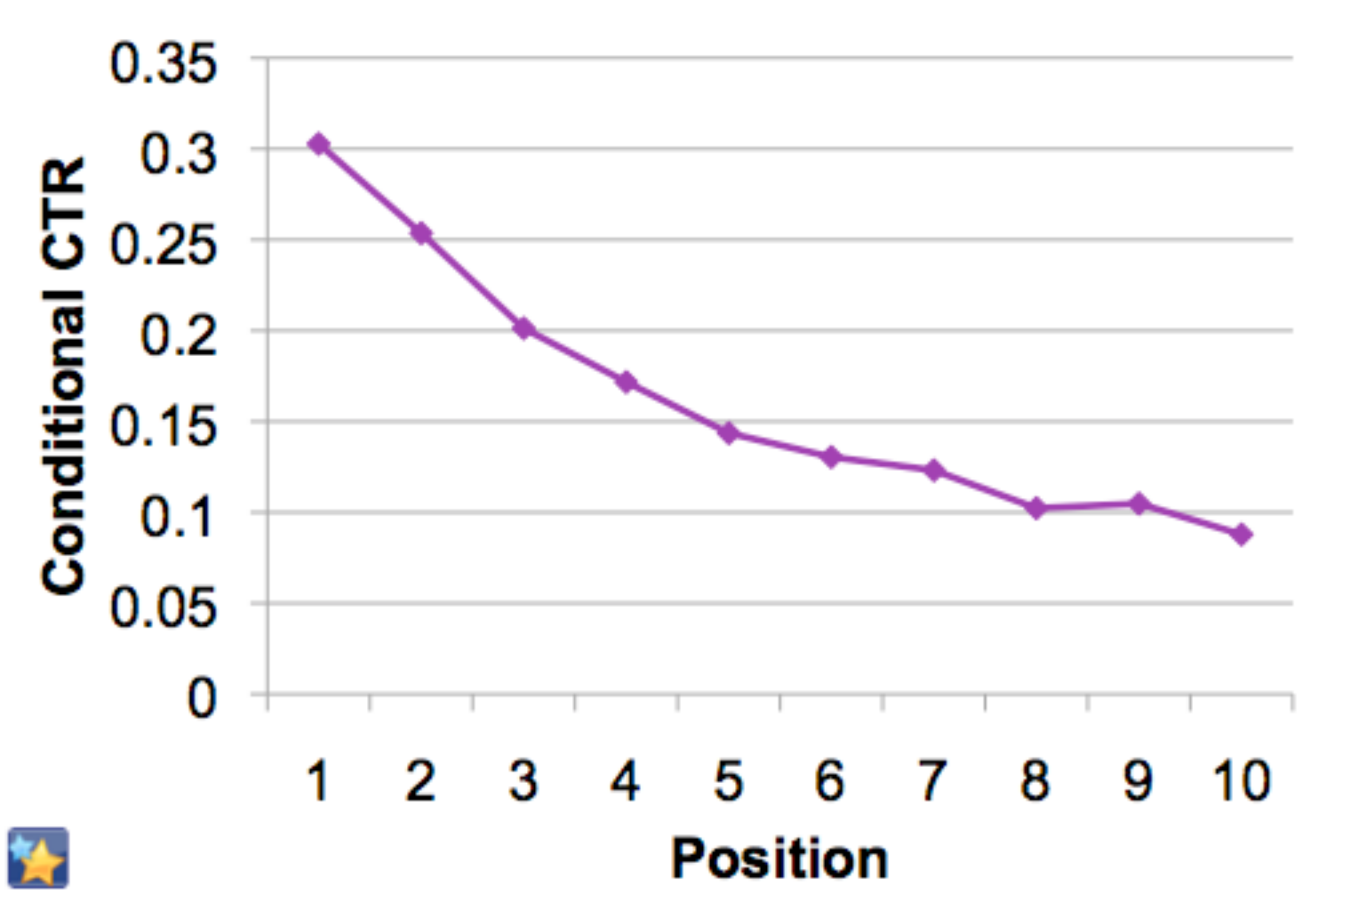

YAHOO!

Evaluation – Overall Performance

	CTR		COEC	
Run	**mDCG**	**mnDCG**	**mDCG**	**mnDCG**
Ideal	2.375	--	2.594	--
Baseline	1.728	0.709	1.812	0.677
GBDT*	2.090	**0.874**	2.436	**0.930**

- Based on the mnDCG computed over the first 10 results.
- Comparing baseline performance against the CTR/ COEC is unfair, due to the grouping of facets by category!

Evaluation – DCG@p

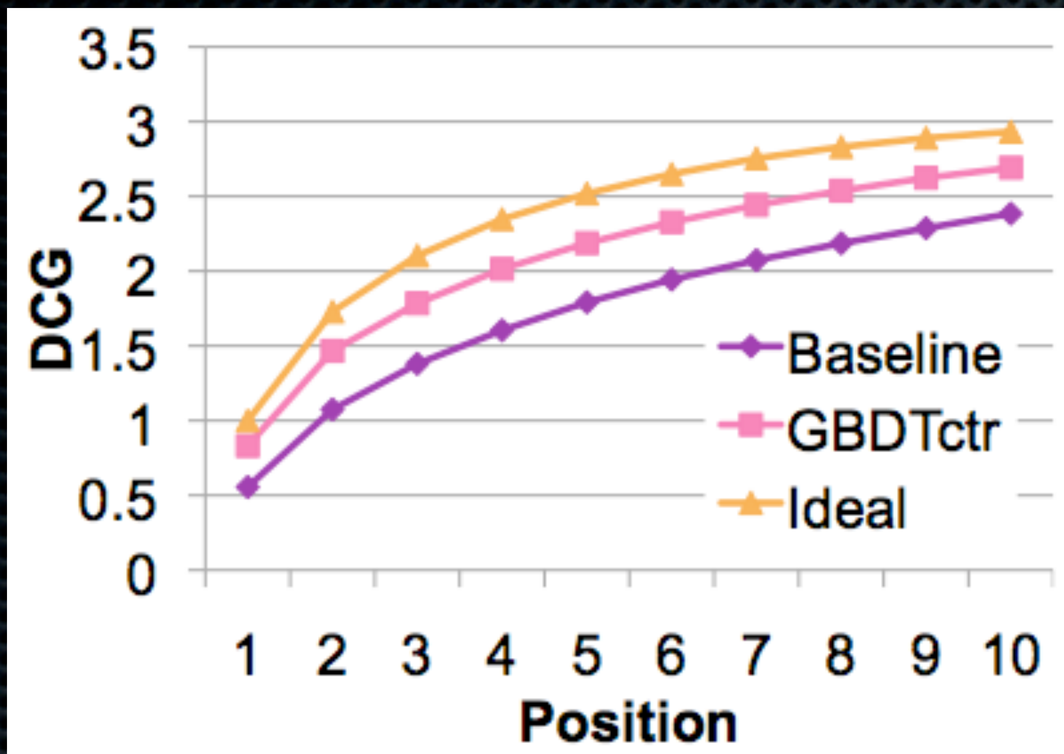

CTR test set

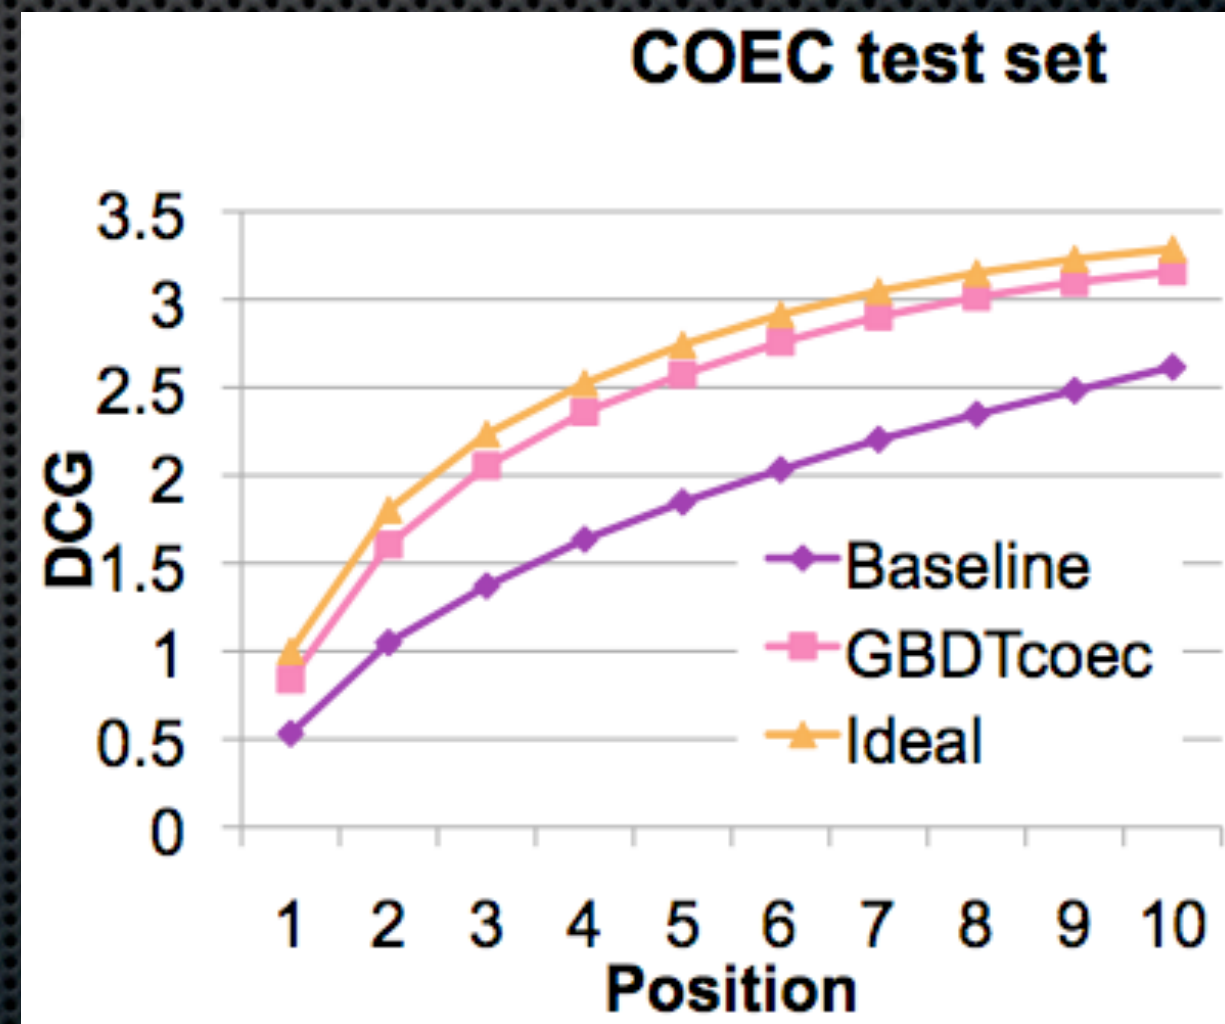

Evaluation – nDCG@p

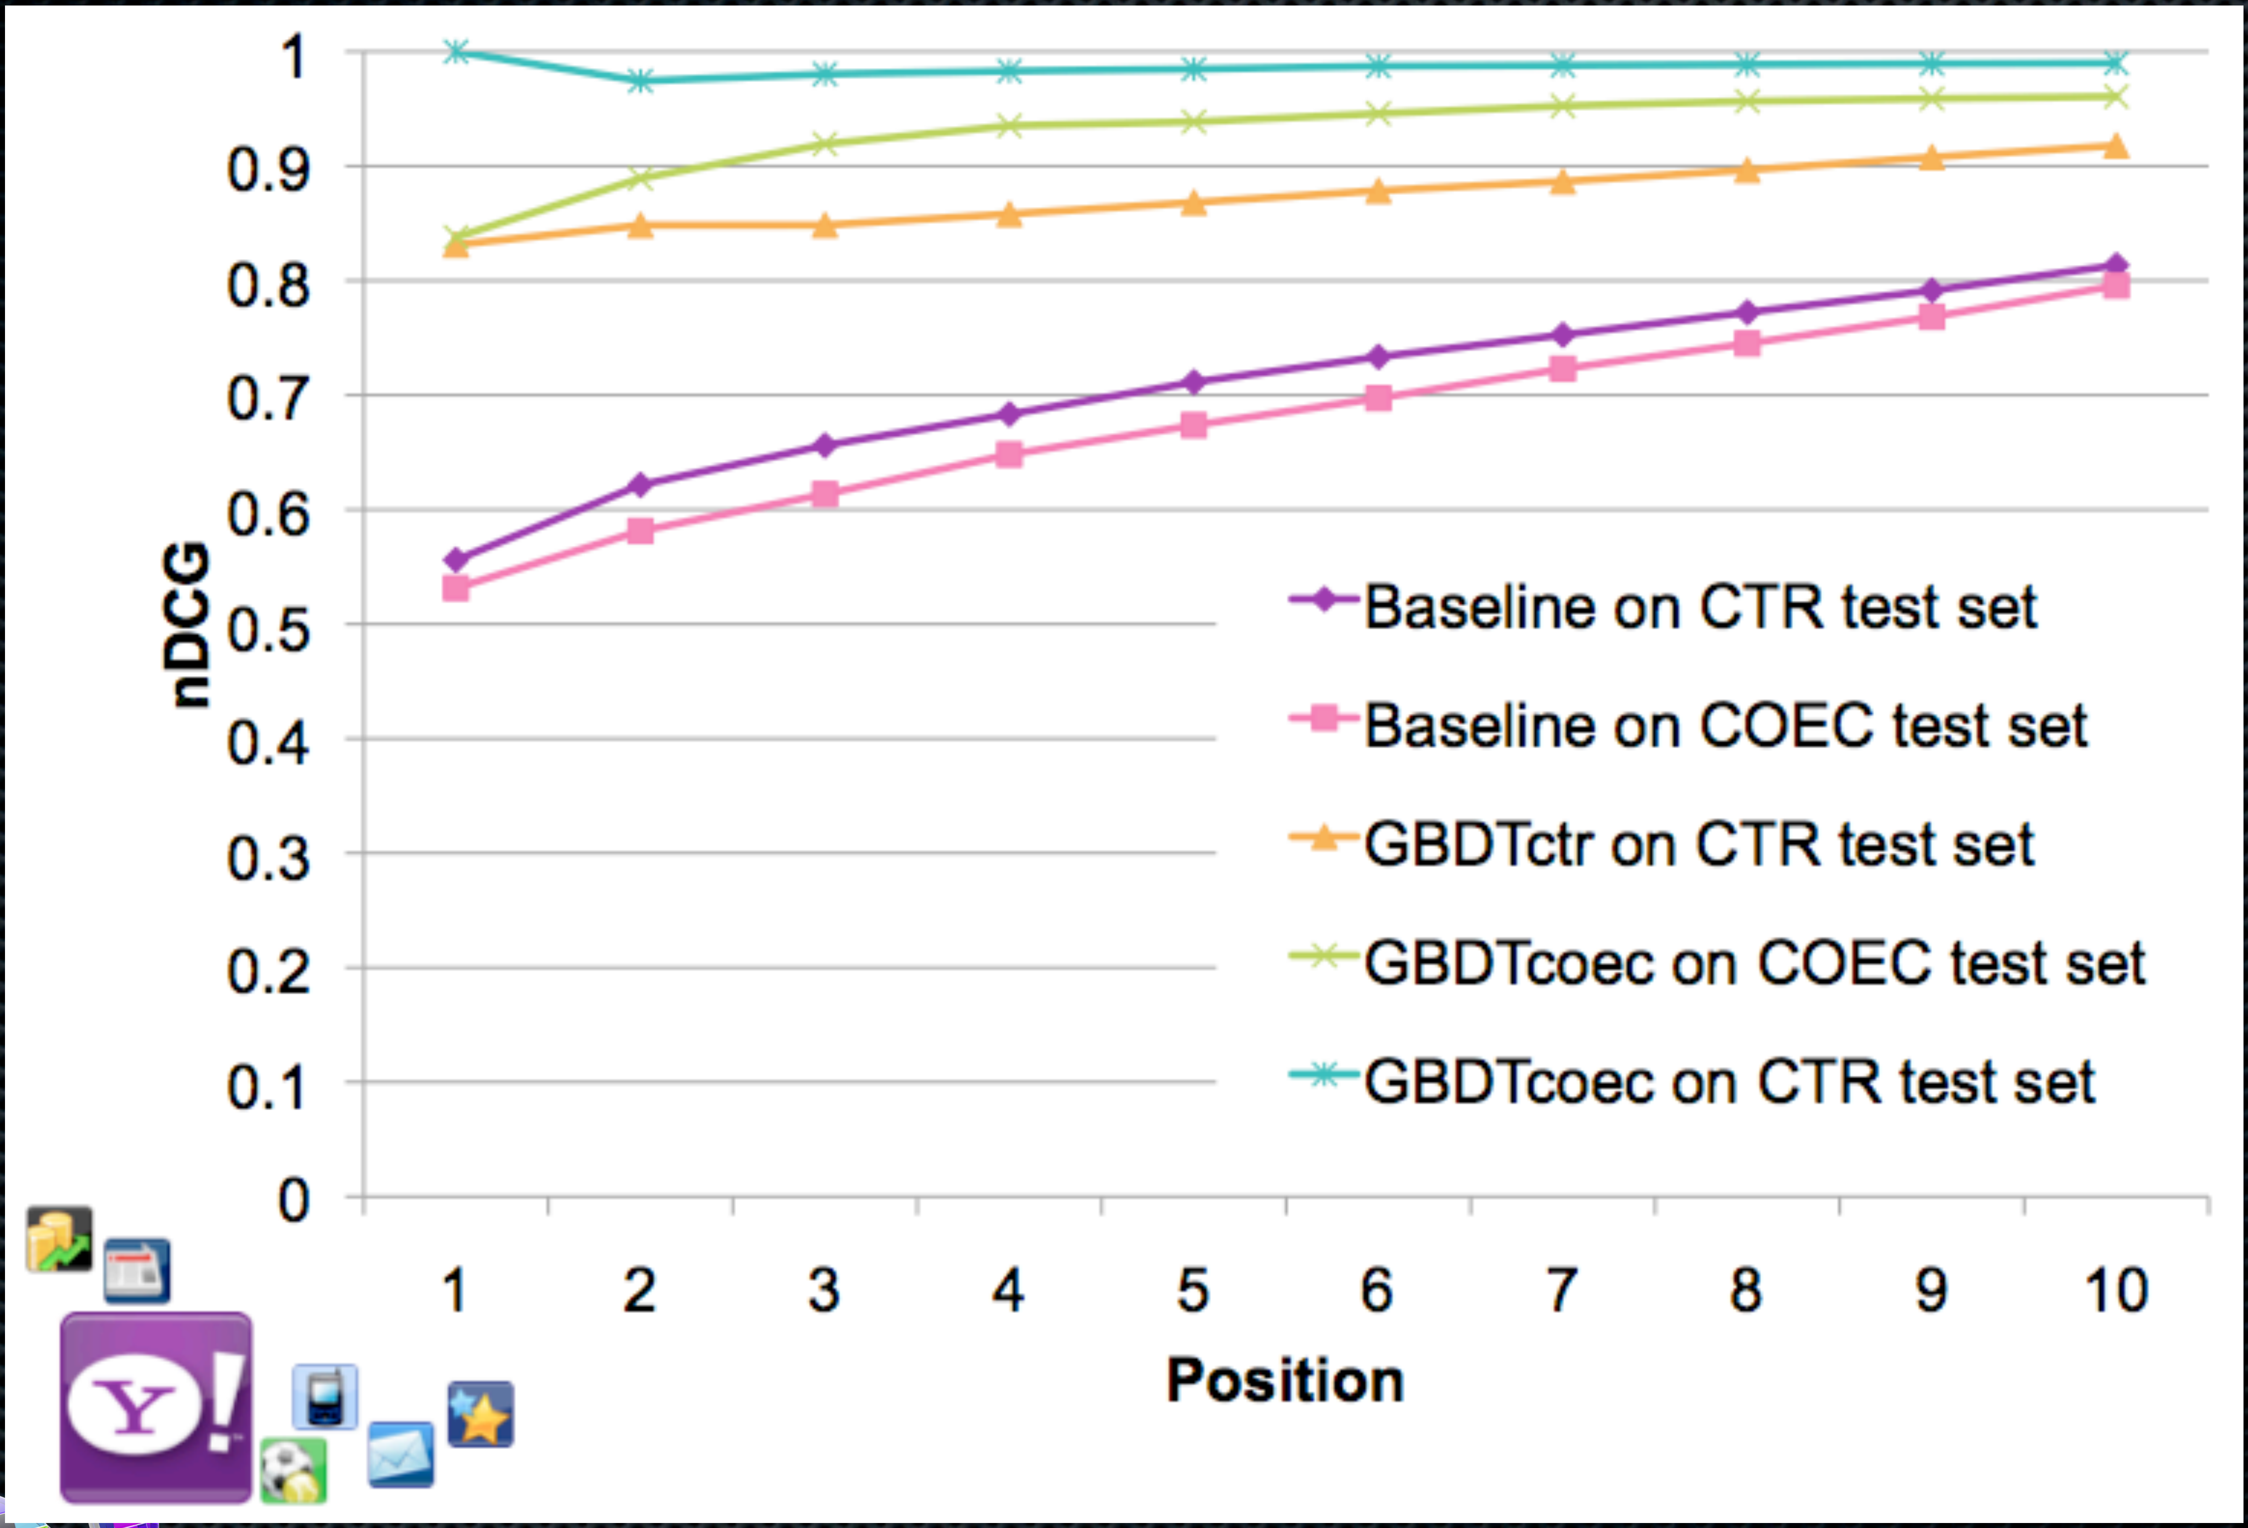

Evaluation – per query

- All results reported are statistically significant ($p<0.001$)
- “Justin Timberlake” example:

Facet	CTR	COEC	Basel.	G_{coec}	G_{ctr}
Jessica Biel	1	7	5	1	1
Jesse Mc Cartney	2	1	12	2	4
Britney Spears	3	5	10	5	3
NSync	4	2	7	3	9
Alpha Dog	5	3	1	6	10
Cameron Dias	6	10	4	4	6
JC Chasez	7	9	2	8	8
T.I.	8	4	13	9	11
Ciara	9	8	9	11	5
Timbaland	10	6	11	10	13
Positional error	**--**	**28**	**41**	**10**	**23**

Evaluation – Feature Importance

GBDT$_{ctr}$		GBDT$_{coec}$	
Feature	**Weight**	**Feature**	**Weight**
QS $P_u(e\|f) * P(f)$	100	QT $P_u(e, f)/P(f)$	100
QT $P(e)$	85.11	FT $P(e)$	11.56
QS $P(e)$	76.88	QT $P(e)$	9.57
QT $P_u(e, f)$	69.32	QS $P(e)$	9.22
QT $P_u(e\|f) * P(f)$	69.21	FT $E(e)$	9.22
QT $P_u(f\|e) * P(f)$	64.38	FT $KL(e)$	8.84
QS $P(e, f)$	59.78	QT $P_u(f\|e) * P(f)$	8.19
QT $P(e, f)$	52.98	QT $P_u(e\|f) * P(f)$	8.14
QS $P_u(e, f)$	48.26	QS $P(f)$	7.53
FT $P(e)$	43.71	QT $P_u(e, f)$	7.25

QT: Query term; QS: Query session; FT: Flickr tag.

Conclusions – Ongoing work

- **Yahoo!** first to introduce the WOO in the SERP
- Proposed a machine learned approach for ranking facets, based on user-click feedback
- Extract features in generic manner from various sources (query term-, query session-, and Flickr tag analysis)
- Enriched feature space (user prone, and combined)
- Adopt/evaluated two click-feedback models
- **2011 is about:**
  - **engaging media, Explore/exploit, Providing context!**

Smarter Thumbnails

Joint work with:
Lyndon Kennedy, Nicolas Torzec, Belle Tseng

YAHOO!

Motivation

- Eye-catchers
- Summary
- “Triggers”

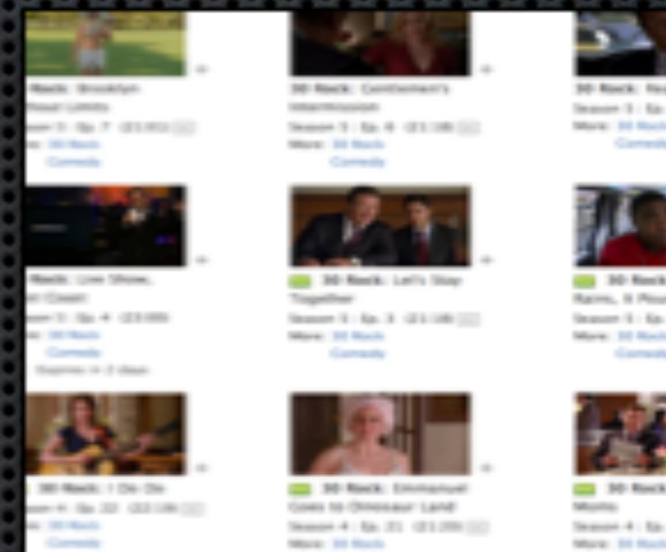

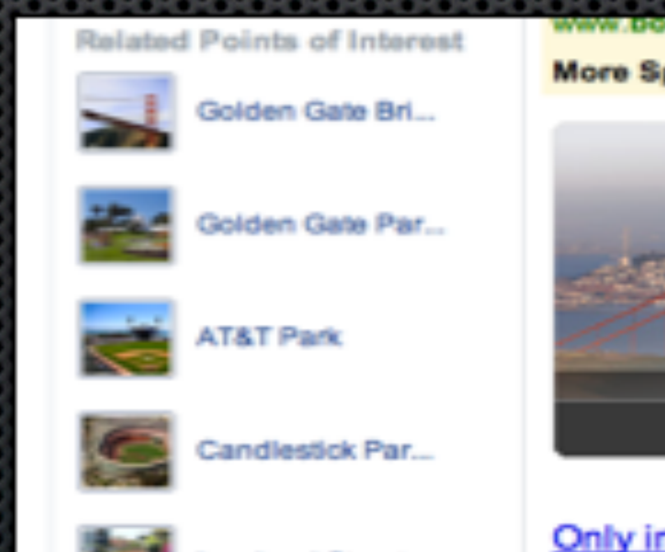

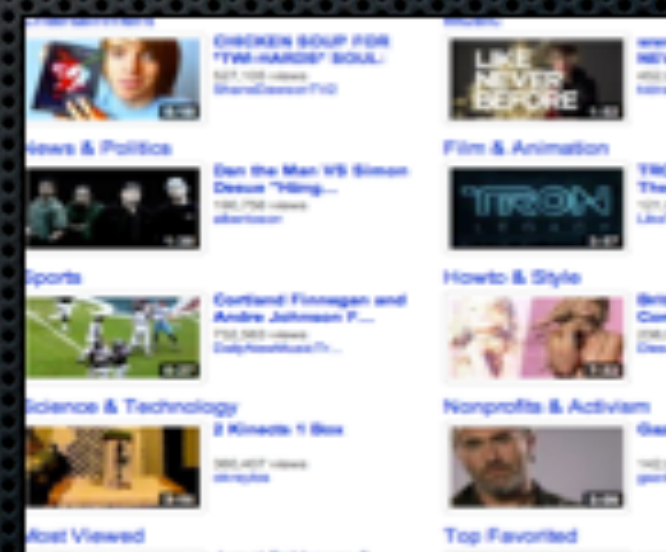


Web
Images
Video
Local
Shopping
News
More
George Clooney
Search
QuickApps
Search Pad
SafeSearch - On
5,260,000 results for
George Clooney
Related People
Brad Pitt
Matt Damon
Renee Zellweger
Frances McDorma...
Meryl Streep
Coen brothers
Latest On: Jennifer Anis... Nicolas Cage Evangeline Li... Abigail Bresl...
Overview
George Clooney Takes On Wall Street And ...
Irish Central - Apr 15 07:23am
George Clooney is set to produce, and possibly direct, a movie about the 2008 financial crisis and the government bailout of troubled financial institutions. Full Story »
More George Clooney Stories
George Clooney Photos
More Photos »
George Clooney - Wikipedia, the free encyclopedia
Early life | Career | Humanitarian work | In the media
George Timothy Clooney is an American actor, film director, producer, and screenwriter. For his work as an actor, he has received two Golden Globe Awards and an Academy Award. Clooney...
en.wikipedia.org/wiki/George_Clooney - Cached
George Clooney: Biography from Answers.com
Born: 6 May 1961 Birthplace: Lexington, Kentucky Best Known As: The international star of

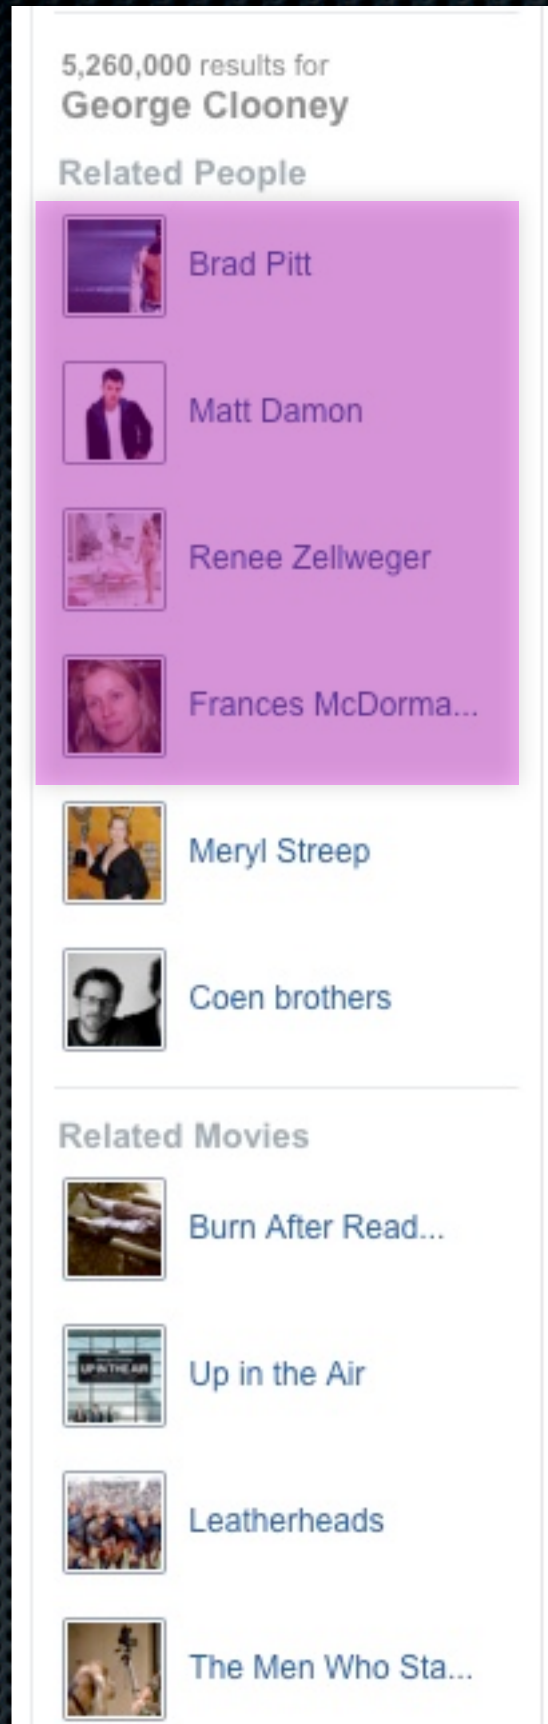
5,260,000 results for
George Clooney
Related People
Brad Pitt
Matt Damon
Renee Zellweger
Frances McDorma...
Meryl Streep
Coen brothers
Related Movies
Burn After Read...
Up in the Air
Leatherheads
The Men Who Sta...

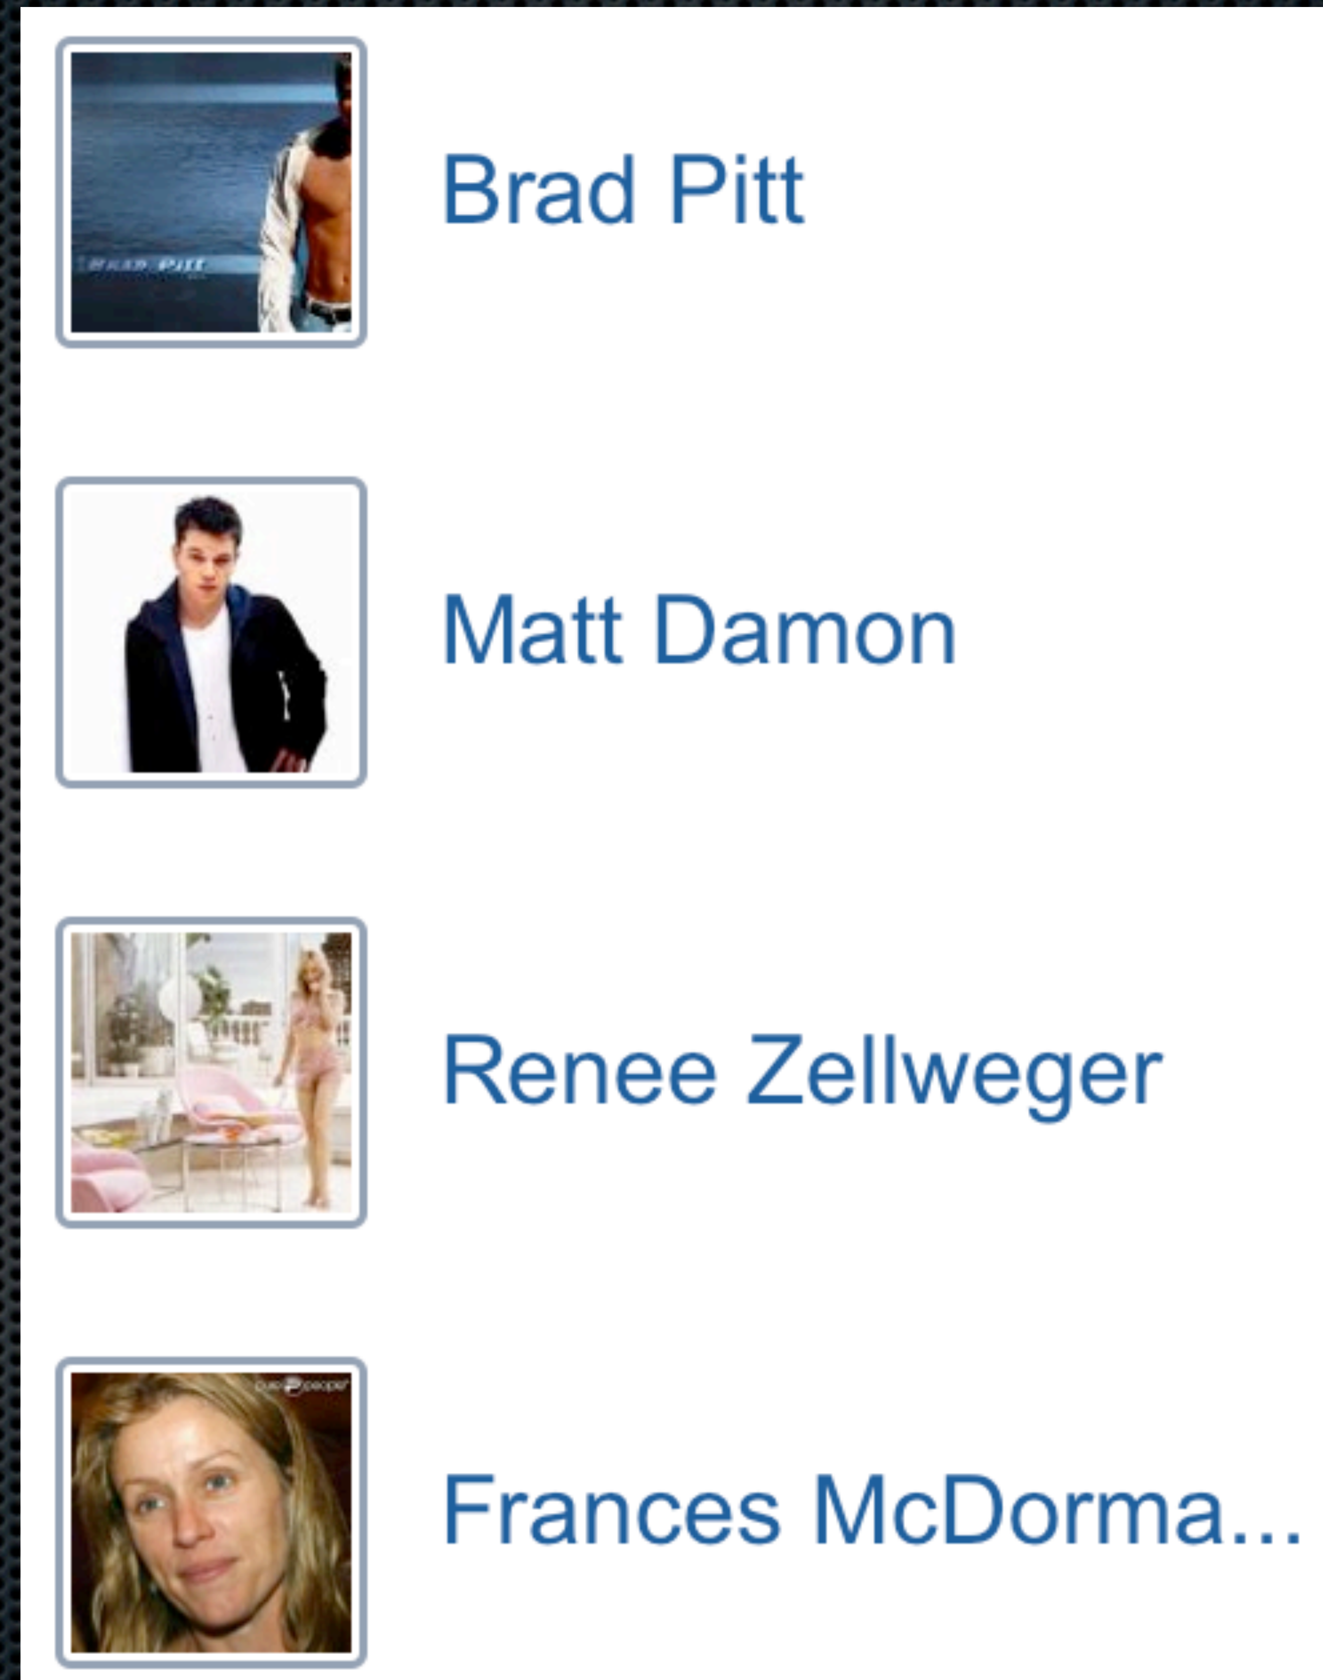
Brad Pitt
Matt Damon
Renee Zellweger
Frances McDorma...

Machine learned thumbnail generation

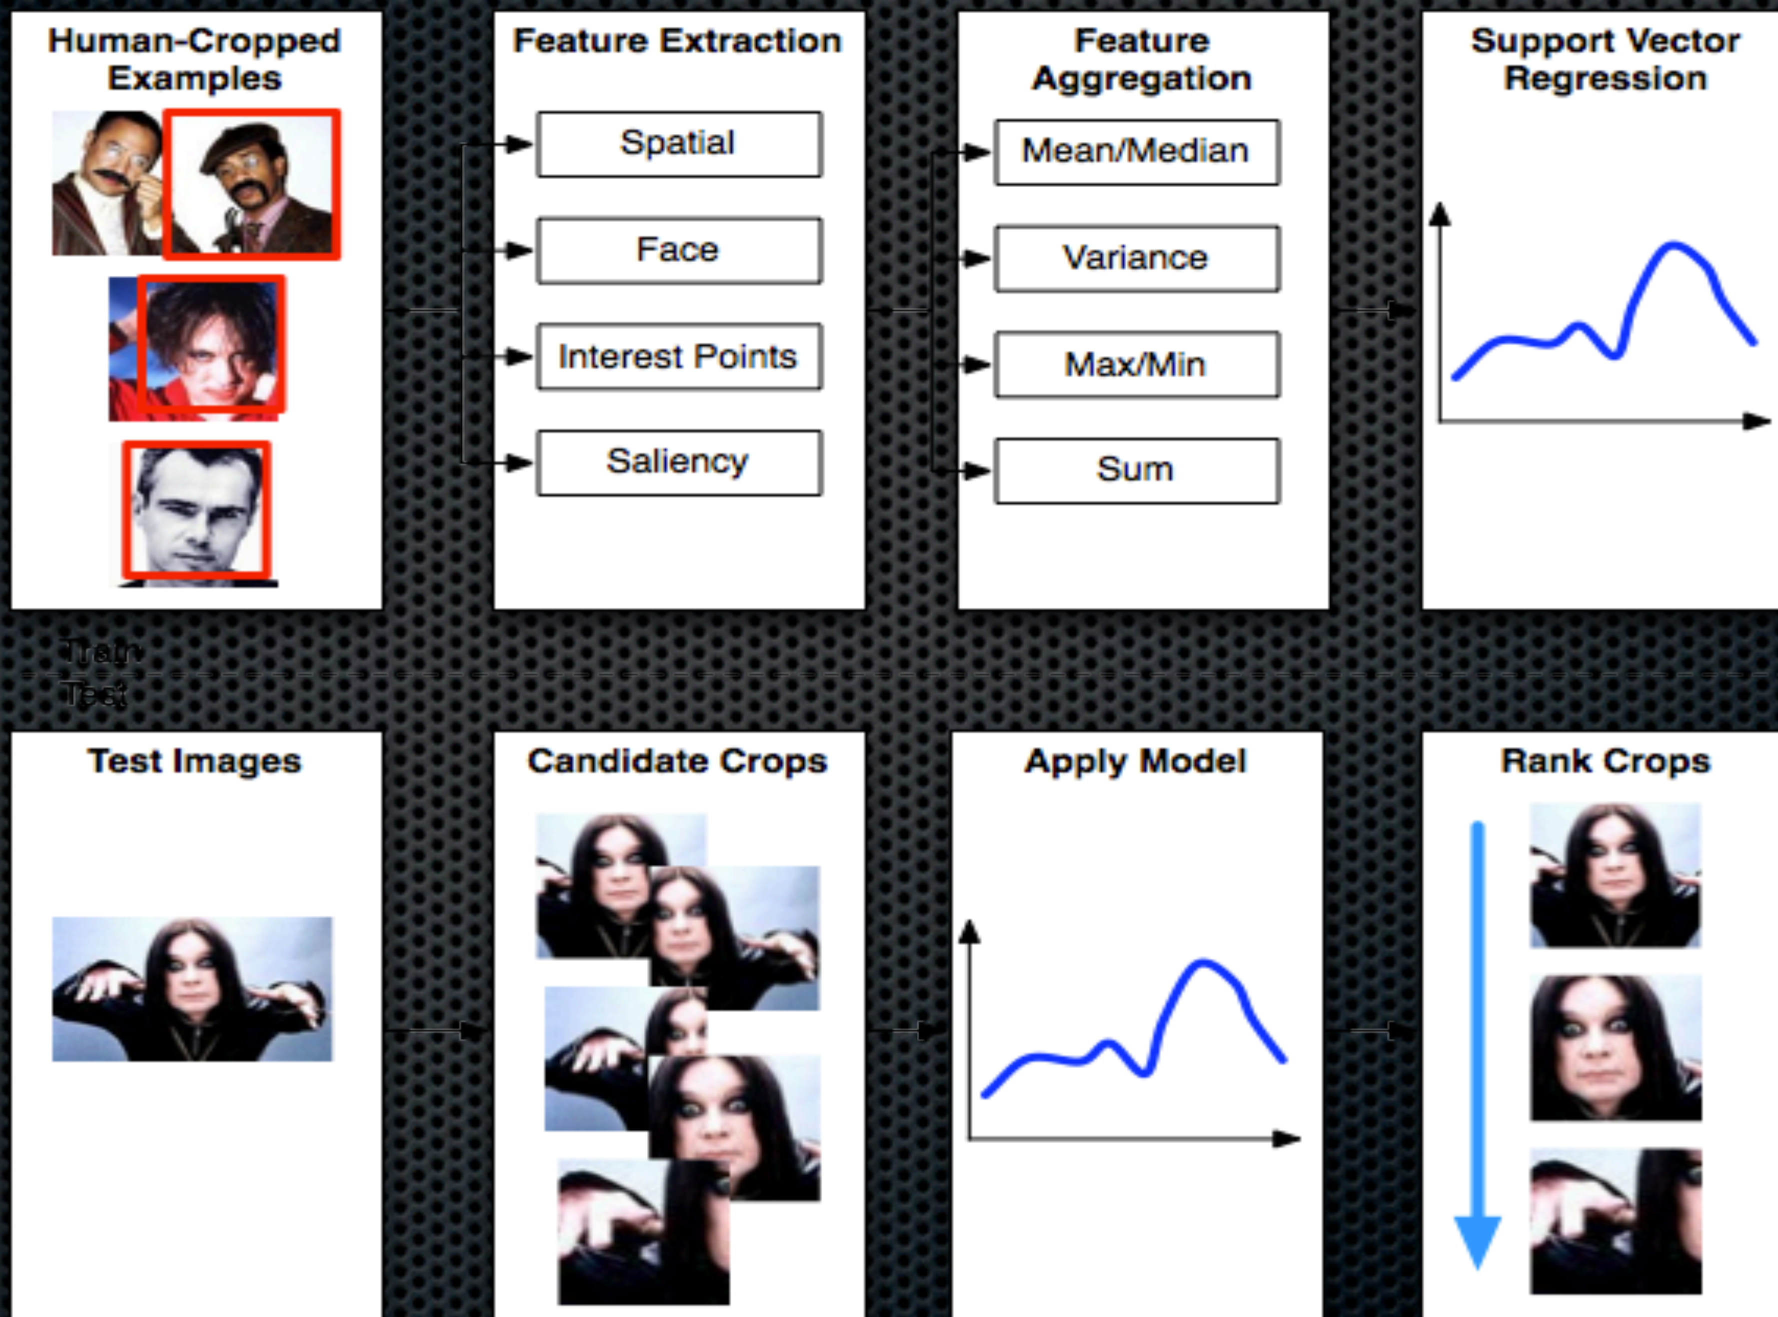

Jerry Weintraub
(person)

Logged in as **lyndonk**. Labeled 2 thumbnails so far.


☐ Accept


☐ Accept


☐ Accept


Your Thumbnail

Thumbnail coordinates:

x1: 9
y1: 7
x2: 93
y2: 91

---- Select one to reject all thumbs ----

Comments:

Submit

Categories and coverage

Category	Judgements	Inter assessor	Total
Person	5076	625	**5701**
Location	3978	625	**4603**
Movies	3000	250	**3250**
TV show	1708	250	**1958**
Sports team	738	125	**863**
Album	1500	125	**1625**
Total	16000	2000	**18000**

Performance metric

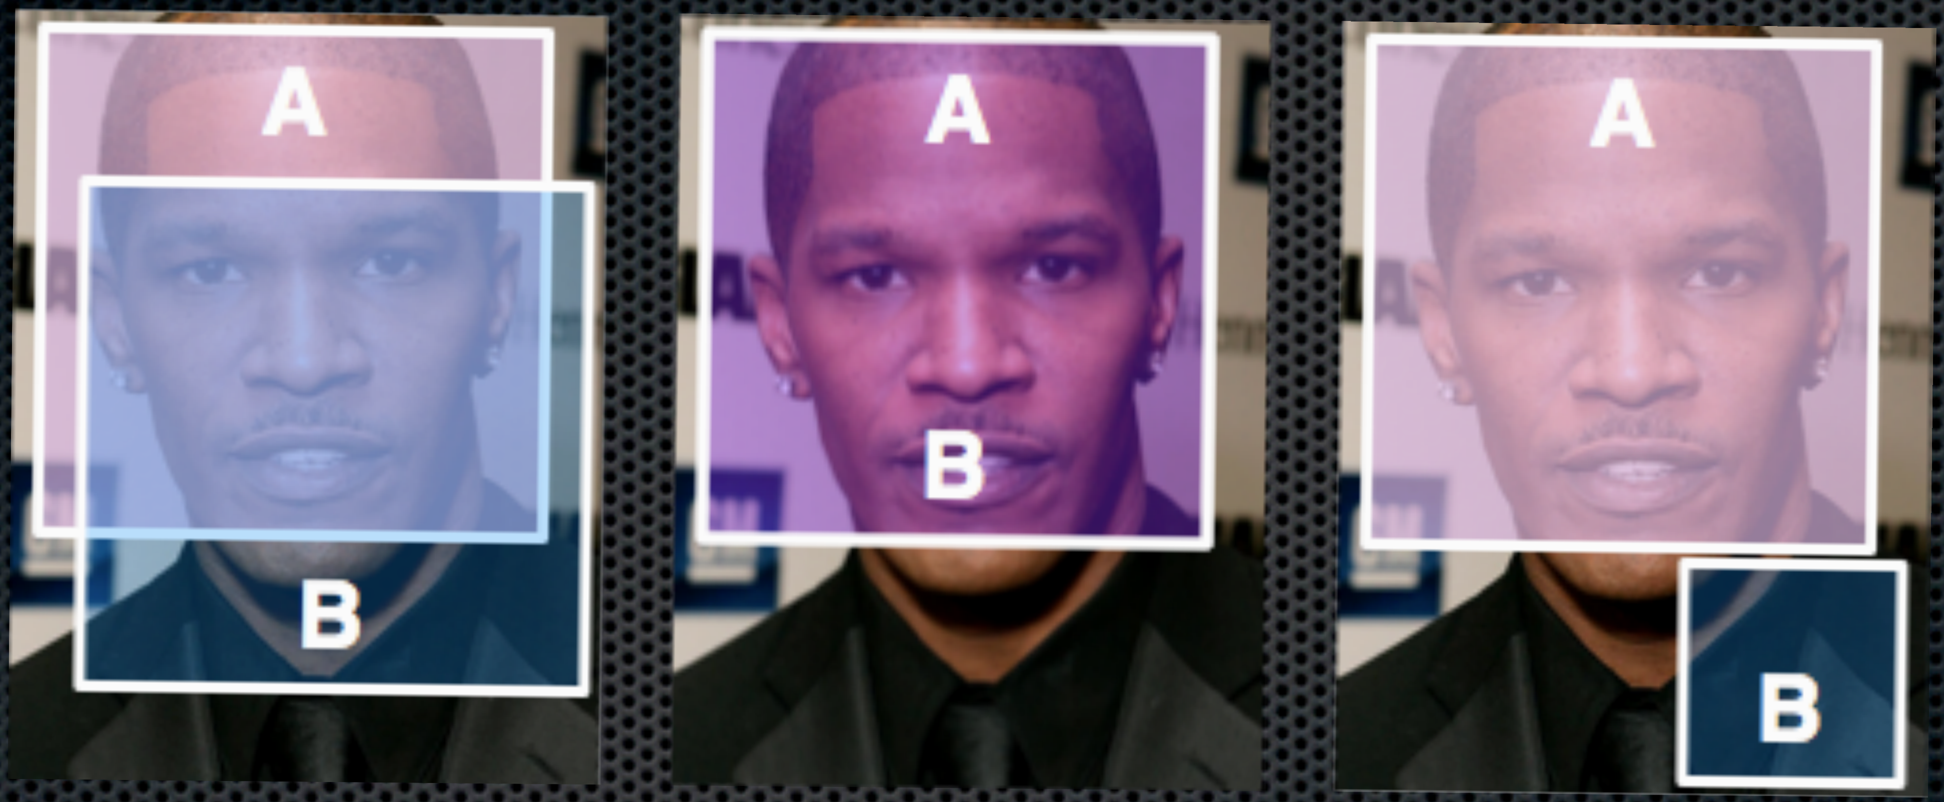

$$S_{A,B} = \frac{|I_A \bigcap I_B|}{|I_A \bigcup I_B|}$$

Feature space

- **At the PIXEL level:**
  - Spatial: horizontal, vertical, and euclidean distance from center
  - Facial: distance to face edges and centers
  - Interest points: Distance to interest point, # of adjacent points
  - Saliency Map (existing feature)

Feature space

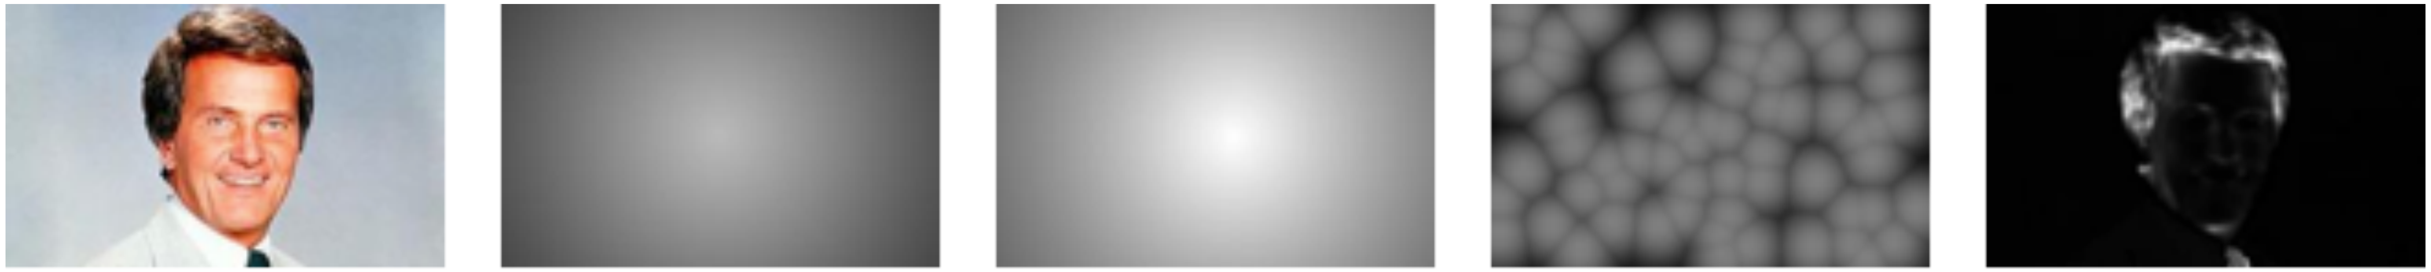

(a) Image (b) Spatial (c) Face (d) Interest Points (e) Saliency

Feature Space

- The learning objective is to predict the best crop region, e.g. not the importance of a pixel in the region.
- Introduce feature aggregation to handle the pixel features in a candidate crop region
- Each feature is represented by:
  - Min, Mean, Max, StdDev, Sum, and Median of the energy contained in crop region.

Machine Learning Setup

- Features:
  - Aggregate energy features of the candidate region
- Learning model:
  - Support Vector Machine (regression)
- Testing performance:
  - Sweep of candidate regions with variable size and steps. Select highest scoring candidate.

Results

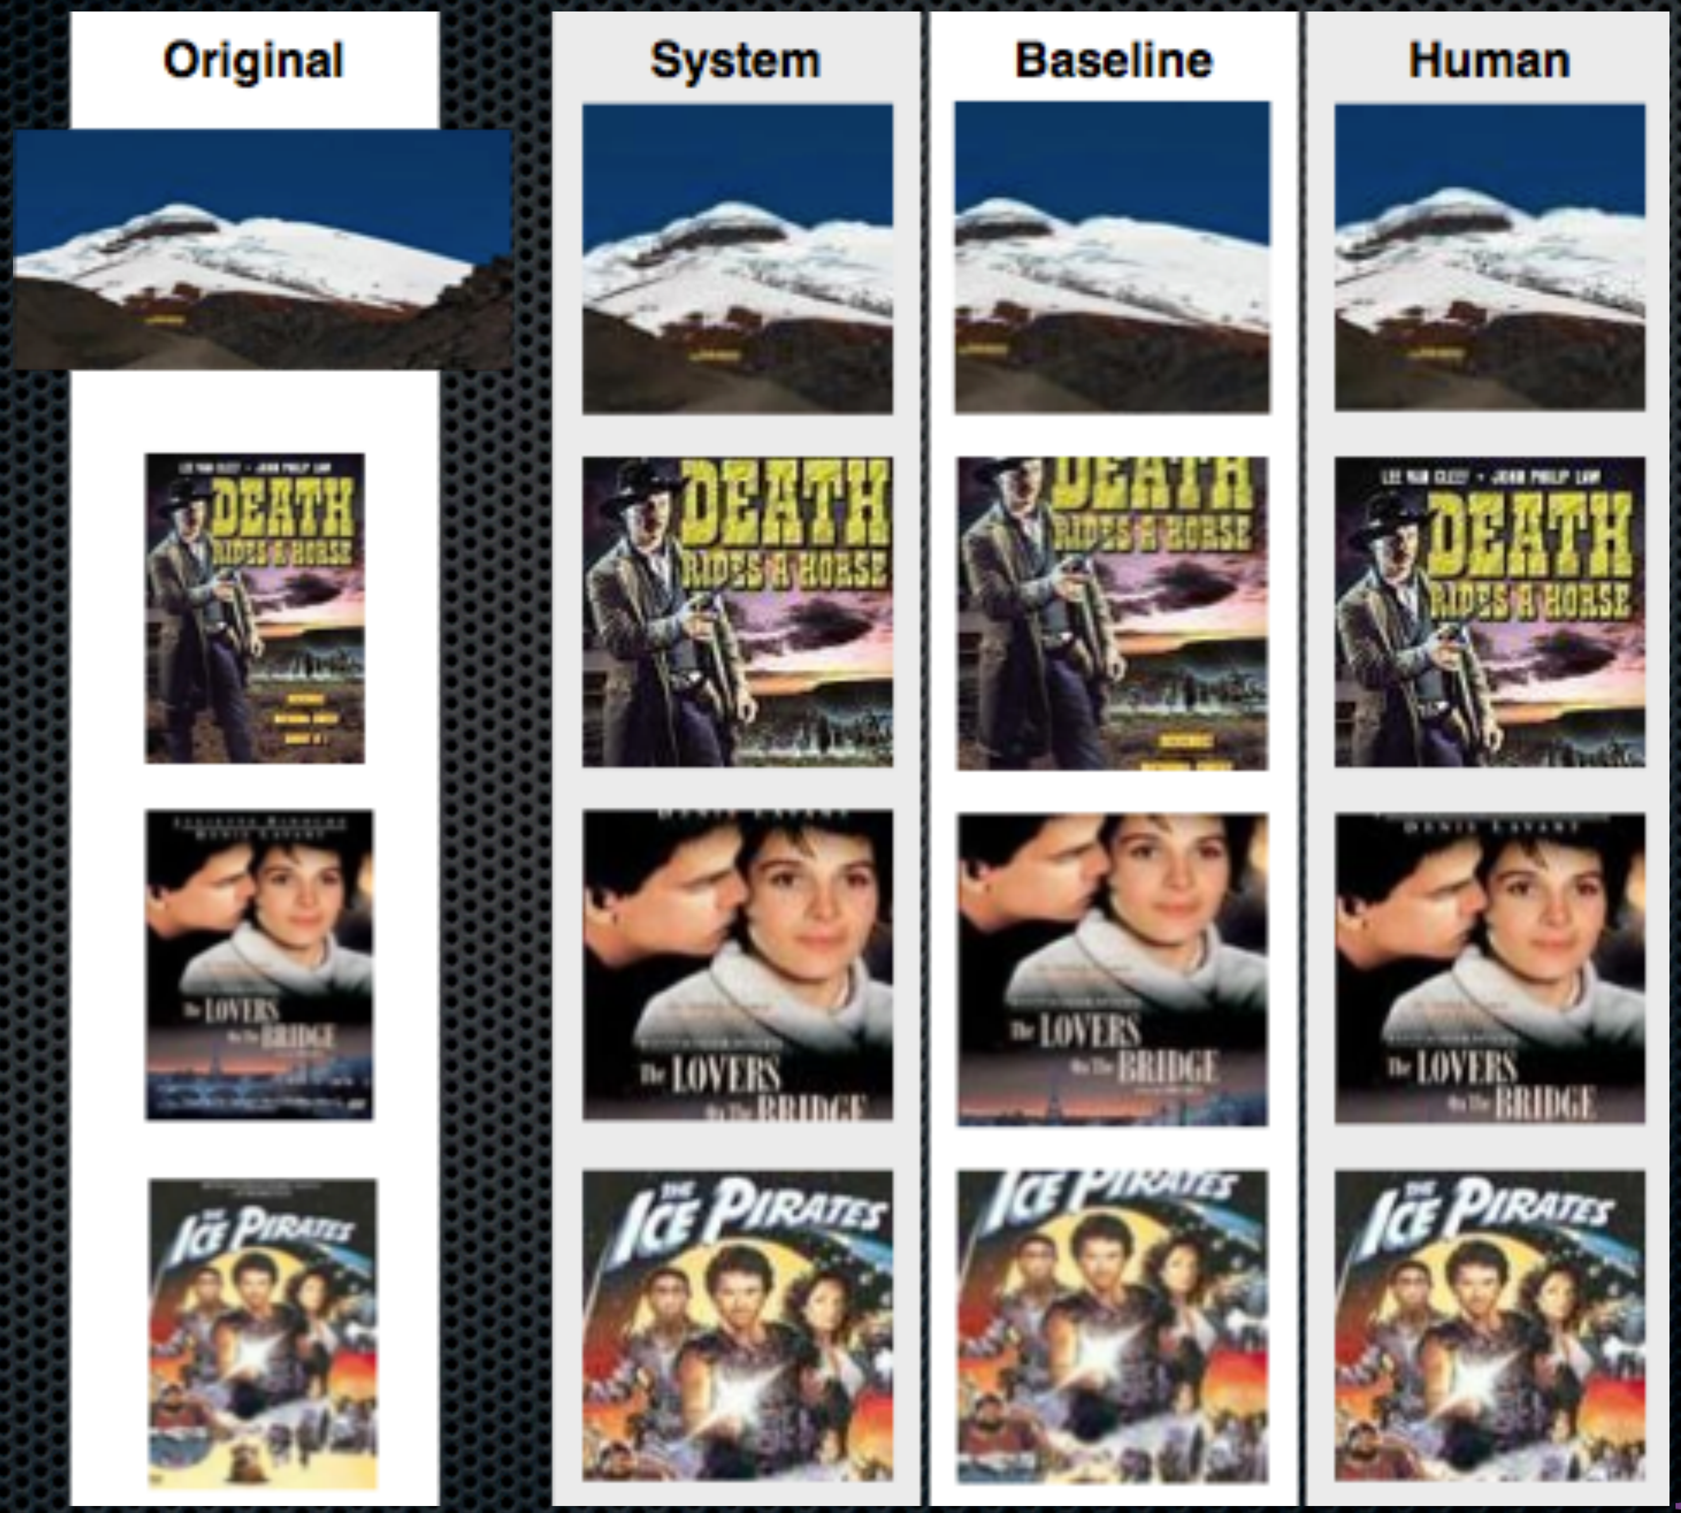
Original
System
Baseline
Human
DEATH
RIDES A HORSE
The LOVERS
on The BRIDGE
ICE PIRATES

Results

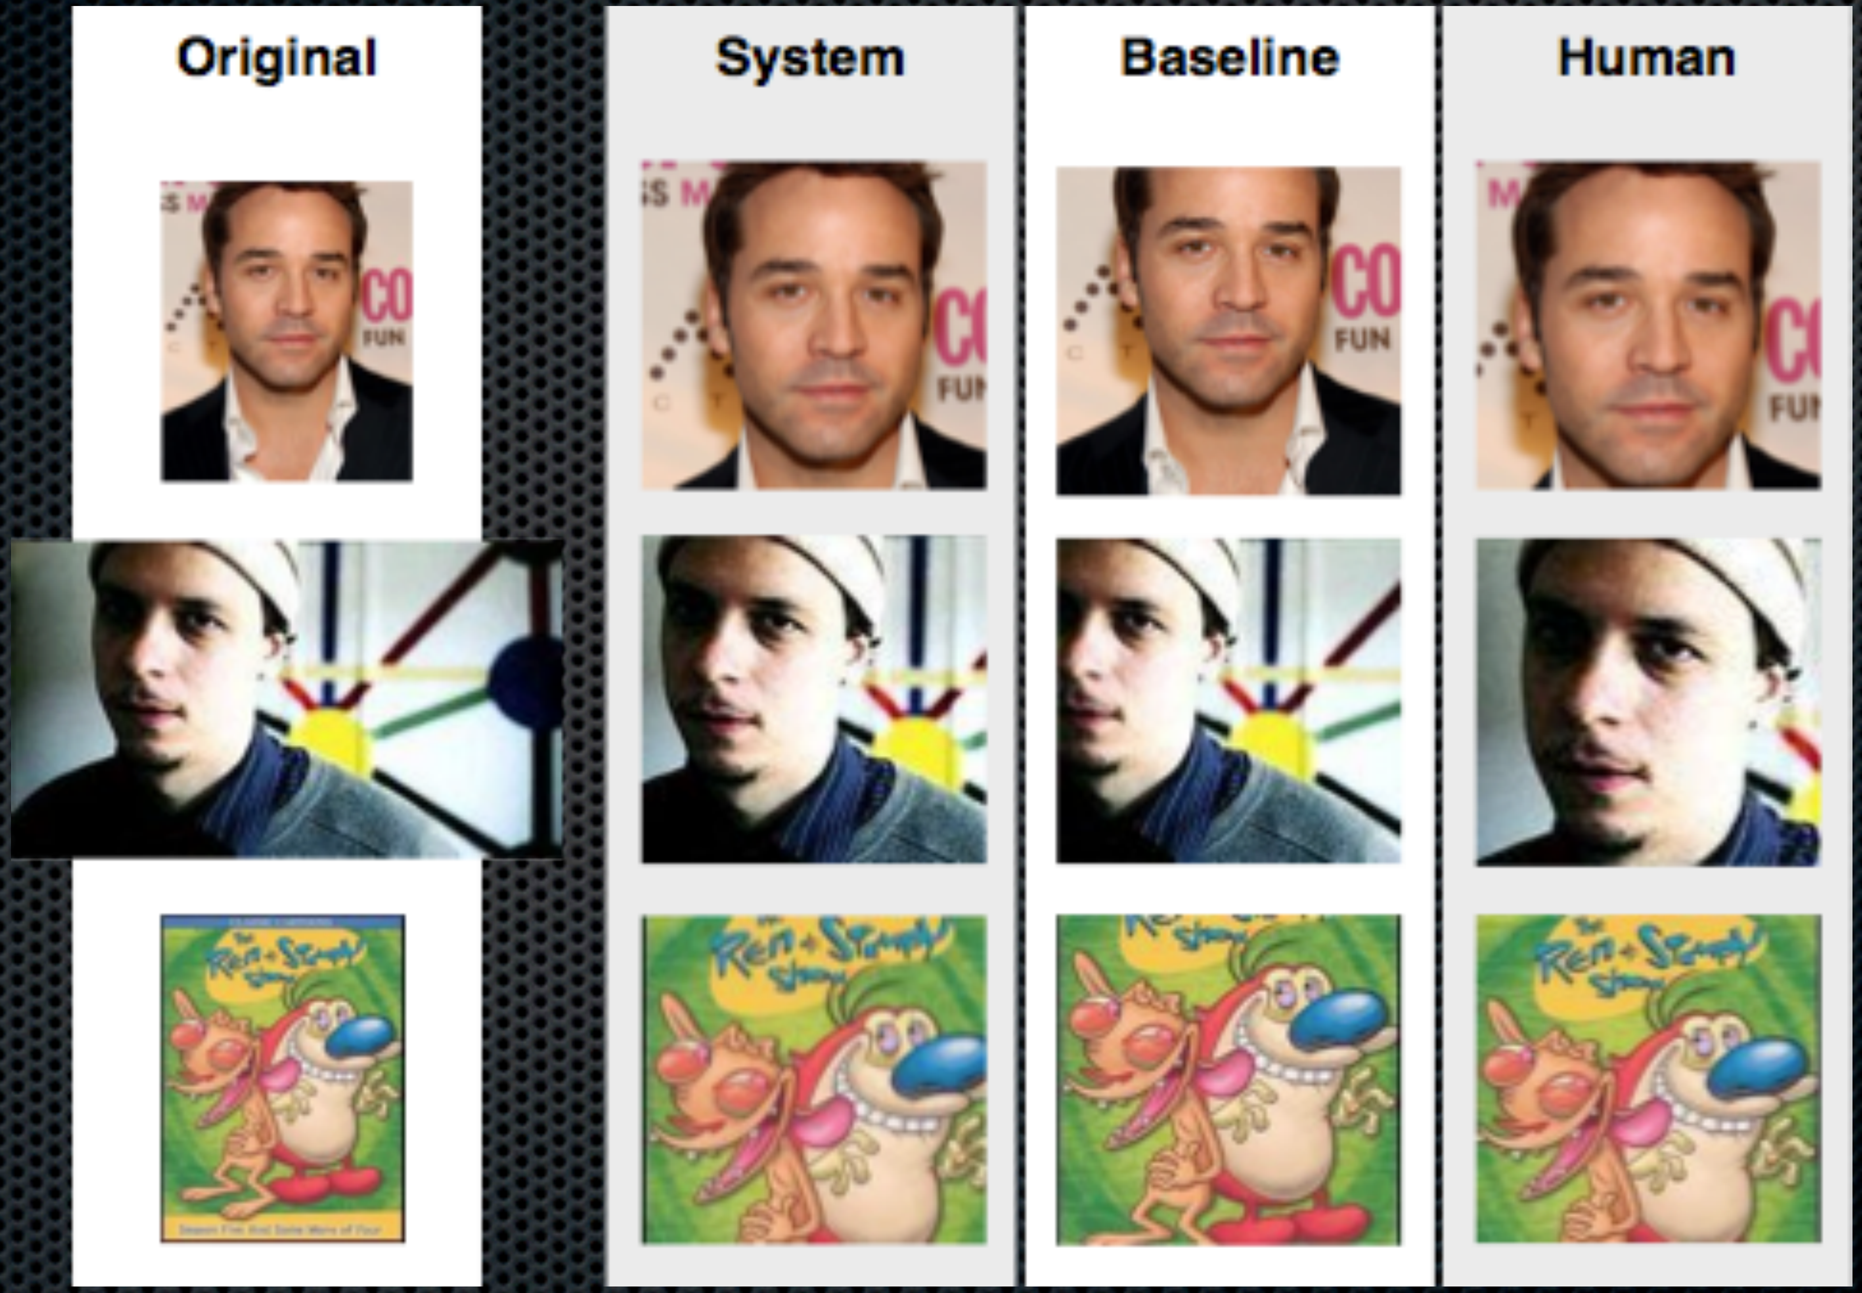

Performance per category

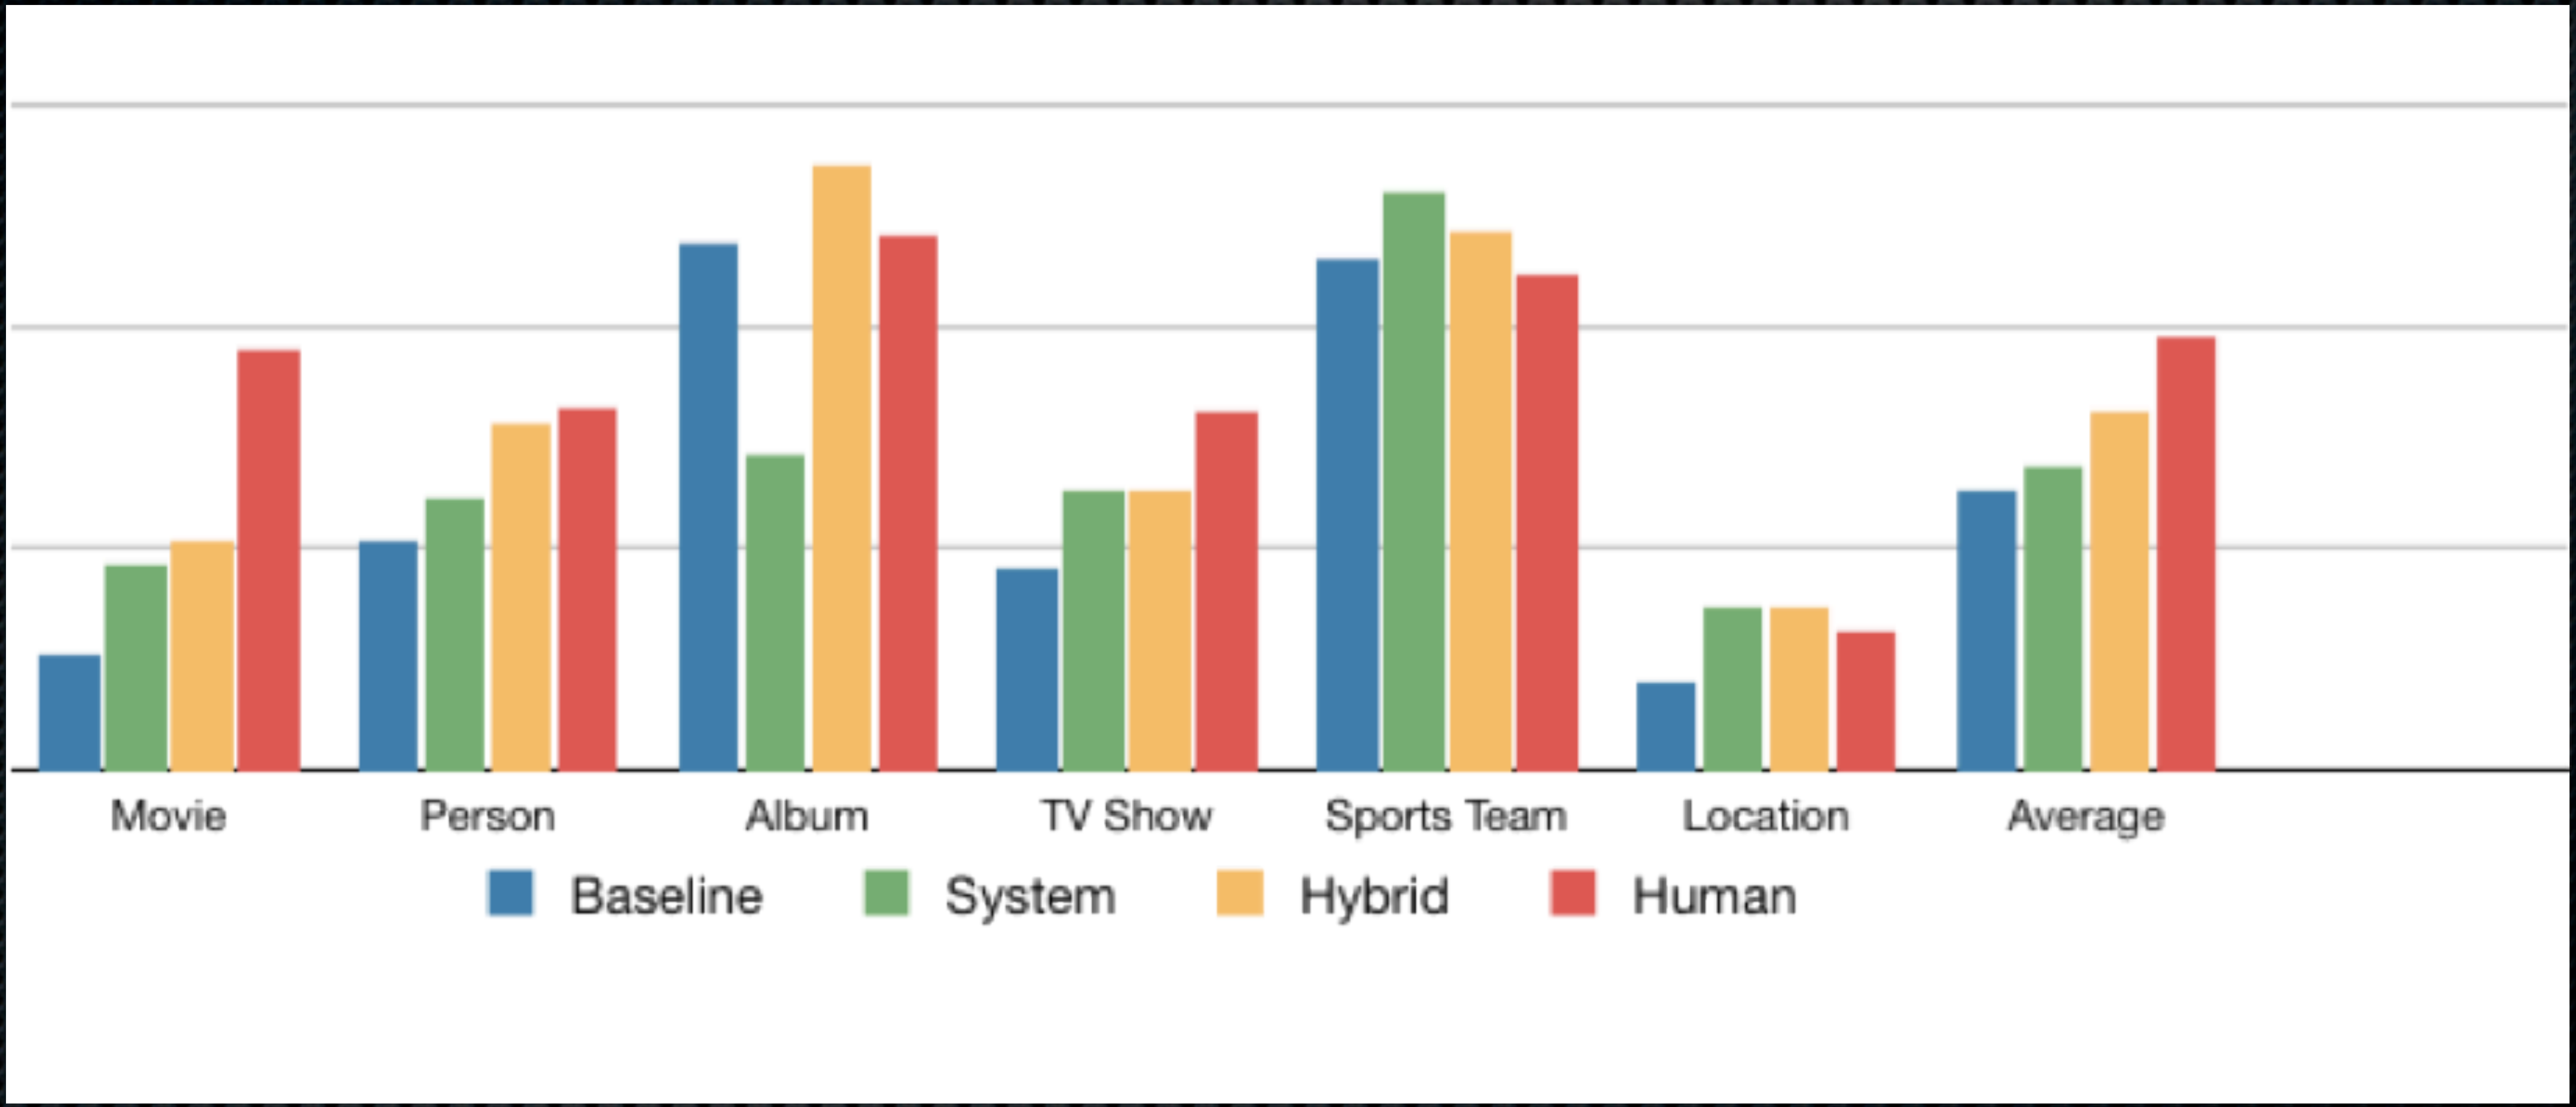

Performance

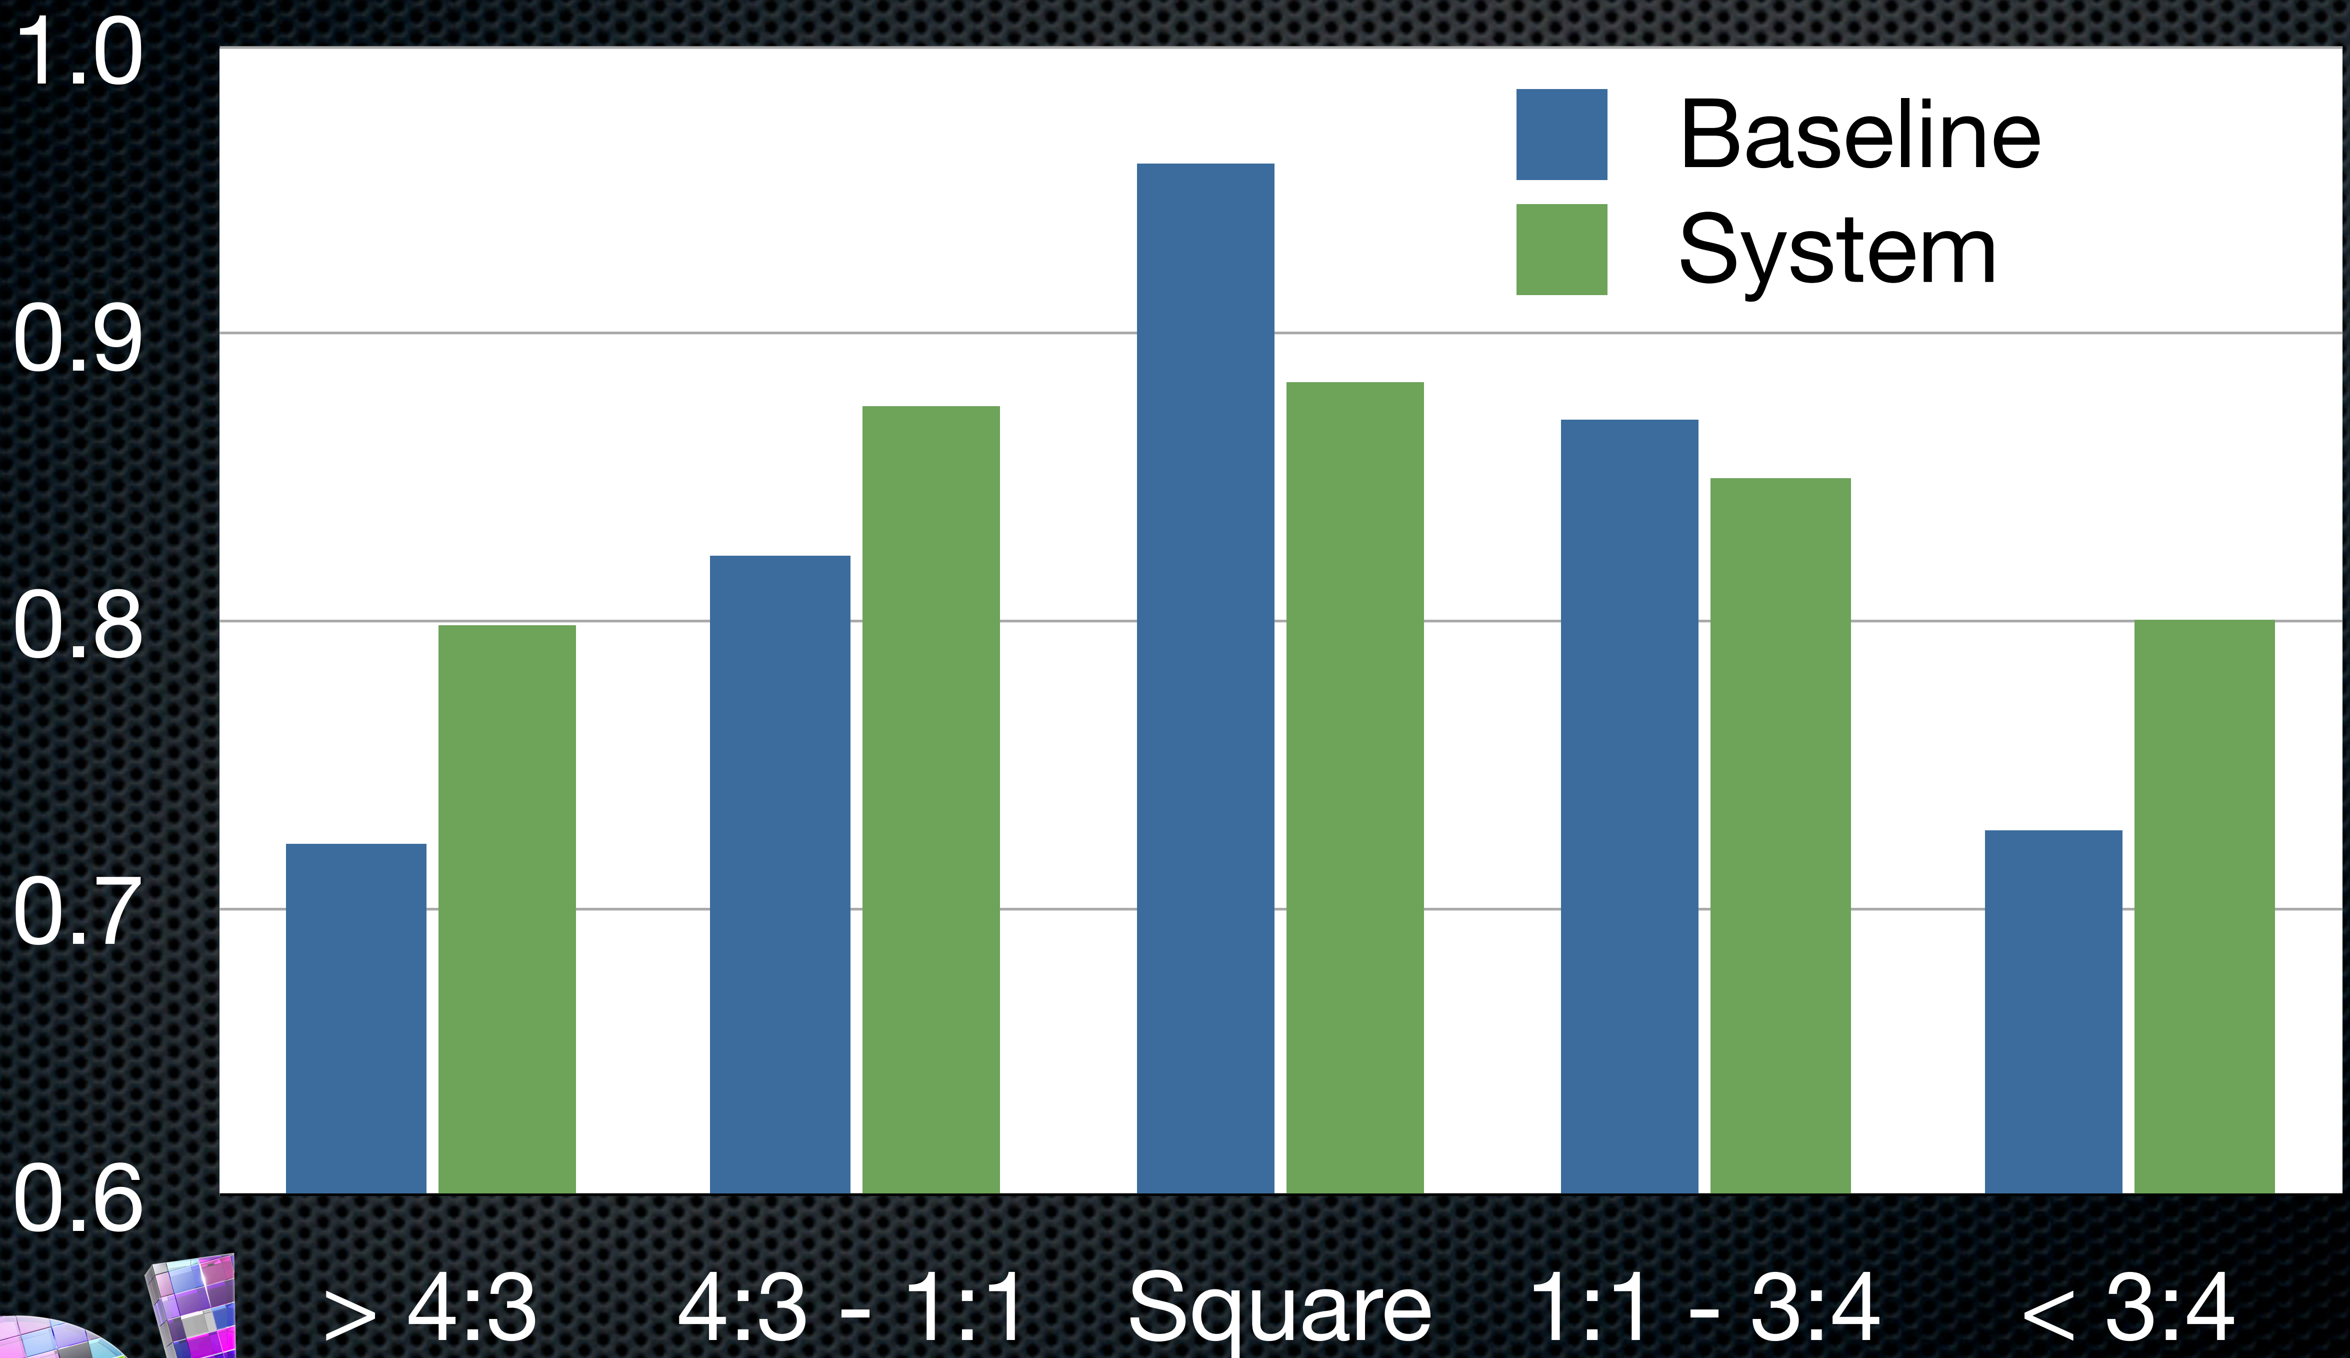

Summary

- Select thumbnail regions by visual/spatial cues
- System out-performs baseline and is inching toward human performance
- Unique aspects
  - Fusion of diverse cues (spatial, face, interest, etc)
  - Learning-driven architecture (approximates editor decisions)

Explore-exploit

Optimizing for engagement

Joint work with:
Lyndon Kennedy,
Malcolm Slaney, and
Nicolas Torzec

YAHOO!


YAHOO!
Web Images Video Local Shopping News More
George Clooney
Search
QuickApps
Search Pad
SafeSearch - On
5,260,000 results for
George Clooney
Related People
Brad Pitt
Matt Damon
Renee Zellweger
Frances McDorma...
Meryl Streep
Coen brothers
Latest On: Jennifer Anis... Nicolas Cage Evangeline Li... Abigail Bresl...
Overview
George Clooney Takes On Wall Street And ...
Irish Central - Apr 15 07:23am
George Clooney is set to produce, and possibly direct, a movie about the 2008 financial crisis and the government bailout of troubled financial institutions. Full Story »
More George Clooney Stories
George Clooney Photos
More Photos »
George Clooney - Wikipedia, the free encyclopedia
Early life | Career | Humanitarian work | In the media
George Timothy Clooney is an American actor, film director, producer, and screenwriter. For his work as an actor, he has received two Golden Globe Awards and an Academy Award. Clooney...
en.wikipedia.org/wiki/George_Clooney - Cached
George Clooney: Biography from Answers.com
Born: 6 May 1961 Birthplace: Lexington, Kentucky Best Known As: The international star of

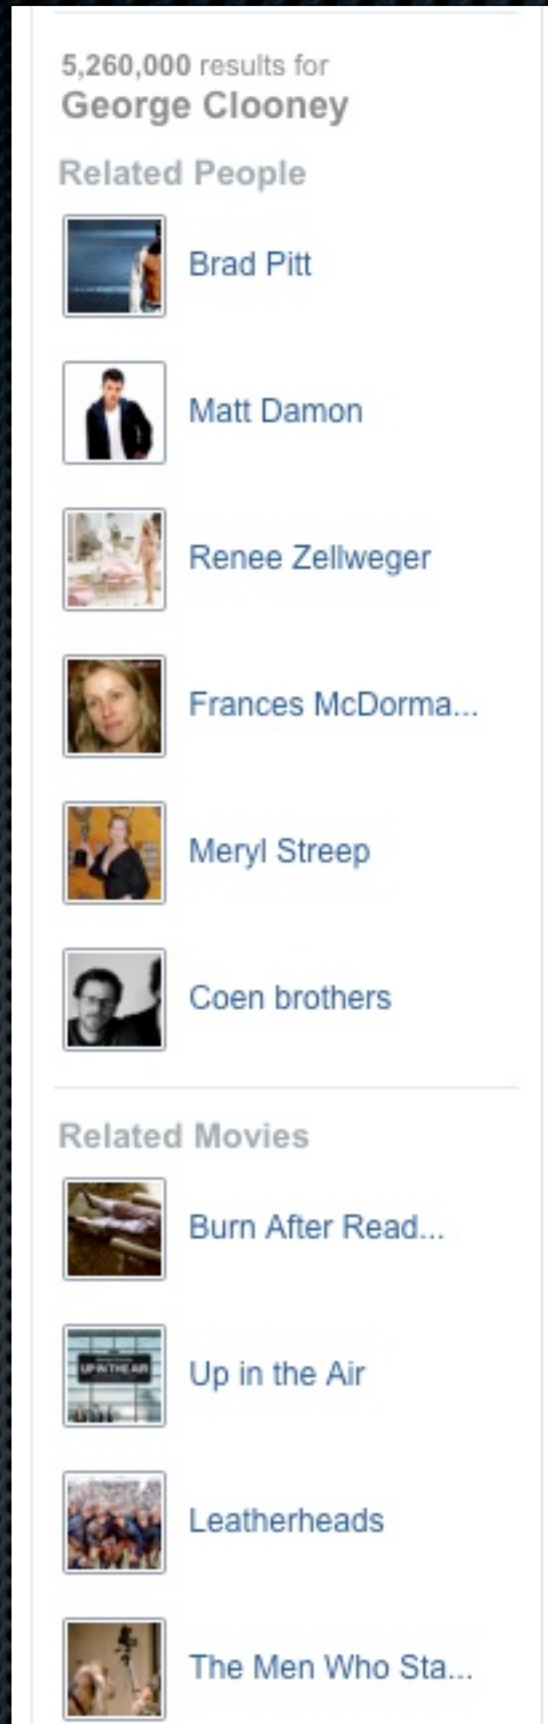
5,260,000 results for
George Clooney
Related People
Brad Pitt
Matt Damon
Renee Zellweger
Frances McDorma...
Meryl Streep
Coen brothers
Related Movies
Burn After Read...
Up in the Air
Leatherheads
The Men Who Sta...


George Clooney
Related People
Brad Pitt
Renee Zellweger
Meryl Streep
Coen brothers
Burn After Read
Up in the Air

Explore/Exploit for Web Content Optimization

- By: Deepak Agarwal, Bee-Chung Chen, Pradheep Elango (Yahoo! Research) at ICDM'09
- **Multi-armed bandit: maximize total clicks on a content module**
- Intuition: "Explore" each candidate item by displaying it to a small fraction of user visits to estimate the item's click-through rate (CTR), and then "exploit" high CTR items in order to maximize clicks.

Explore/exploit for images

- Collect click-feedback on left-rail search suggestions
- Random rotation of image thumbnails (***3 per facet***, ***35x35 px***)
- During 47 days, using fraction of traffic
- Ensured that ranking of facets is constant, and image pool is constant (e.g. facets can be suggested for different queries, at different positions)
- **Objective: Maximize the number of clicks on the**

Click-feedback collected

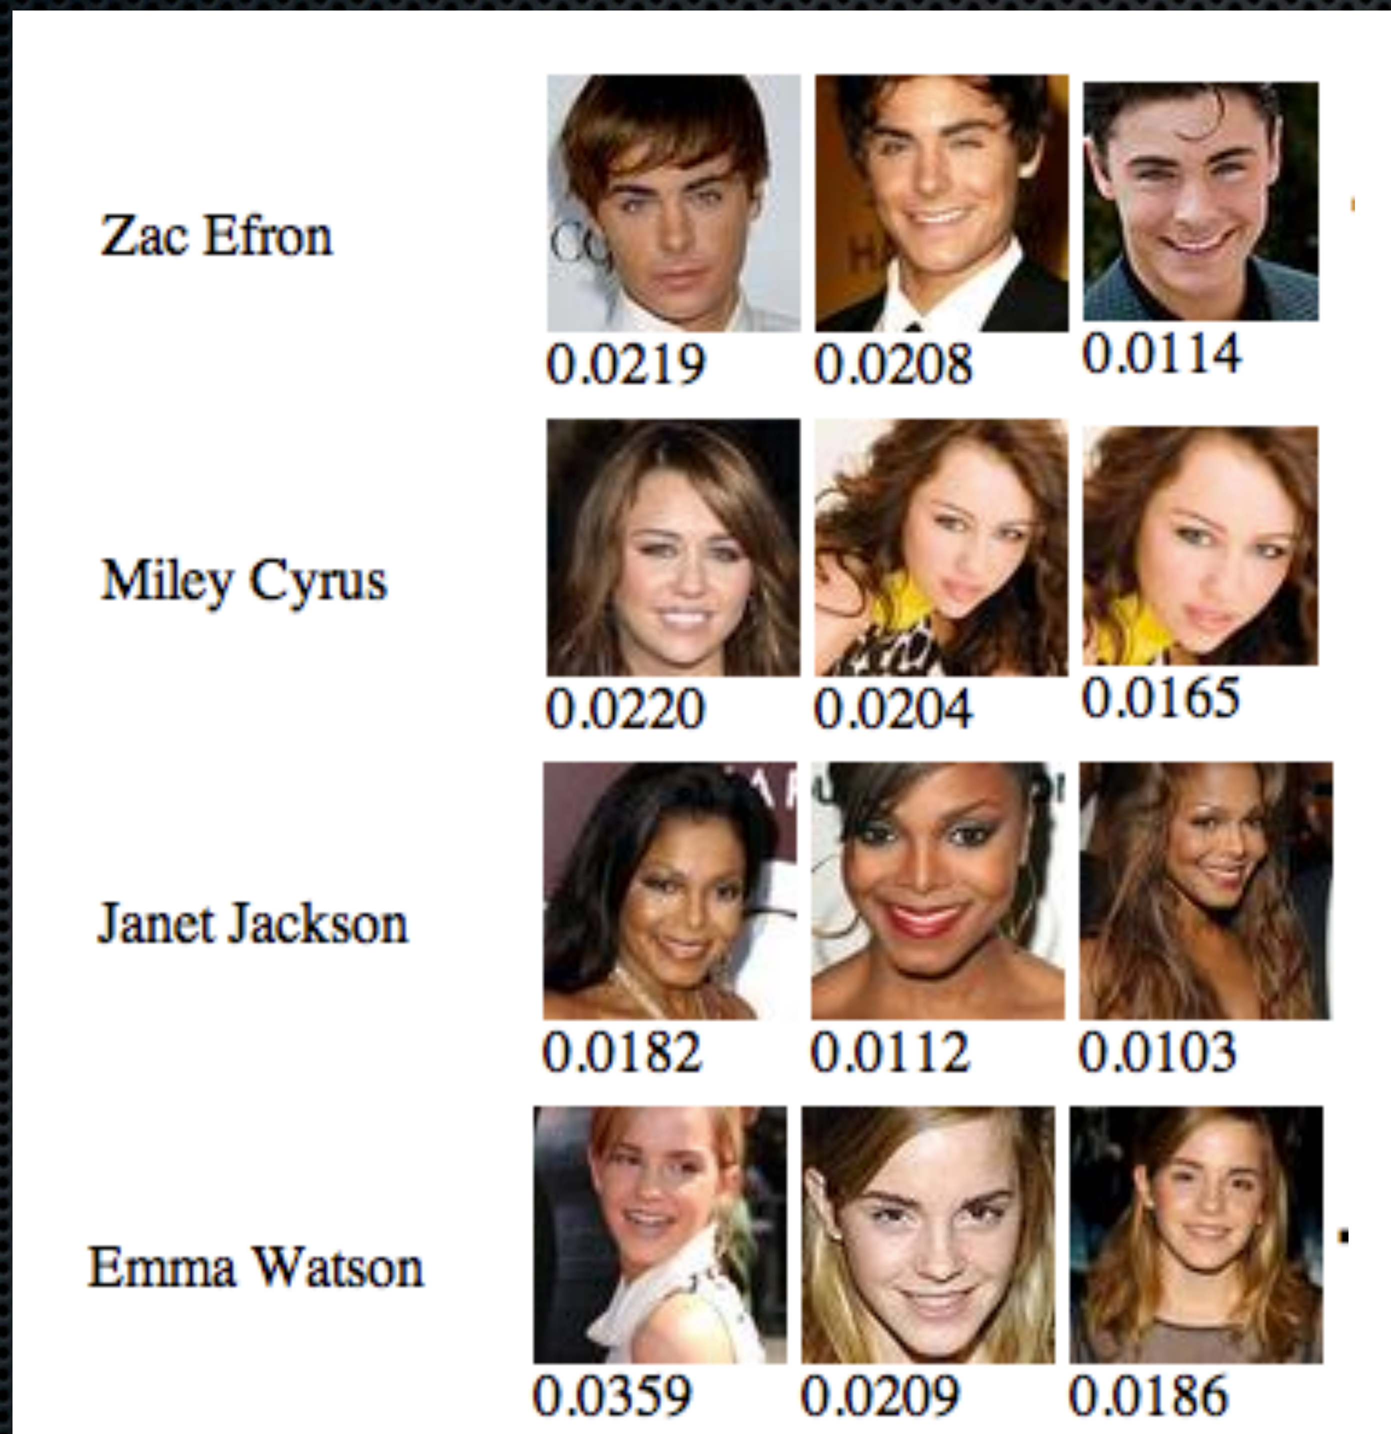

YAHOO!

Gain

Method	Clicks	
Bayes-relax	102.1	winner (38% lift)
Bayes-relax (post)	109.7	cheating
Random	73.6	baseline

Method	Clicks	
Bayes-relax	391.4	winner (20%)
Bayes-relax (post)	408.5	cheating
Random	324.1	baseline

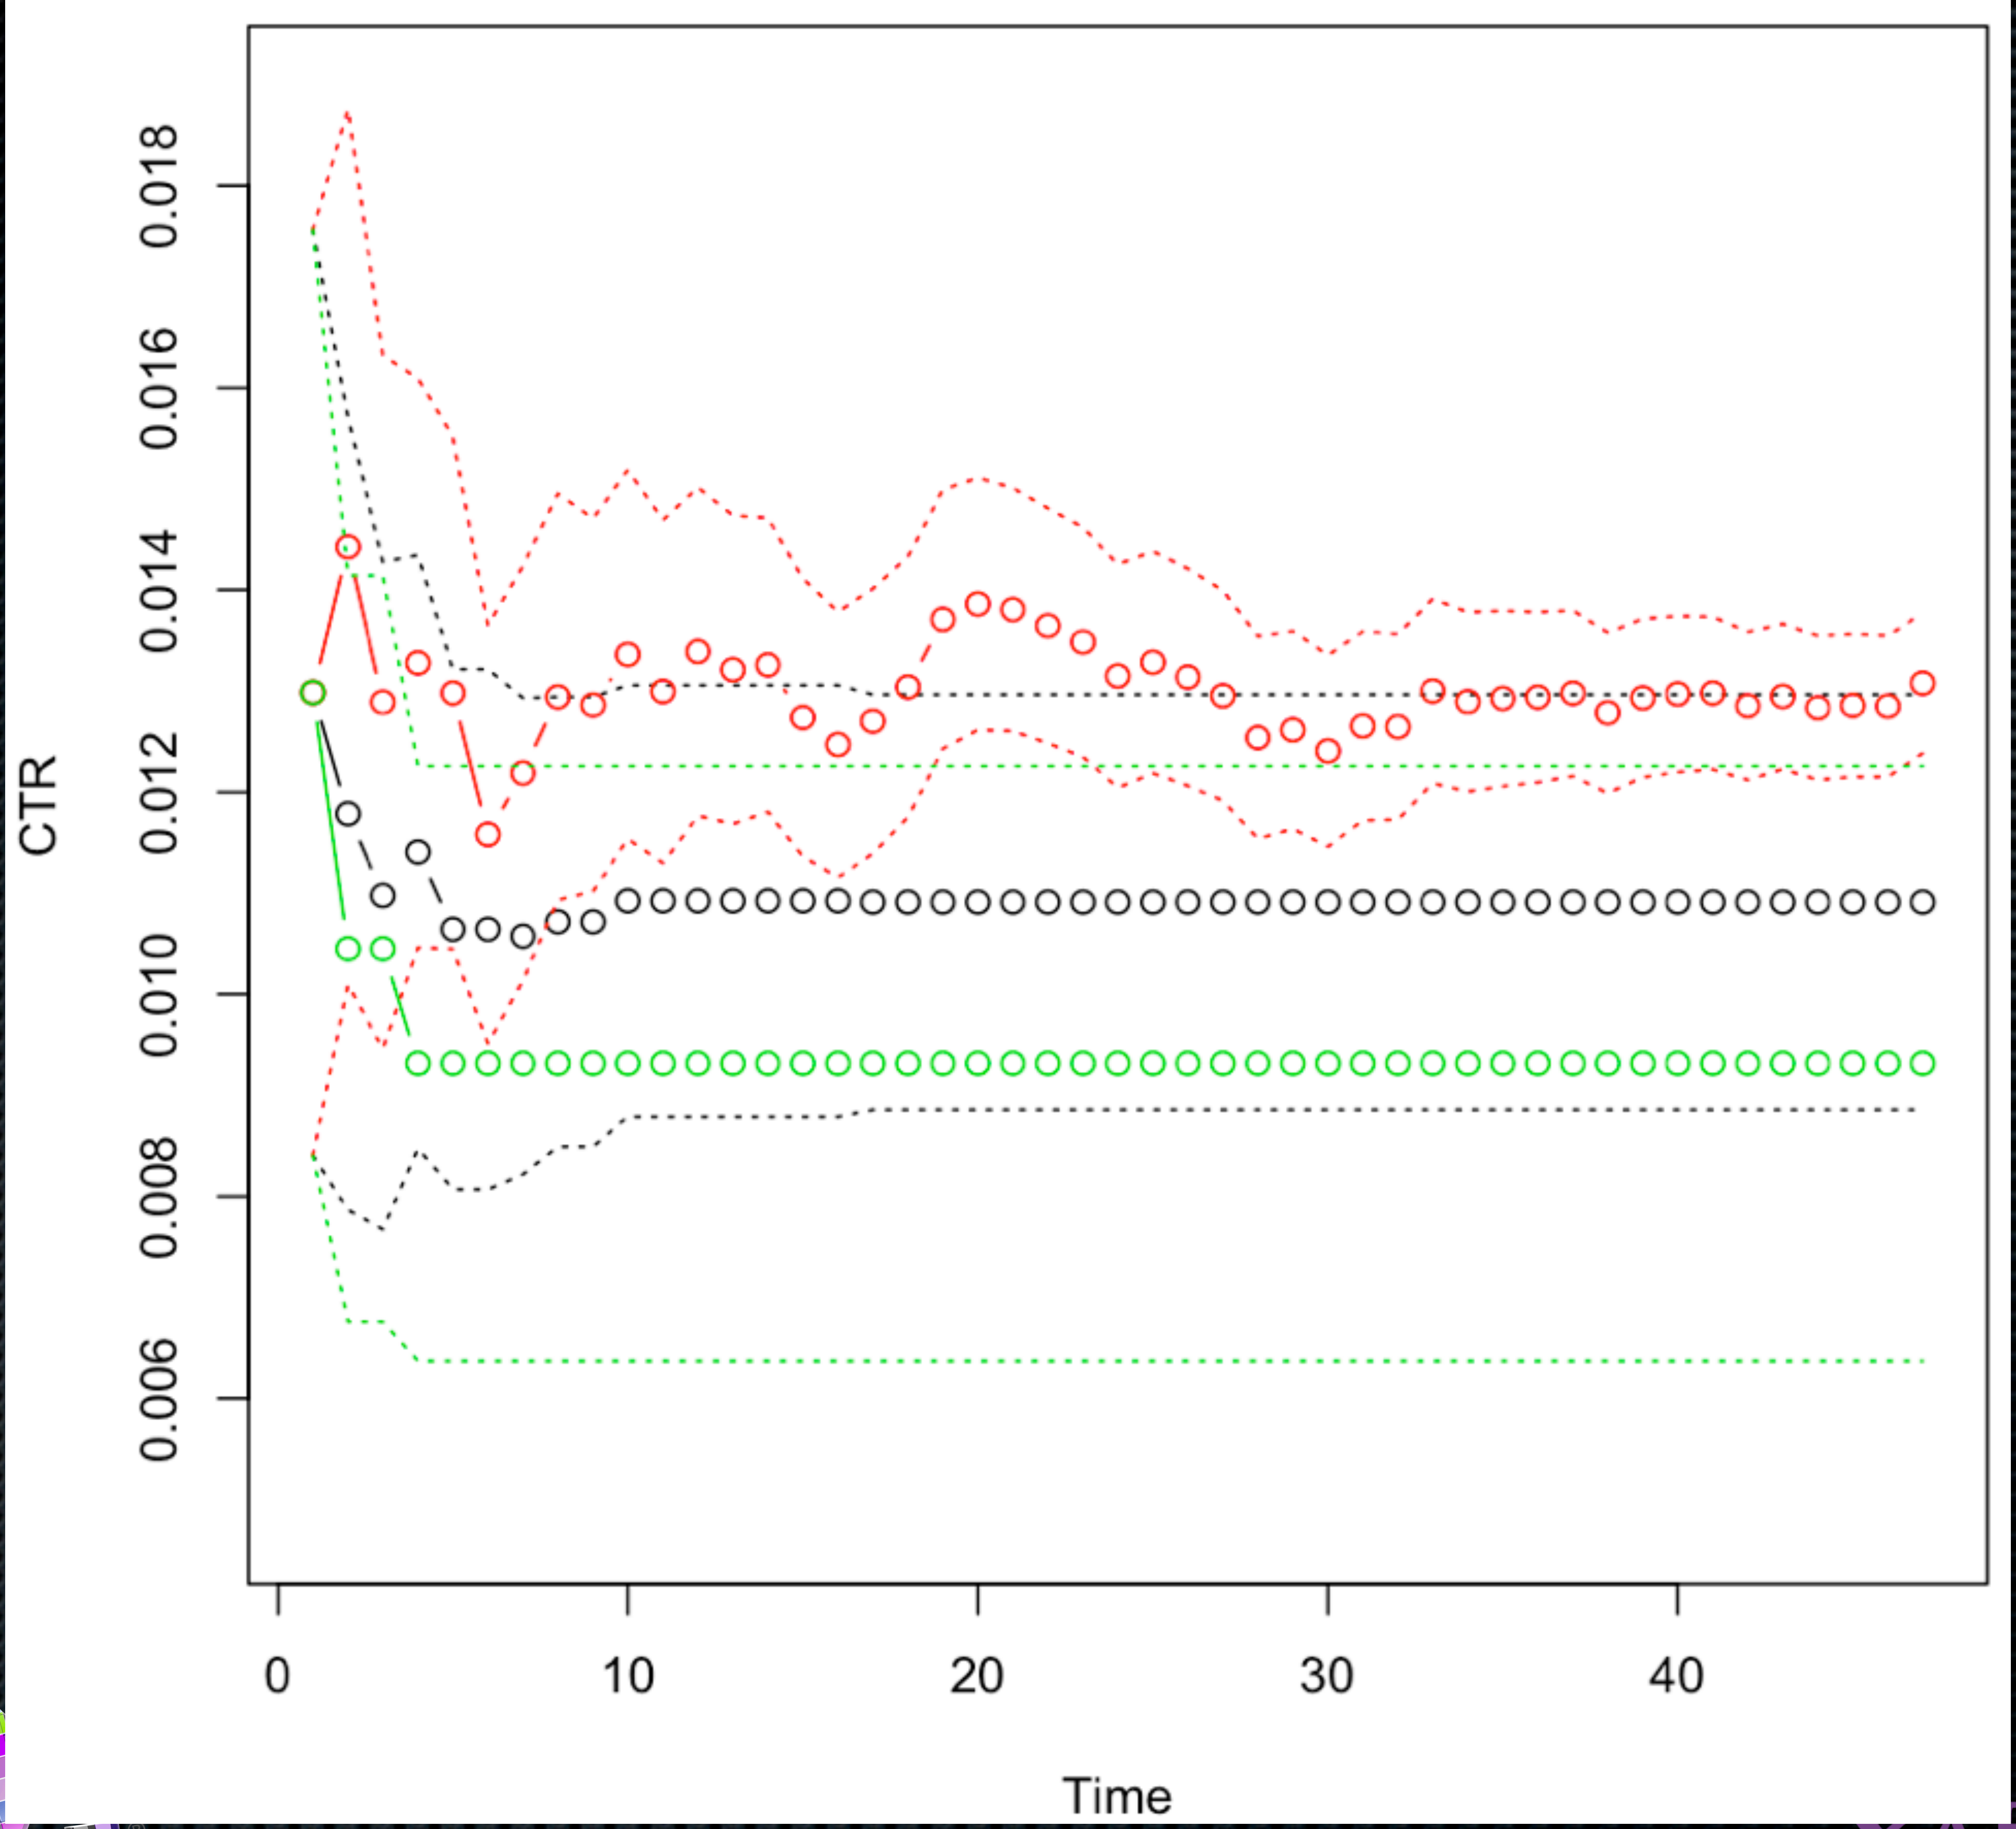
CTR
0.006
0.008
0.010
0.012
0.014
0.016
0.018
0
10
20
30
40
Time

Predicting Flickr Favorites

Joint work with: Lluis Garcia, and Adam Rae

What triggers you to call a photo a favourite on Flickr?

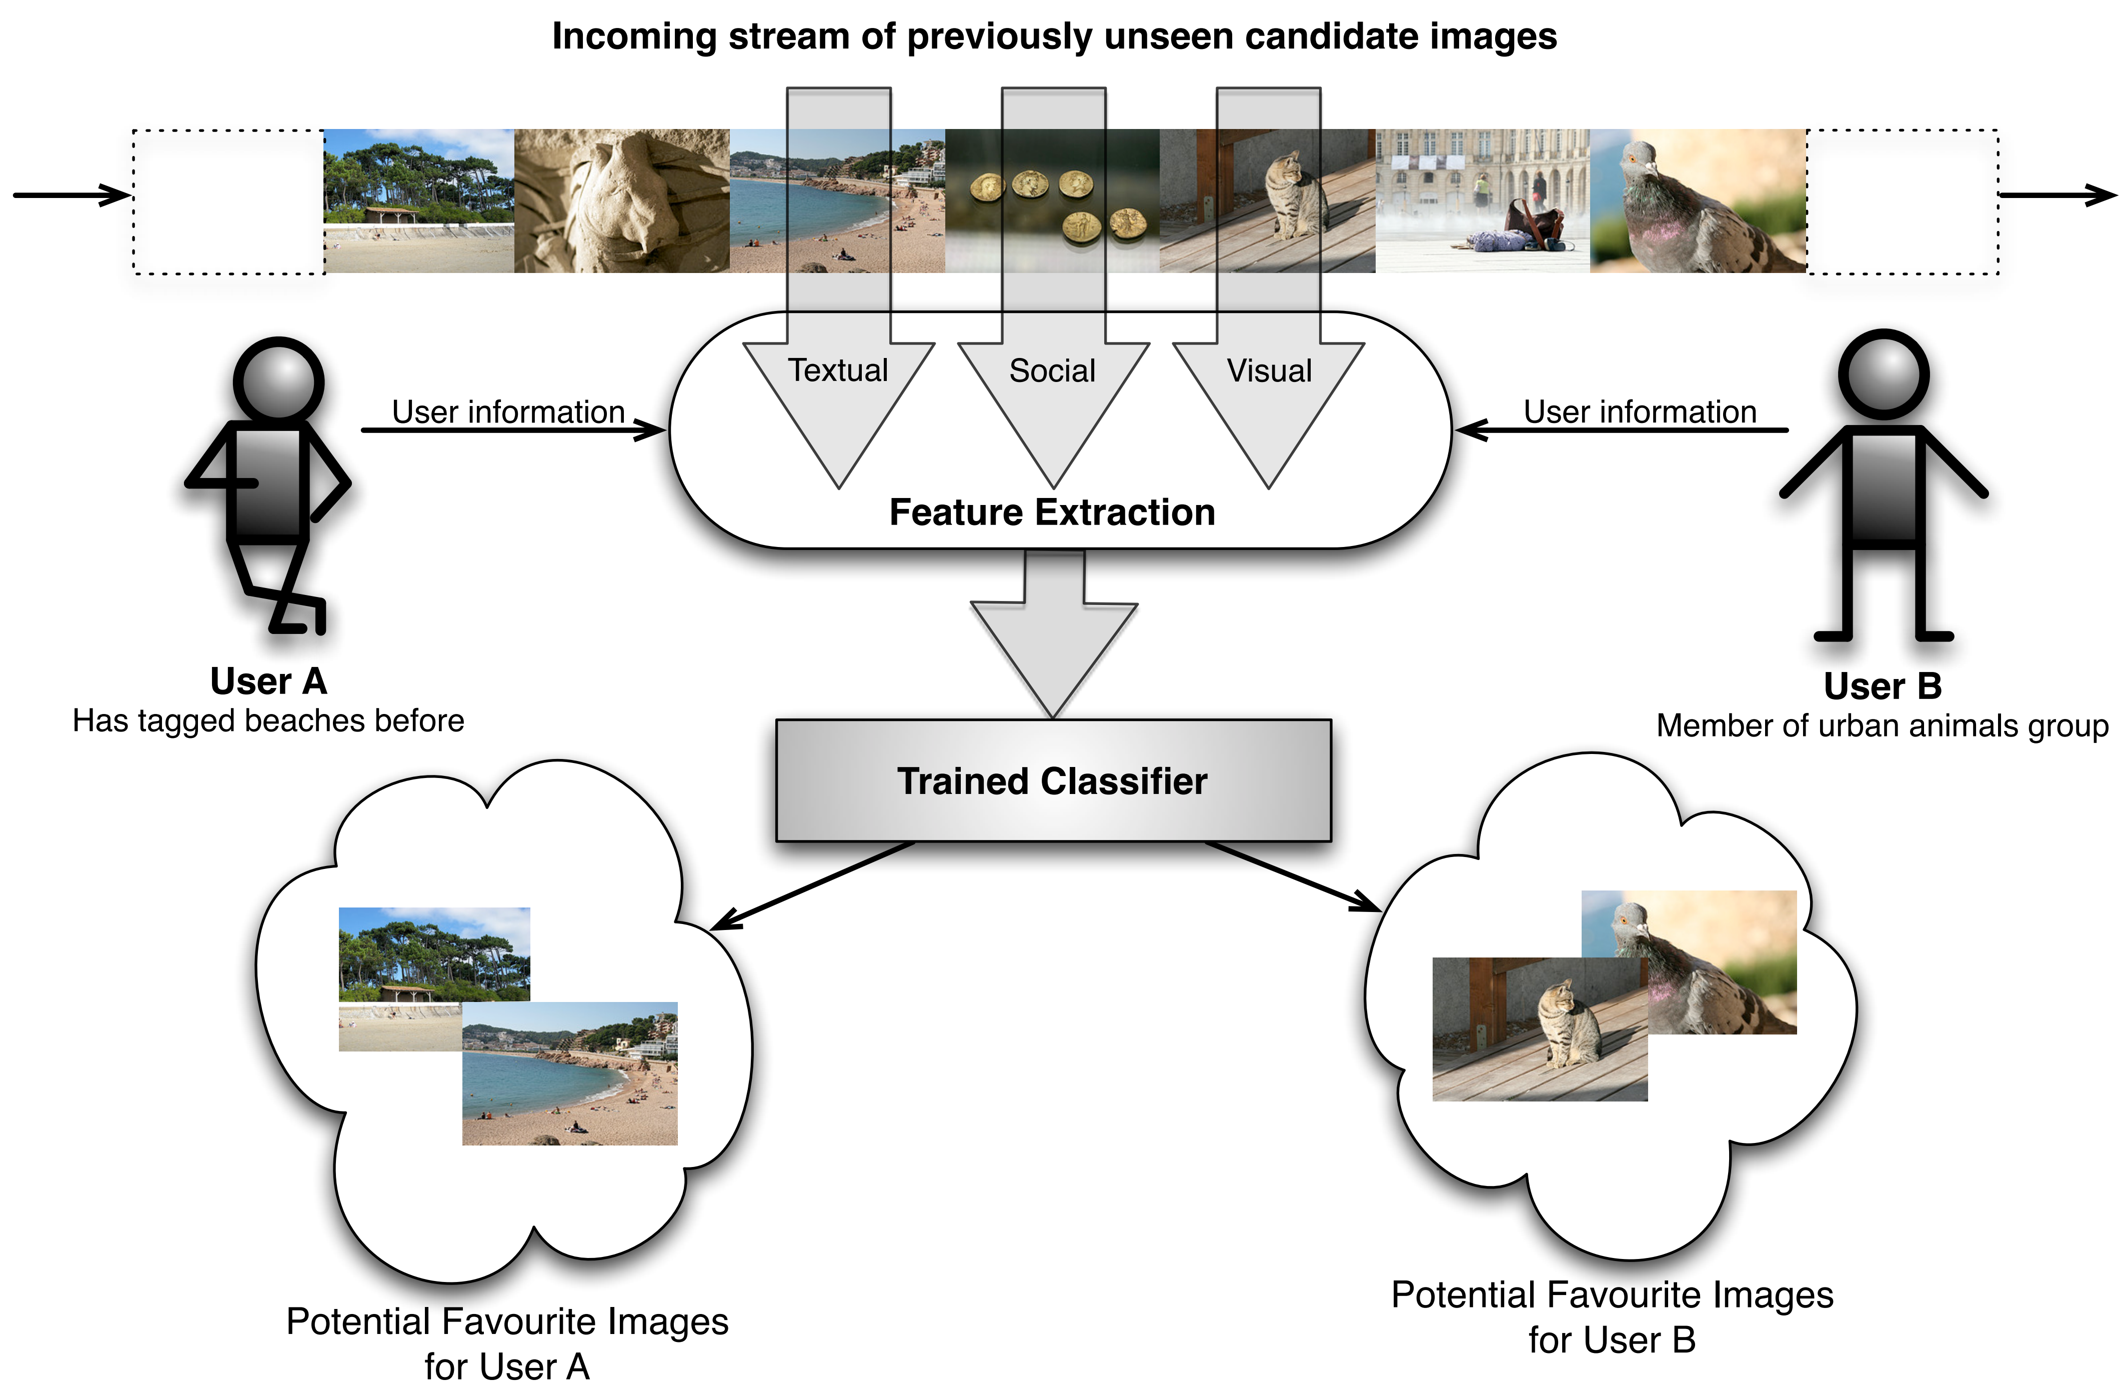
Incoming stream of previously unseen candidate images
Textual
Social
Visual
Feature Extraction
User information
User information
User A
Has tagged beaches before
User B
Member of urban animals group
Trained Classifier
Potential Favourite Images
for User A
Potential Favourite Images
for User B

Distribution of favorites

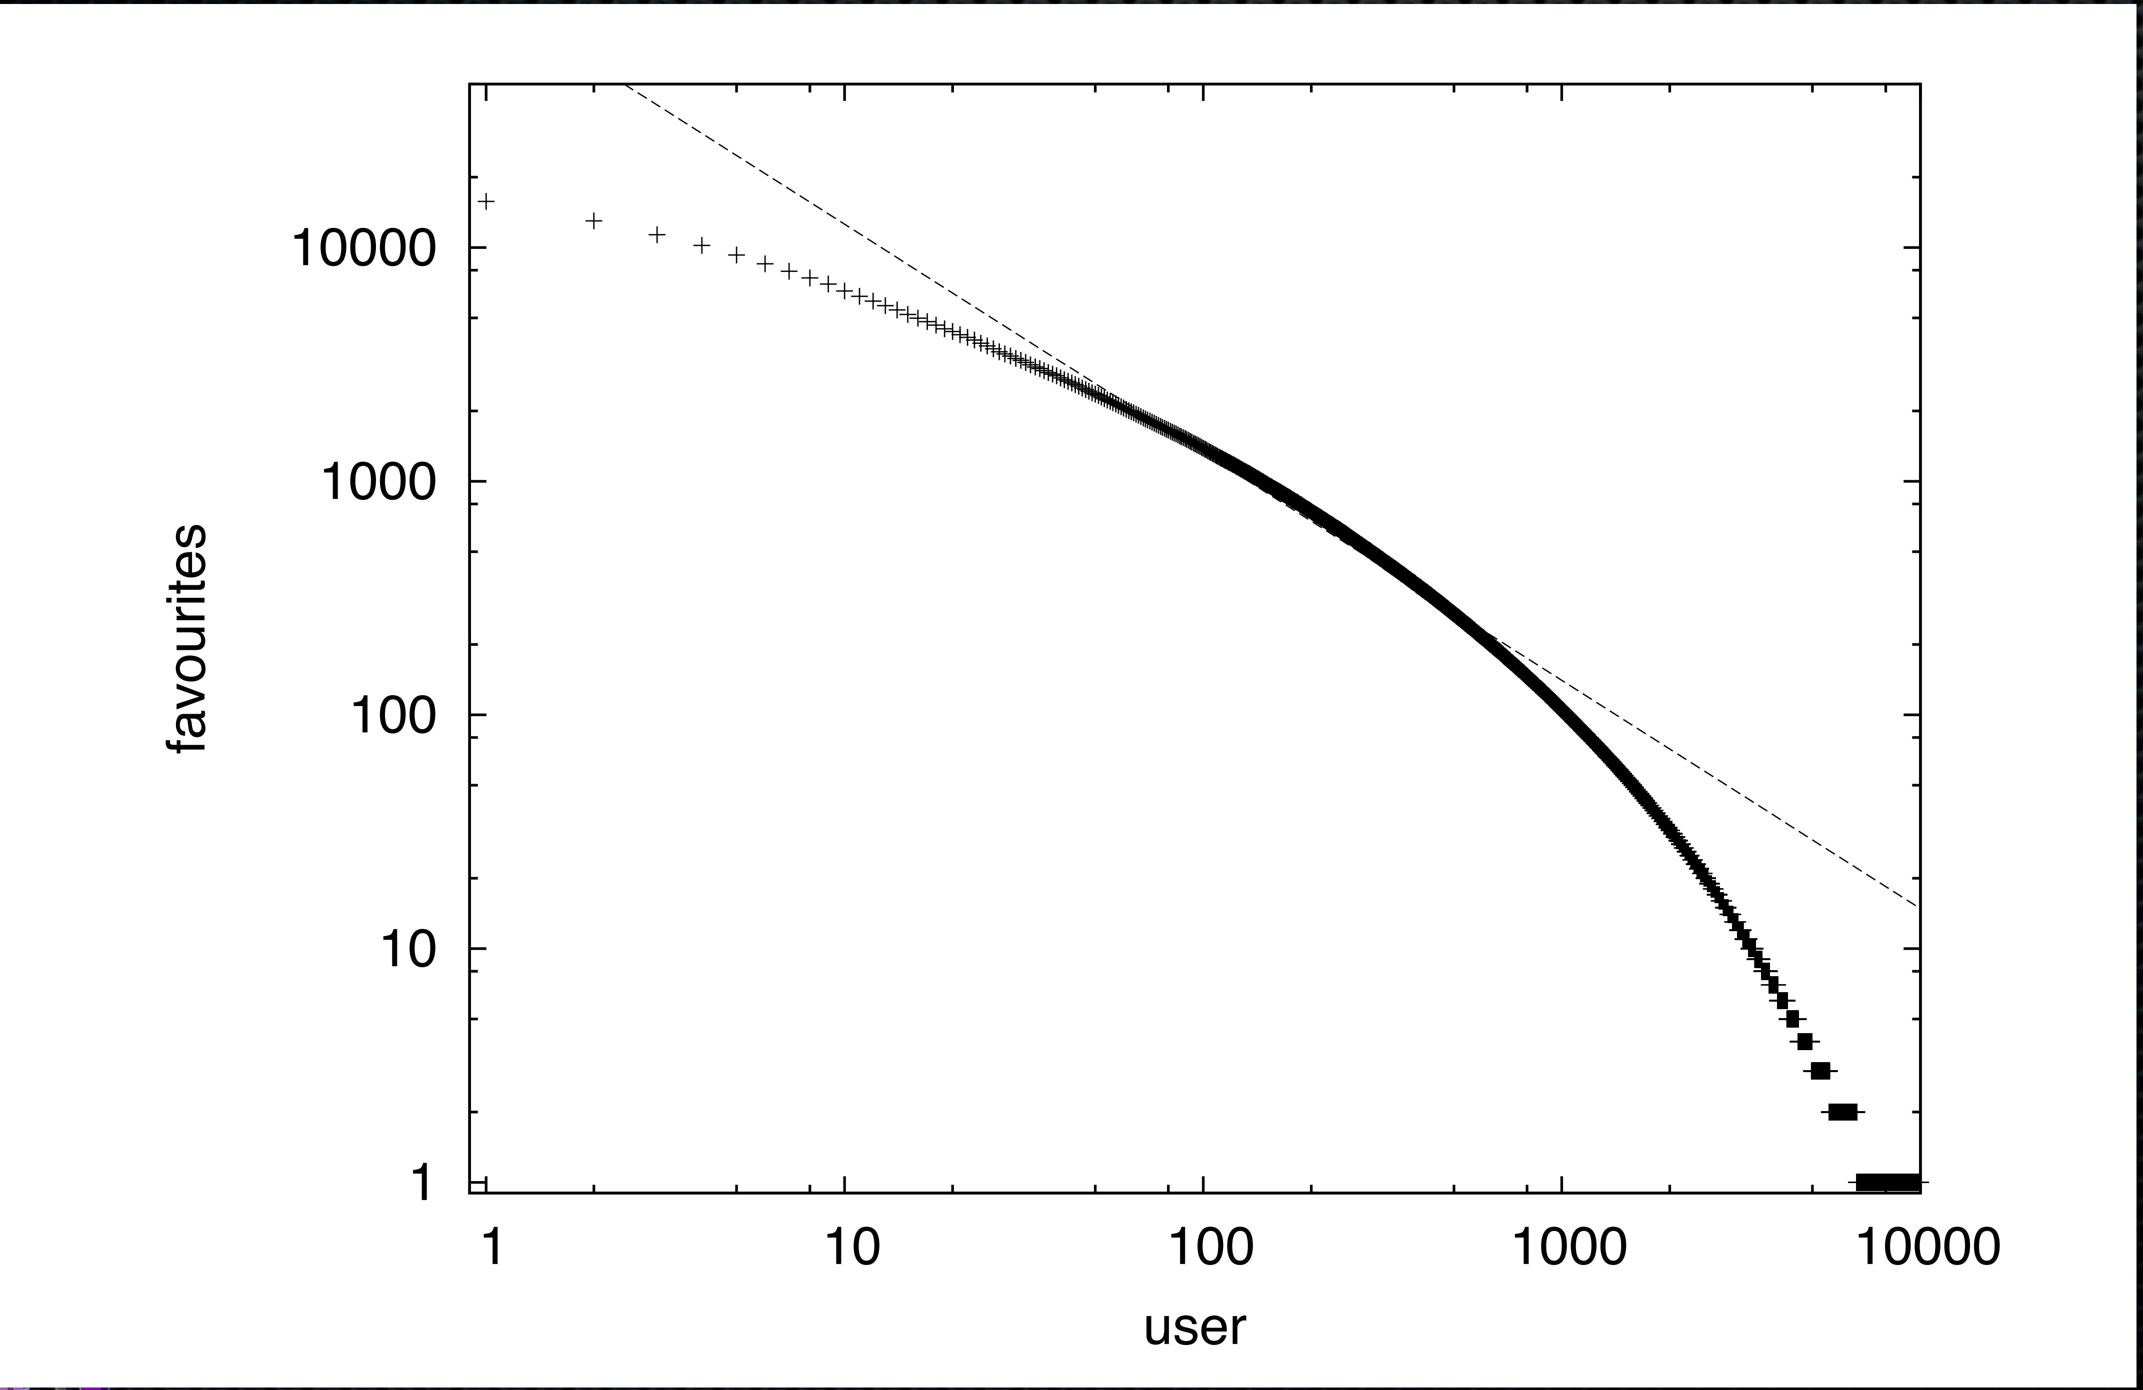

Features

- Types of Signals:
  - Textual (tags, title, description)
  - Visual (contrast, saturation, sharpness, naturalness, brightness, etc.)
  - Social (of a friend, posted in group, groups shared, number of comments)
  - Spatial (geo tagged)
  - Temporal (reflecting change of interest)

Learning

- Approach:
  - 74 features defined for the different types of signals.
- Machine learned approach based on:
  - **Gradient-boosted decision trees**
    - **logistic regression**

Approach

- Approach:
  - Data collected from 100 active users on Flickr, having 100+ favourites:
  - 44,322 photos (positive judgements)
  - 288,139 photos (negative judgements)
  - Random Context: pool at random from Flickr
  - Social Random Context: pool at random from a user's contacts and groups.

Evaluation (random context)

System	Accuracy	Pos. prec.	Pos. rec.	Neg. prec.	Neg. rec.
Textual	0.92	0.84	0.49	0.93	0.99
Visual	0.88	0.78	0.15	0.88	0.99
Social	**0.97**	0.94	0.82	0.97	0.99
Textual+Visual	0.93	0.88	0.57	0.94	0.99
Textual+Social	**0.97**	**0.95**	0.83	0.97	0.99
Visual+Social	**0.98**	**0.95**	**0.91**	0.99	0.99
Text+Vis+Soc	**0.98**	**0.95**	**0.91**	0.99	0.99

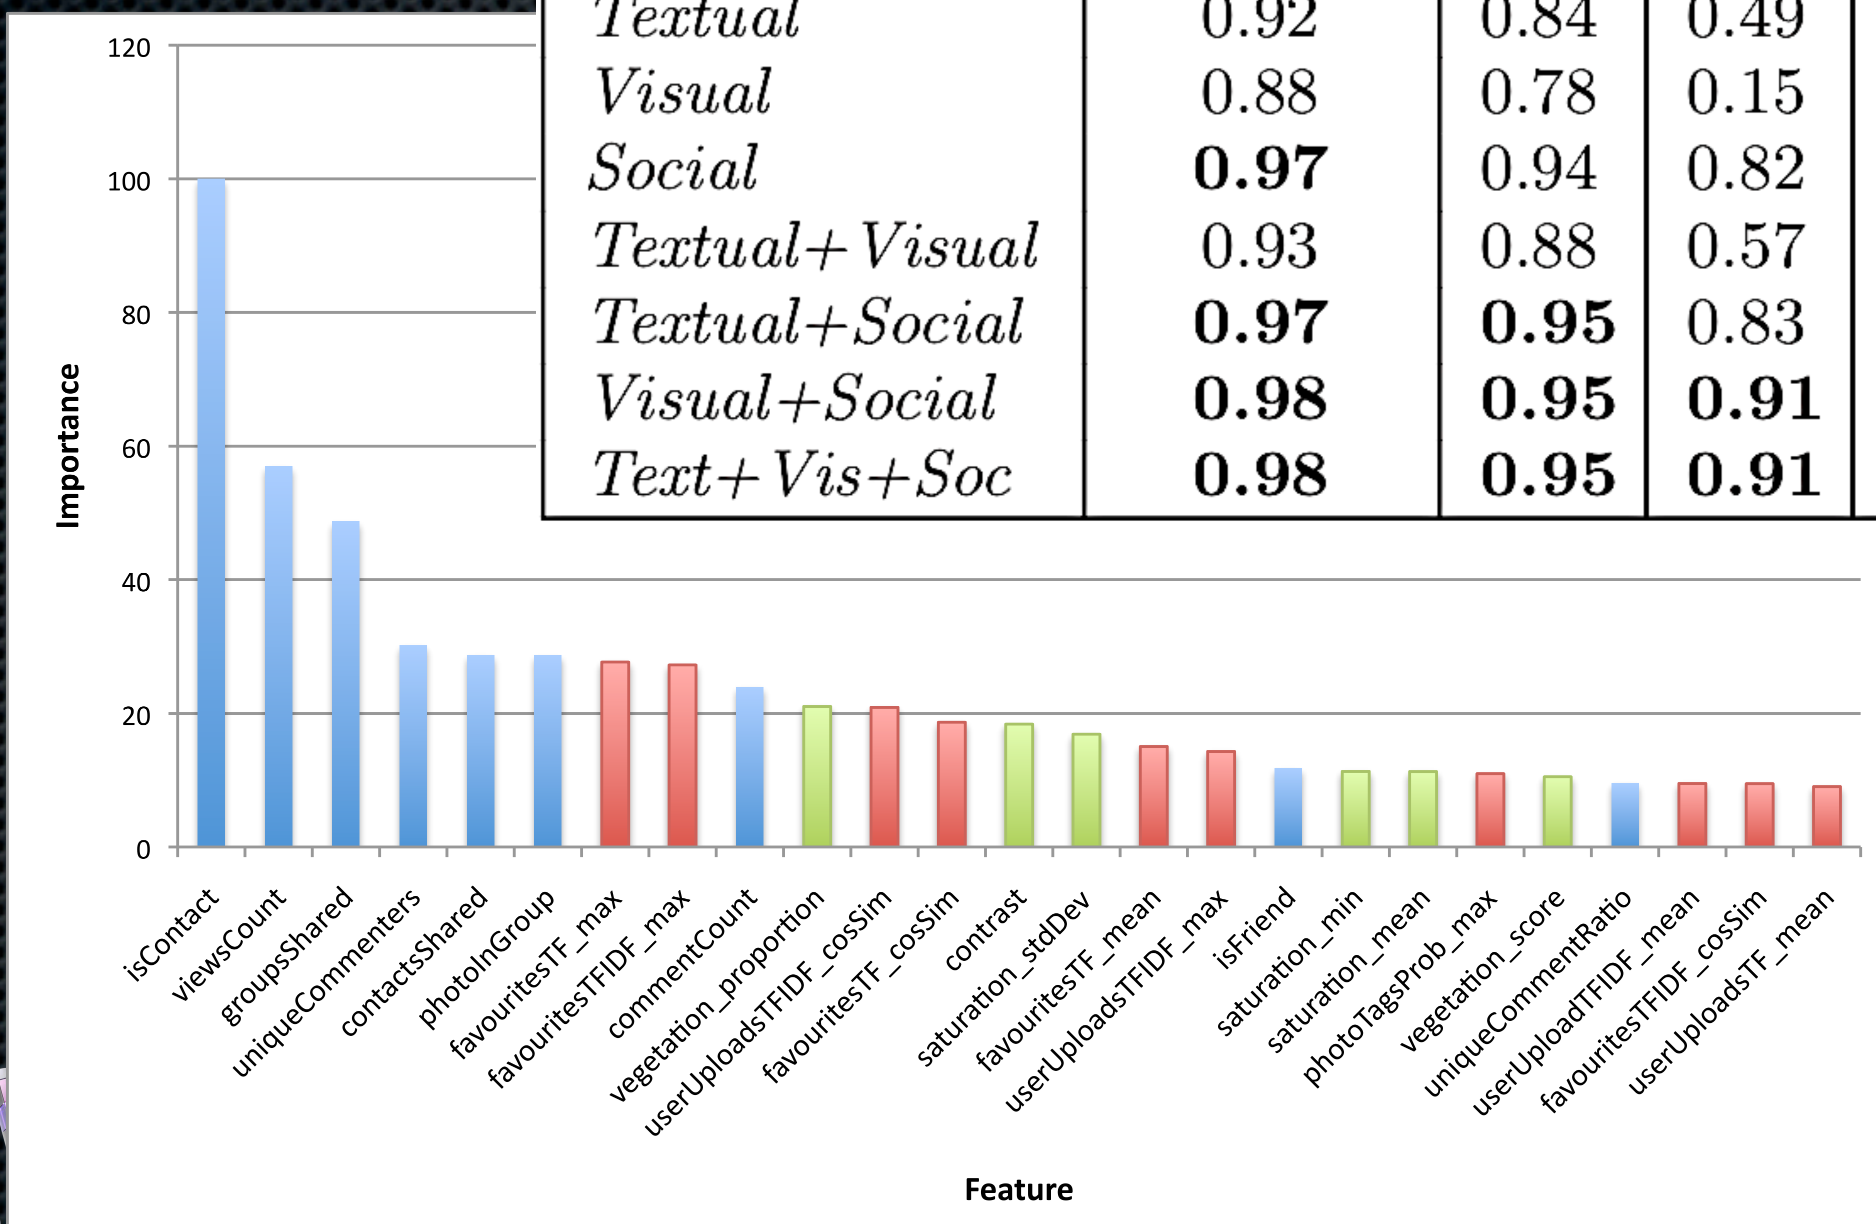

Evaluation (random context)

System	Accuracy	Pos. prec.	Pos. rec.	Neg. prec.	Neg. rec.
Textual	0.87	0.48	0.18	0.88	0.97
Visual	0.88	**1.0**	**0.085**	0.88	1.0
Social	0.92	0.80	0.56	0.94	0.98
Textual+Visual	0.88	0.62	0.27	0.90	0.97
Textual+Social	0.92	0.77	0.62	0.94	0.97
Visual+Social	**0.93**	**0.89**	**0.56**	0.94	0.99
Text+Vis+Soc	**0.93**	**0.84**	**0.62**	0.94	0.98

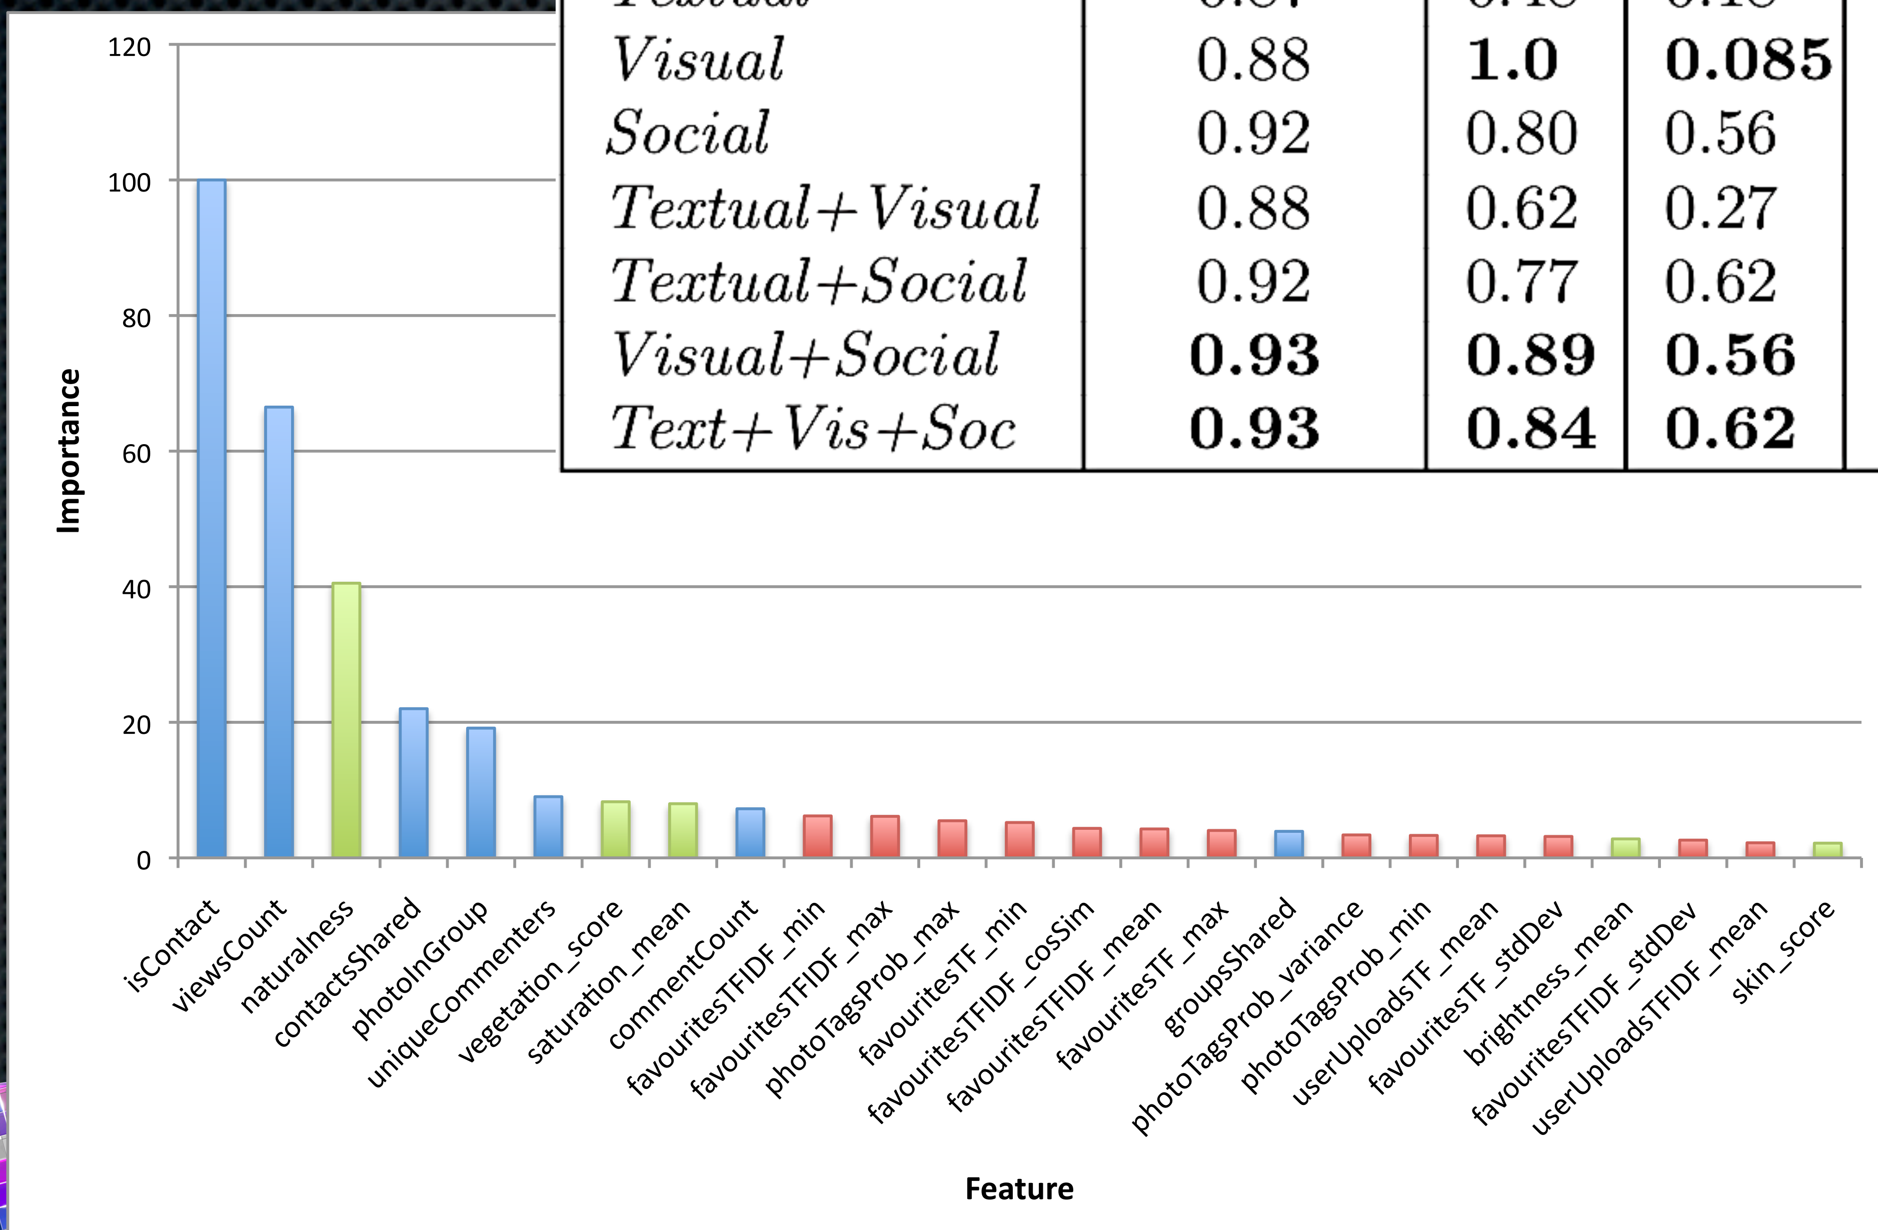

Pesonalization in Social Media

By: Eli Pariser

Eli Parisher is the former executive director of MoveOn.org, and the organization's current board president.

Who of you in the audience is:

(1) Working on personalization? (2) Prefers their results to

Automated Slideshows

A Classification Based Framework For Concept Summarization. *Dhruv Mahajan, Sundararajan Sellamanickam, Subhajit Sanyal, Amit Madaan.*

In WWW, 2010.

Slideshow generation

- Input:
  - Professionally produced content from Getty, AP, Reuters and about 500 other publishers
  - High quality, high resolution
  - Well annotated:title, caption, tags, entites, etc (NewsML)
  - On many topics we have 1000+ images

YAHOO!

Objective

- To produce an effective summarization of a collec- tion, which is characterized by ***some important properties*** that the summary should possess

- Dimensions: visual and semantic features

What Properties?

1. **Diversity** - Two images in the summary should not be similar to each other visually or semantically.
2. **Coverage** - The summary should cover all interesting and important visual and semantic aspects. Visual and semantic aspects with high likelihood should be present in the summary.
3. **Balance** - The various vi- sual and semantic aspects should be present in a balanced way to avoid any misunderstanding of summary.

Practical issues

- How to get data? What to **show**?
- How to measure **succes**?
- Requirement specific slideshows?
  - ***Regional bias***: Indian fans could be shown more news related to team India.
  - ***Temporal bias***: Recent images could be given preference.

Cricket example

- 1000+ images from Y! News, images are time stamped

San Francisco example

- 40.000 images from Flickr, time stamped, and geo-tagged

Approach (1)

- Given a collection of images for a single entity of interest:
    - Discover topics in the collection, based on LDA over the textual meta-data.
    - Each image is probabilistically assigned to K topics.
    - Different topics represent different aspects of the query/concept

Topic Discovery

howard
council
icc
president
john
minister

Topic - 0

cardiff
wales
watson
tim
shane
broad

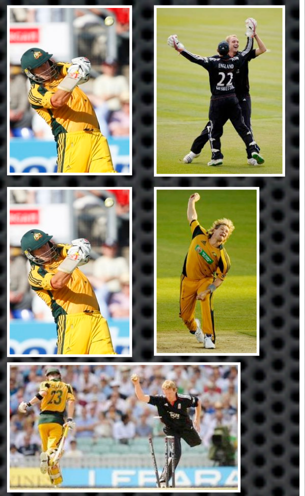

Topic - 1

africa's
barbados
Bridgetown
indies'
indies
africa

Topic - 4

pakistan
hussey
umar
series
akmal
runs

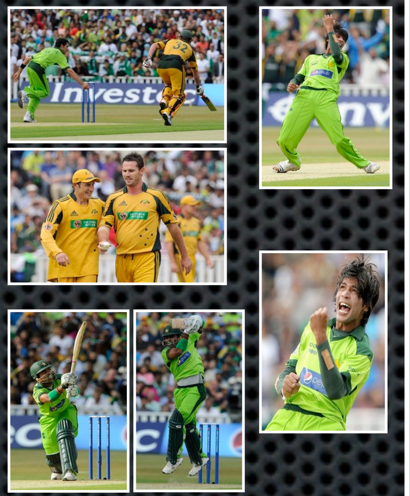

Topic - 5

indian
league
bcci
board
murali
premier

Topic - 8

finals
lanka's
tournament
lankan
final
dhoni

Topic - 9

Editors get to see only unorganized collection of images

- They might miss certain aspects

Selecting representative images

- Criteria
  - **Likelihood**: The representative images should be similar to many other images in the dataset.
  - **Diversity**: Any two representative images should not be similar to each other.
  - **Coverage**: The representative images should capture most of the dataset.
  - **Sparsity**: The set of representative images should be as sparse as possible

Sparse Models

- Kernel Logistic Regression (KLR)?
- Replace Hinge loss by logistic loss function.
- Gives natural estimate of data probability.
- Can handle multi–category classification naturally
- Model no longer sparse
-

$$\min_{\mathbf{a}} \frac{1}{N} \sum_{i=1}^{N} \log(1 + e^{-y_i f(\mathbf{x_i})}) + \frac{\lambda}{2} \mathbf{a}^T \mathbf{K} \mathbf{a}$$

$$f(\mathbf{x}) = \sum_{i=1}^{N} a_i K(\mathbf{x}, \mathbf{x_i}).$$

YAHOO!

Sparse Models

- Kernel Logistic Regression (KLR)?
- Replace Hinge loss by logistic loss function.
- Gives natural estimate of data probability.
- Can handle multi–category classification na

$$p(\mathbf{x}) = \frac{e^{f(\mathbf{x})}}{1+e^{f(\mathbf{x})}}$$

- Model no longer sparse
-

$$\min_{\mathbf{a}} \frac{1}{N}\sum_{i=1}^{N} \log(1+e^{-y_i f(\mathbf{x_i})}) + \frac{\lambda}{2}\mathbf{a}^T\mathbf{K}\mathbf{a}$$

$$f(\mathbf{x}) = \sum_{i=1}^{N} a_i K(\mathbf{x}, \mathbf{x_i}).$$

Import Vector Machines (IVM)

Full Model: $$f(\mathbf{x}) = \sum_{i=1}^{N} a_i K(\mathbf{x}, \mathbf{x_i}).$$

Import Vector Machines (IVM)

All Points

Full Model: $$f(\mathbf{x}) = \boxed{\sum_{i=1}^{N}} a_i K(\mathbf{x}, \mathbf{x_i}).$$

Import Vector Machines (IVM)

All Points

Full Model: $f(\mathbf{x}) = \sum_{i=1}^{N} a_i K(\mathbf{x}, \mathbf{x_i}).$

Sparse Model: $f(\mathbf{x}) = \sum_{\mathbf{x_i} \in \mathbf{S}} a_i K(\mathbf{x}, \mathbf{x_i}).$

Import Vector Machines (IVM)

All Points

Full Model: $$f(\mathbf{x}) = \boxed{\sum_{i=1}^{N}} a_i K(\mathbf{x}, \mathbf{x_i}).$$

Sparse Model: $$f(\mathbf{x}) = \boxed{\sum_{\mathbf{x_i} \in \mathbf{S}}} a_i K(\mathbf{x}, \mathbf{x_i}).$$

Import Vector Machines (IVM)

All Points

Full Model: $f(\mathbf{x}) = \boxed{\sum_{i=1}^{N}} a_i K(\mathbf{x}, \mathbf{x_i}).$

Sparse Model: $f(\mathbf{x}) = \boxed{\sum_{\mathbf{x_i} \in \mathbf{S}}} a_i K(\mathbf{x}, \mathbf{x_i}).$

Small Subset
(Import Vectors)

Import Vector Machines (IVM)

All Points

Full Model: $$f(\mathbf{x}) = \boxed{\sum_{i=1}^{N}} a_i K(\mathbf{x}, \mathbf{x_i}).$$

Sparse Model: $$f(\mathbf{x}) = \boxed{\sum_{\mathbf{x_i} \in \mathbf{S}}} a_i K(\mathbf{x}, \mathbf{x_i}).$$

Small Subset
(Import Vectors)

Use Greedy algorithm to add elements to set S one by one.

Multi – Category IVM

$$p_j(\mathbf{x}) = \frac{e^{f_j(\mathbf{x})}}{\sum_{c=1}^{C} e^{f_c(\mathbf{x})}}, \quad f_c(\mathbf{x}) = \sum_{\mathbf{x_i} \in \mathbf{S}} a_{ic} K(\mathbf{x}, \mathbf{x_i})$$

Multi – Category IVM

$$p_j(\mathbf{x}) = \frac{e^{f_j(\mathbf{x})}}{\sum_{c=1}^{C} e^{f_c(\mathbf{x})}}, \quad f_c(\mathbf{x}) = \sum_{\mathbf{x_i} \in \mathbf{S}} a_{ic} K(\mathbf{x}, \mathbf{x_i})$$

Probability that a sample belongs to class j

Multi – Category IVM

$$p_j(\mathbf{x}) = \frac{e^{f_j(\mathbf{x})}}{\sum_{c=1}^{C} e^{f_c(\mathbf{x})}}, \quad f_c(\mathbf{x}) = \sum_{\mathbf{x_i} \in \mathbf{S}} a_{ic} K(\mathbf{x}, \mathbf{x_i})$$

Probability that a sample belongs to class j

$$\min_{\mathbf{S}} \quad -\frac{1}{N} \sum_{i=1}^{N} y_{i,c_i} \log(p_{c_i}(\mathbf{x_i})) + \frac{\lambda}{2} \sum_{c=1}^{C} \mathbf{a_{:,c}}^T \mathbf{K} \mathbf{a_{:,c}}.$$

Multi – Category IVM

$$p_j(\mathbf{x}) = \frac{e^{f_j(\mathbf{x})}}{\sum_{c=1}^{C} e^{f_c(\mathbf{x})}}, \quad f_c(\mathbf{x}) = \sum_{\mathbf{x_i} \in \mathbf{S}} a_{ic} K(\mathbf{x}, \mathbf{x_i})$$

Probability that a sample belongs to class j

$$\min_{\mathbf{S}} \; -\frac{1}{N}\sum_{i=1}^{N} y_{i,c_i} \log(p_{c_i}(\mathbf{x_i})) + \frac{\lambda}{2}\sum_{c=1}^{C} \mathbf{a_{:,c}}^T \mathbf{K} \mathbf{a_{:,c}}.$$

Negative log-likelihood

Multi – Category IVM

$$p_j(\mathbf{x}) = \frac{e^{f_j(\mathbf{x})}}{\sum_{c=1}^{C} e^{f_c(\mathbf{x})}}, \quad f_c(\mathbf{x}) = \sum_{\mathbf{x_i} \in \mathbf{S}} a_{ic} K(\mathbf{x}, \mathbf{x_i})$$

Probability that a sample belongs to class j

$$\min_{\mathbf{S}} \quad -\frac{1}{N} \sum_{i=1}^{N} y_{i,c_i} \log(p_{c_i}(\mathbf{x_i})) + \frac{\lambda}{2} \sum_{c=1}^{C} \mathbf{a}_{:,\mathbf{c}}{}^T \mathbf{K} \mathbf{a}_{:,\mathbf{c}}.$$

Negative log-likelihood

Probability distribution qi given by LDA is not binary.
More than one category can have non-zero values

Multi – Category IVM

$$p_j(\mathbf{x}) = \frac{e^{f_j(\mathbf{x})}}{\sum_{c=1}^{C} e^{f_c(\mathbf{x})}}, \quad f_c(\mathbf{x}) = \sum_{\mathbf{x_i} \in \mathbf{S}} a_{ic} K(\mathbf{x}, \mathbf{x_i})$$

Probability that a sample belongs to class j

Binary

$$\min_{\mathbf{S}} \; -\frac{1}{N} \sum_{i=1}^{N} \boxed{y_{i,c_i}} \log(p_{c_i}(\mathbf{x_i})) + \frac{\lambda}{2} \sum_{c=1}^{C} \mathbf{a}_{:,\mathbf{c}}{}^{T} \mathbf{K} \mathbf{a}_{:,\mathbf{c}}.$$

Negative log-likelihood

Probability distribution qi given by LDA is not binary.
More than one category can have non-zero values

Multi – Category IVM

$$p_j(\mathbf{x}) = \frac{e^{f_j(\mathbf{x})}}{\sum_{c=1}^{C} e^{f_c(\mathbf{x})}}, \quad f_c(\mathbf{x}) = \sum_{\mathbf{x_i} \in \mathbf{S}} a_{ic} K(\mathbf{x}, \mathbf{x_i})$$

Multi – Category IVM

$$p_j(\mathbf{x}) = \frac{e^{f_j(\mathbf{x})}}{\sum_{c=1}^{C} e^{f_c(\mathbf{x})}}, \quad f_c(\mathbf{x}) = \sum_{\mathbf{x_i} \in \mathbf{S}} a_{ic} K(\mathbf{x}, \mathbf{x_i})$$

Probability that a sample belongs to class j

Multi – Category IVM

$$p_j(\mathbf{x}) = \frac{e^{f_j(\mathbf{x})}}{\sum_{c=1}^{C} e^{f_c(\mathbf{x})}}, \quad f_c(\mathbf{x}) = \sum_{\mathbf{x_i} \in \mathbf{S}} a_{ic} K(\mathbf{x}, \mathbf{x_i})$$

Probability that a sample belongs to class j

$$\min_{\mathbf{S}} \quad -\frac{1}{N} \sum_{i=1}^{N} y_{i,c_i} \log(p_{c_i}(\mathbf{x_i})) + \frac{\lambda}{2} \sum_{c=1}^{C} \mathbf{a_{:,c}}^T \mathbf{K} \mathbf{a_{:,c}}.$$

Multi – Category IVM

$$p_j(\mathbf{x}) = \frac{e^{f_j(\mathbf{x})}}{\sum_{c=1}^{C} e^{f_c(\mathbf{x})}}, \quad f_c(\mathbf{x}) = \sum_{\mathbf{x_i} \in \mathbf{S}} a_{ic} K(\mathbf{x}, \mathbf{x_i})$$

Probability that a sample belongs to class j

$$\min_{\mathbf{S}} \quad -\frac{1}{N} \sum_{i=1}^{N} y_{i,c_i} \log(p_{c_i}(\mathbf{x_i})) + \frac{\lambda}{2} \sum_{c=1}^{C} \mathbf{a}_{:,\mathbf{c}}{}^T \mathbf{K} \mathbf{a}_{:,\mathbf{c}}.$$

Negative log-likelihood

Multi – Category IVM

$$p_j(\mathbf{x}) = \frac{e^{f_j(\mathbf{x})}}{\sum_{c=1}^{C} e^{f_c(\mathbf{x})}}, \quad f_c(\mathbf{x}) = \sum_{\mathbf{x_i} \in \mathbf{S}} a_{ic} K(\mathbf{x}, \mathbf{x_i})$$

Probability that a sample belongs to class j

Binary

$$\min_{\mathbf{S}} \quad -\frac{1}{N} \sum_{i=1}^{N} \boxed{y_{i,c_i}} \log(p_{c_i}(\mathbf{x_i})) + \frac{\lambda}{2} \sum_{c=1}^{C} \mathbf{a}_{:,\mathbf{c}}{}^{T} \mathbf{K} \mathbf{a}_{:,\mathbf{c}}.$$

Negative log-likelihood

YAHOO!

Multi – Category IVM

$$p_j(\mathbf{x}) = \frac{e^{f_j(\mathbf{x})}}{\sum_{c=1}^{C} e^{f_c(\mathbf{x})}}, \quad f_c(\mathbf{x}) = \sum_{\mathbf{x_i} \in \mathbf{S}} a_{ic} K(\mathbf{x}, \mathbf{x_i})$$

Probability that a sample belongs to class j

Binary

$$\min_{\mathbf{S}} \; -\frac{1}{N} \sum_{i=1}^{N} \boxed{y_{i,c_i}} \log(p_{c_i}(\mathbf{x_i})) + \frac{\lambda}{2} \sum_{c=1}^{C} \mathbf{a}_{:,\mathbf{c}}{}^T \mathbf{K} \mathbf{a}_{:,\mathbf{c}}.$$

REPLACE

Negative log-likelihood

YAHOO!

Multi – Category IVM

$$p_j(\mathbf{x}) = \frac{e^{f_j(\mathbf{x})}}{\sum_{c=1}^{C} e^{f_c(\mathbf{x})}}, \quad f_c(\mathbf{x}) = \sum_{\mathbf{x_i} \in \mathbf{S}} a_{ic} K(\mathbf{x}, \mathbf{x_i})$$

Probability that a sample belongs to class j

Binary

$$\min_{\mathbf{S}} \; -\frac{1}{N} \sum_{i=1}^{N} \boxed{y_{i,c_i}} \log(p_{c_i}(\mathbf{x_i})) + \frac{\lambda}{2} \sum_{c=1}^{C} \mathbf{a}_{:,\mathbf{c}}{}^T \mathbf{K} \mathbf{a}_{:,\mathbf{c}}.$$

REPLACE

Negative log-likelihood

$$\min_{\mathbf{S}} \; -\frac{1}{N} \sum_{i=1}^{N} \sum_{j=1}^{C} q_{ij} \log(p_j(\mathbf{x_i})) + \frac{\lambda}{2} \sum_{c=1}^{C} \mathbf{a}_{:,\mathbf{c}}{}^T \mathbf{K} \mathbf{a}_{:,\mathbf{c}}.$$

Topic probabilities given by LDA

Multi – Category IVM

$$p_j(\mathbf{x}) = \frac{e^{f_j(\mathbf{x})}}{\sum_{c=1}^{C} e^{f_c(\mathbf{x})}}, \quad f_c(\mathbf{x}) = \sum_{\mathbf{x_i} \in \mathbf{S}} a_{ic} K(\mathbf{x}, \mathbf{x_i})$$

Probability that a sample belongs to class j

Binary

$$\min_{\mathbf{S}} \; -\frac{1}{N} \sum_{i=1}^{N} \boxed{y_{i,c_i}} \log(p_{c_i}(\mathbf{x_i})) + \frac{\lambda}{2} \sum_{c=1}^{C} \mathbf{a}_{:,\mathbf{c}}{}^T \mathbf{K} \mathbf{a}_{:,\mathbf{c}}.$$

Negative log-likelihood

REPLACE

KL - divergence

$$\min_{\mathbf{S}} \; -\frac{1}{N} \sum_{i=1}^{N} \boxed{\sum_{j=1}^{C} q_{ij} \log(p_j(\mathbf{x_i}))} + \frac{\lambda}{2} \sum_{c=1}^{C} \mathbf{a}_{:,\mathbf{c}}{}^T \mathbf{K} \mathbf{a}_{:,\mathbf{c}}.$$

Topic probabilities given by LDA

- Run LDA on textual meta data of images to discover topics / classes and probability distribution qi over them.
- Run the optimization below to get a sparse set S of import vectors.

$$p_j(\mathbf{x}) = \frac{e^{f_j(\mathbf{x})}}{\sum_{c=1}^{C} e^{f_c(\mathbf{x})}}, \quad f_c(\mathbf{x}) = \sum_{\mathbf{x_i} \in \mathbf{S}} a_{ic} K(\mathbf{x}, \mathbf{x_i})$$

Probability that sample x belongs to class

Import Vectors

$$\min_{\mathbf{S}} \; -\frac{1}{N} \sum_{i=1}^{N} \sum_{j=1}^{C} q_{ij} \log(p_j(\mathbf{x_i})) + \frac{\lambda}{2} \sum_{c=1}^{C} \mathbf{a}_{:,c}{}^T \mathbf{K} \mathbf{a}_{:,c}.$$

Topic probabilities given by LDA

Toy Example

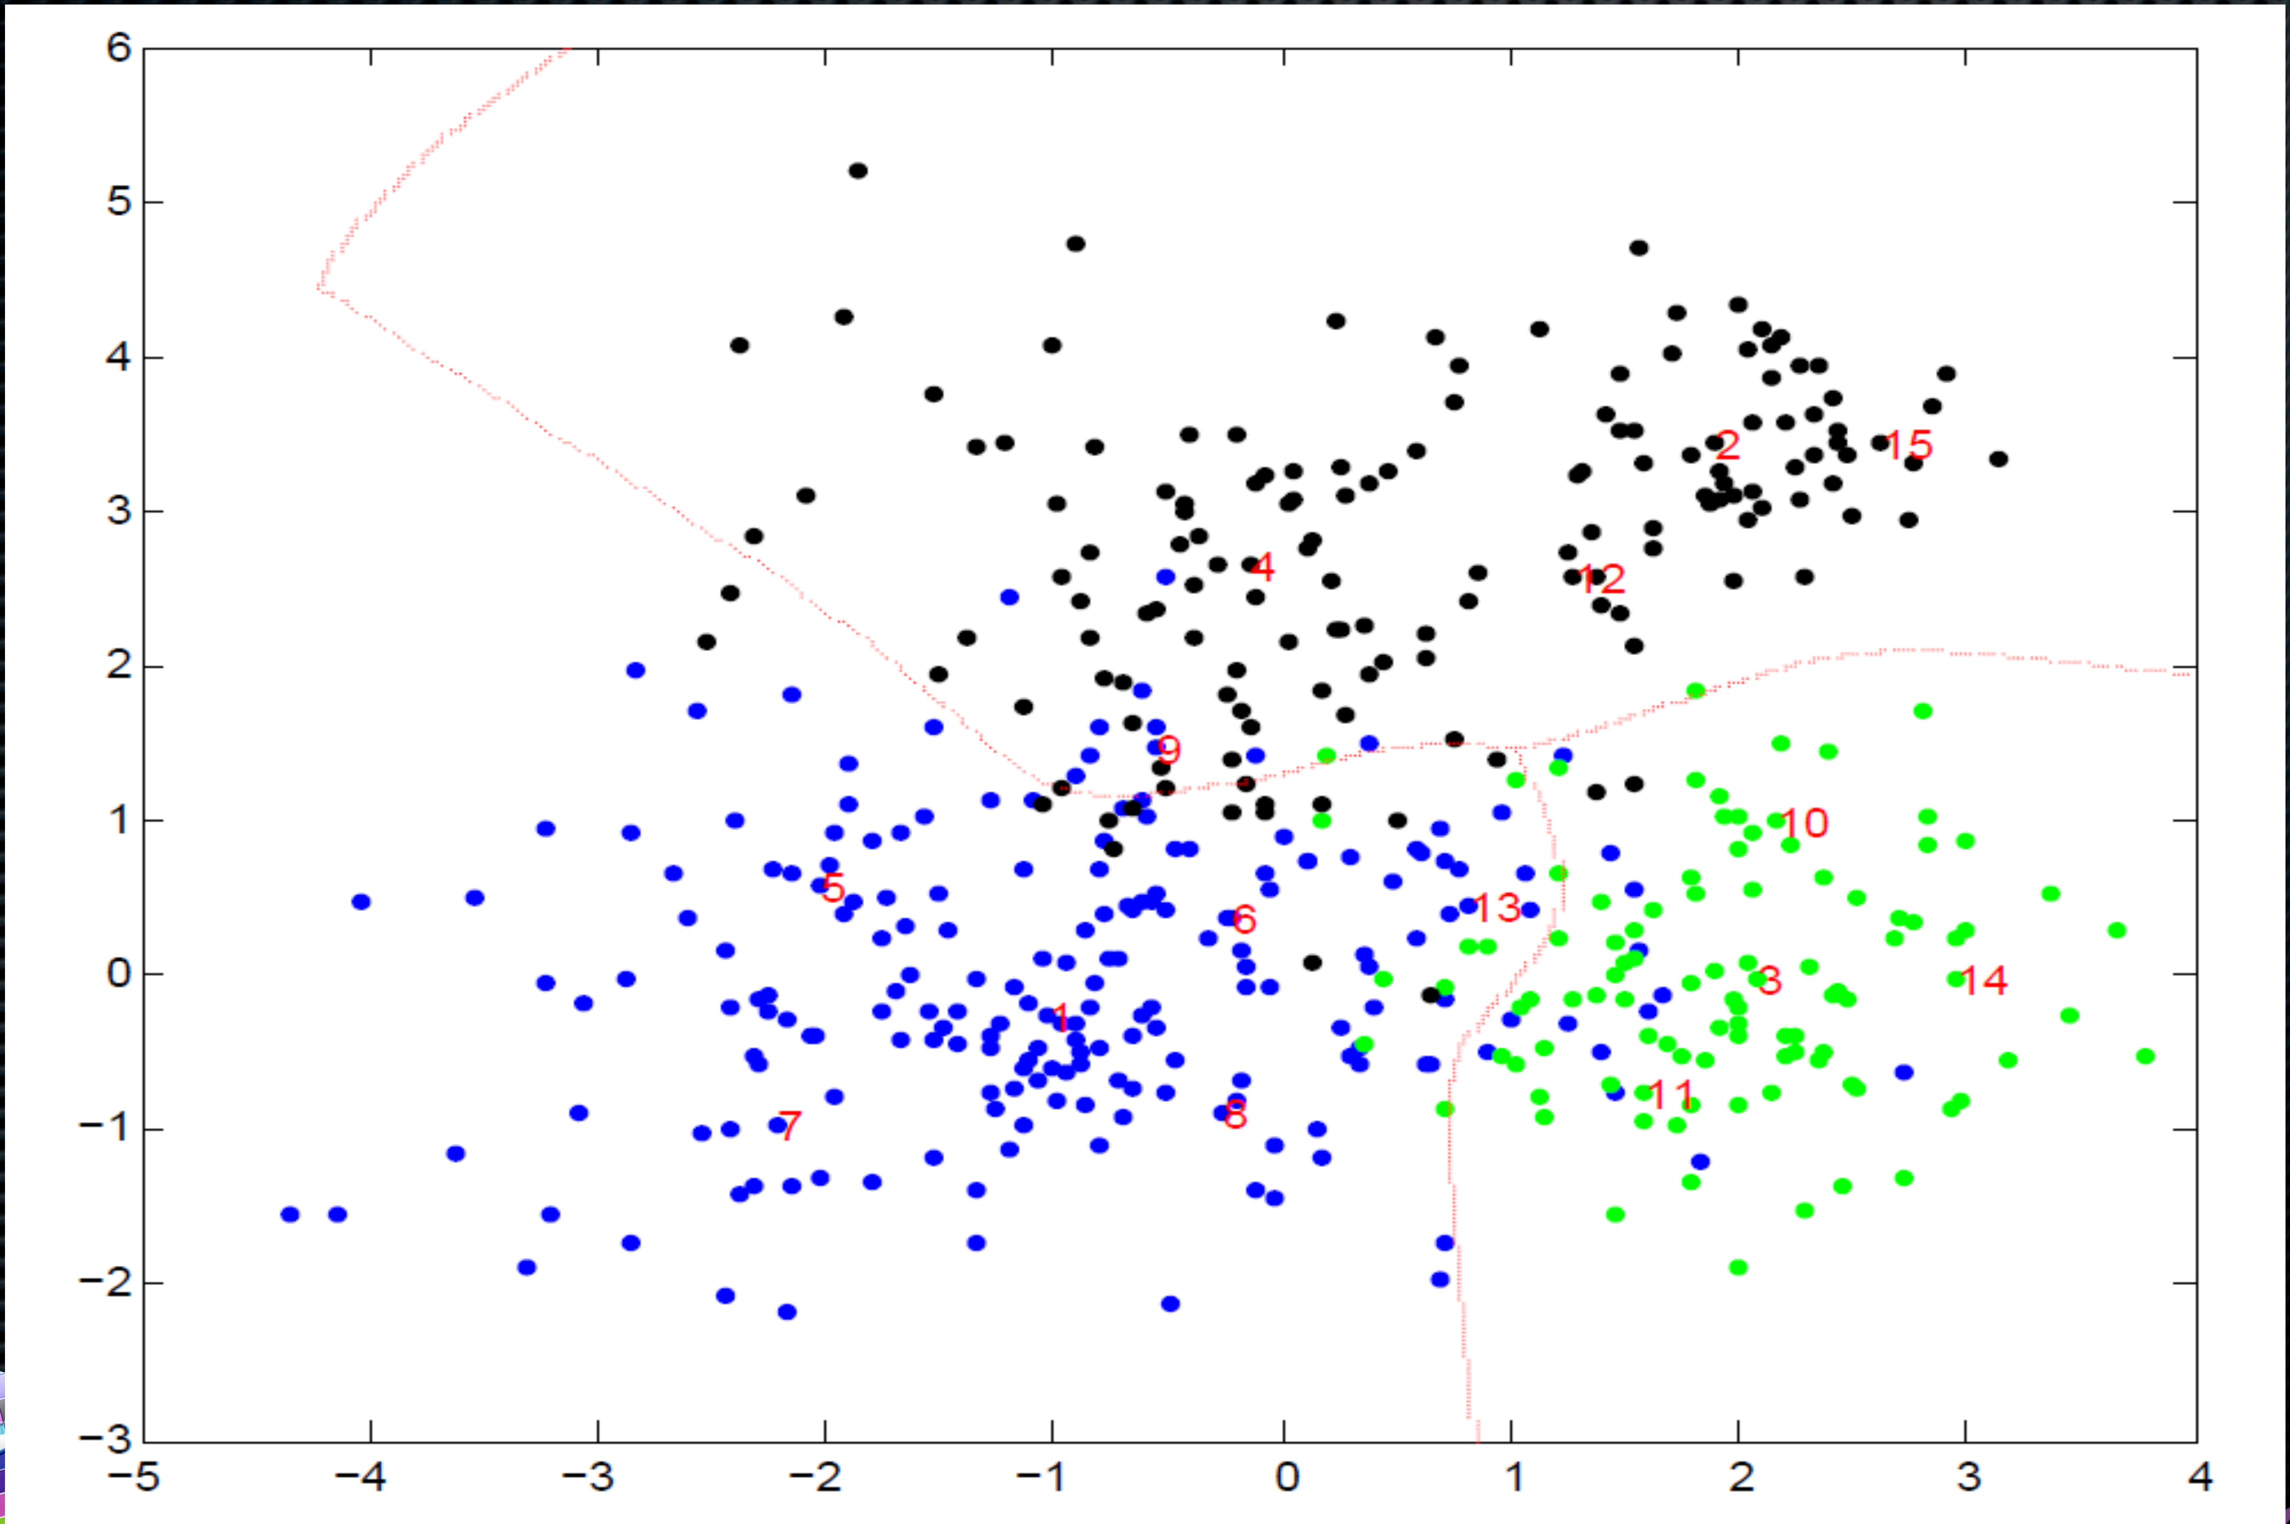

Cricket - Results

Topic 0

Topic 9

Topic 8

Topic 4

Topic 1, Topic 2, Topic 6, Topic 7

Topic 3, Topic 5

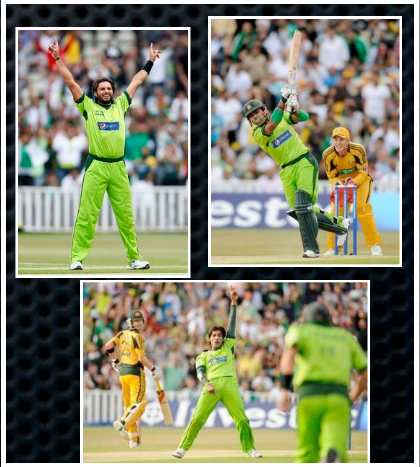

Distribution Constraints

- User might want some flexibility in terms of information needed.
  - **Temporal Bias:** We might want to look only at recent, older or nicely spanned images of Britney Spears over time.
  - **Category Bias**: User in Pakistan could be shown more images related to Pakistani Cricket team.
- Involve some kind of distribution over meta data
  - Based on **Time**, **Categories**, **Geo-locations**

Distribution Constraints

- We can specify desired distributions of representative images over such meta data in a rather simplistic way.
-

Distribution Constraints

- We can specify desired distributions of representative images over such meta data in a rather simplistic way.

  Pt(c) - Desired distribution over categories

-

Distribution Constraints

- We can specify desired distributions of representative images over such meta data in a rather simplistic way.

Pt(c) - Desired distribution over categories

Add extra KL-divergence term

$$\eta KL(P_t(c), \frac{1}{K} \sum_{\mathrm{x_i} \in S} \mathbf{q}_i)$$

-

YAHOO!

Distribution Constraints

- We can specify desired distributions of representative images over such meta data in a rather simplistic way.

Pt(c) - Desired distribution over categories

Add extra KL-divergence term

$$\eta KL(P_t(c), \tfrac{1}{K}\sum_{\mathbf{x}_i \in S} \mathbf{q}_i)$$

Distribution weight

-

Distribution Constraints

- We can specify desired distributions of representative images over such meta data in a rather simplistic way.

Pt(c) - Desired distribution over categories

Add extra KL-divergence term

$$\eta KL(P_t(c), \frac{1}{K} \sum_{\mathrm{x_i} \in S} \mathrm{q}_i)$$

Distribution weight

Number of Representative images

-

Distribution Constraints

- We can specify desired distributions of representative images over such meta data in a rather simplistic way.

Pt(c) - Desired distribution over categories

Add extra KL-divergence term

$$\eta KL(P_t(c), \frac{1}{K} \sum_{\mathrm{x_i} \in S} \mathrm{q}_i)$$

Distribution weight

Number of Representative images

Representative images set

-

Distribution Constraints

- We can specify desired distributions of representative images over such meta data in a rather simplistic way.

Pt(c) - Desired distribution over categories

Add extra KL-divergence term

$$\eta KL\left(P_t(c), \frac{1}{K} \sum_{\mathbf{x}_i \in S} \mathbf{q}_i\right)$$

Distribution weight

Number of Representative images

Representative images set

LDA probability vectors of representative images

-

Distribution Constraints

- We can specify desired distributions of representative images over such meta data in a rather simplistic way.

Pt(c) - Desired distribution over categories

Add extra KL-divergence term

$$\eta KL(P_t(c), \frac{1}{K} \sum_{\mathbf{x}_i \in S} \mathbf{q}_i)$$

Distribution weight

Number of Representative images

Representative images set

LDA probability vectors of representative images

-

Oil Spill - Uniform Distribution


Baby, immature and oil stained


Workers shovel oil off Fourchon Beach


Natural gas from damaged Deepwater Horizon


Greenpeace worker Lindsey Allen


Oil is seen on surface


BP Operations over Discoverer Enterprise drill


Workers clean up beach


Pelicans nest on island on coast


Oil is seen on island on coast

Salazar Ban Arctic Drilling


BP CEO Tony Hayward


Oil stained pelican is seen at its nest

Emary Billiot


Shrimp boats anchor


BP cleanup crew


BP CEO Tony Hayward

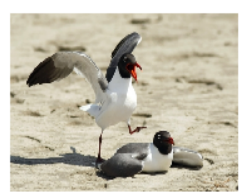
Laughing gull hops on laughing gull


Worker tries to clean up oil along beach

Ranzel Billiot

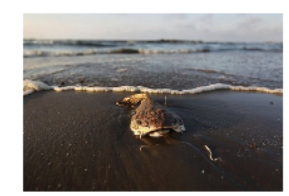
Dead oil-covered fish lies on beach

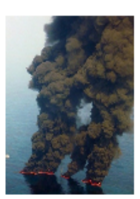
Controlled burns taking place


Nesting pelicans are seen on island


Janet Napolitano

Greenpeace activists

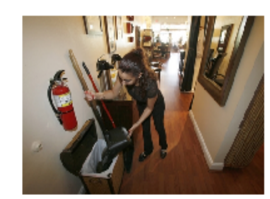
Ashley Perez

Oil Spill - Political view

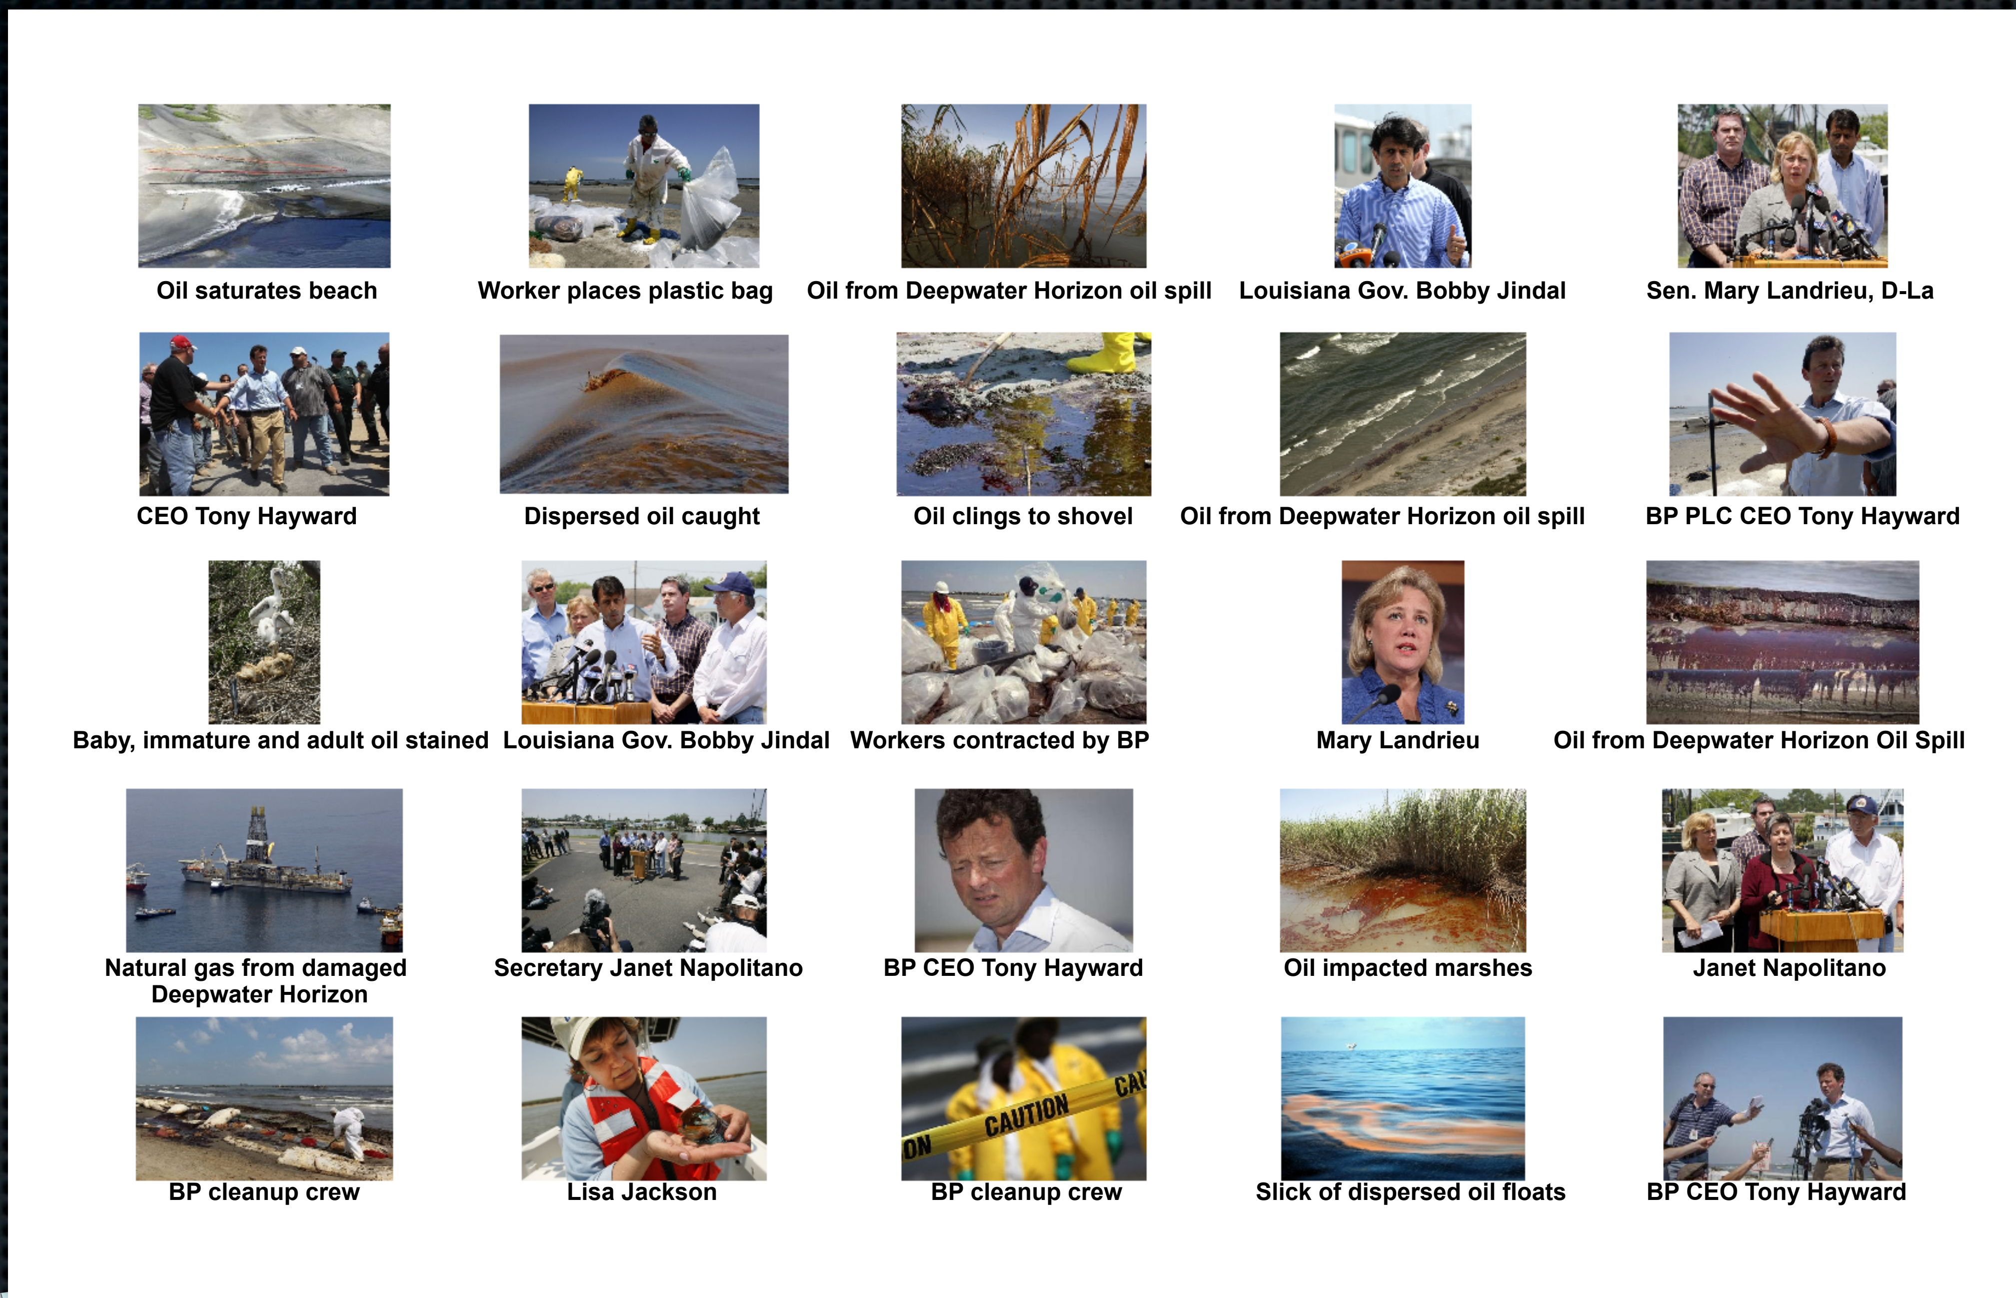

Oil Spill - Environmental view


Workers with Louisiana Department of Wildlife


Workers build a land bridge


Rescue workers move pelican


Workers clean up beach

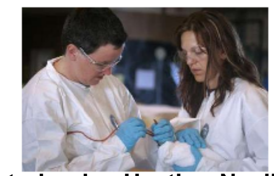
Veterinarian Heather Nevill


Workers contracted by BP


Nesting pelicans stand in oil


Boats motor through oily water


BP CEO Tony Hayward

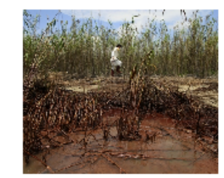
Greenpeace worker Lindsey

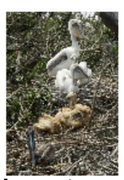
Baby, immature and adult oil stained


U.S. Fish and Wildlife officer Raul Sanchez


BP Operations over Discoverer Enterprise

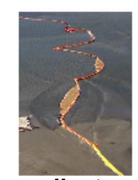
Oil booms collect weathered oil


Pelicans nest on island on coast


Laughing gull hops on laughing gull

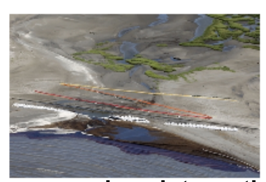
Oil is seen on beach to wetlands


Unidentified Coast Guard personnel

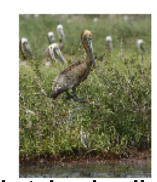
Oil stained pelican is seen at its nest


Worker tries to clean up oil along beach


Sheen is seen as Boat crews work


Nesting pelicans are seen landing on island


Unidentified Coast Guard personnel

Salazar Ban Arctic Drilling


Greenpeace Senior Campaigner Lindsey Allen

YAHOO!

Evaluation?

On facebook now:

Sponsored Slide

- How... ... ...
  - Present automated slideshow to editors for approval.
  - More thinking still to do.

Summary

The true challenge for multimedia is to find a balance between **relevancy**, **freshness**, **quality**, **interestingness** and **diversity** in order to provide an **engaging** rich **media experience** to the user.

Thank you, Roelof