Extracting Emergent Semantics from Large-Scale User-Generated Content

Yiannis Kompatsiaris
Informatics and Telematics Institute
Centre for Research and Technology - Hellas

http://mklab.iti.gr

Contents

Web 2.0 content
(July 2010)

flickr

- 3,190 uploads in the last minute
- 3.2 million things geotagged this month
- 4,754,012,299 photos (2 July 2010)

YouTube

- 24h of video content uploaded every minute
- 2 billion movies watched every day

facebook

- More than 400 million active users
- More than 200 million users log on at least once each day
- 2.5 billion photos uploaded each month

Social networks and media

- Users upload, tag, share, connect and search
- Emphasis is on uploading, visualization of results and interfaces
- Single media item analysis
- Limited usage of the **Collective nature** of Social Networks

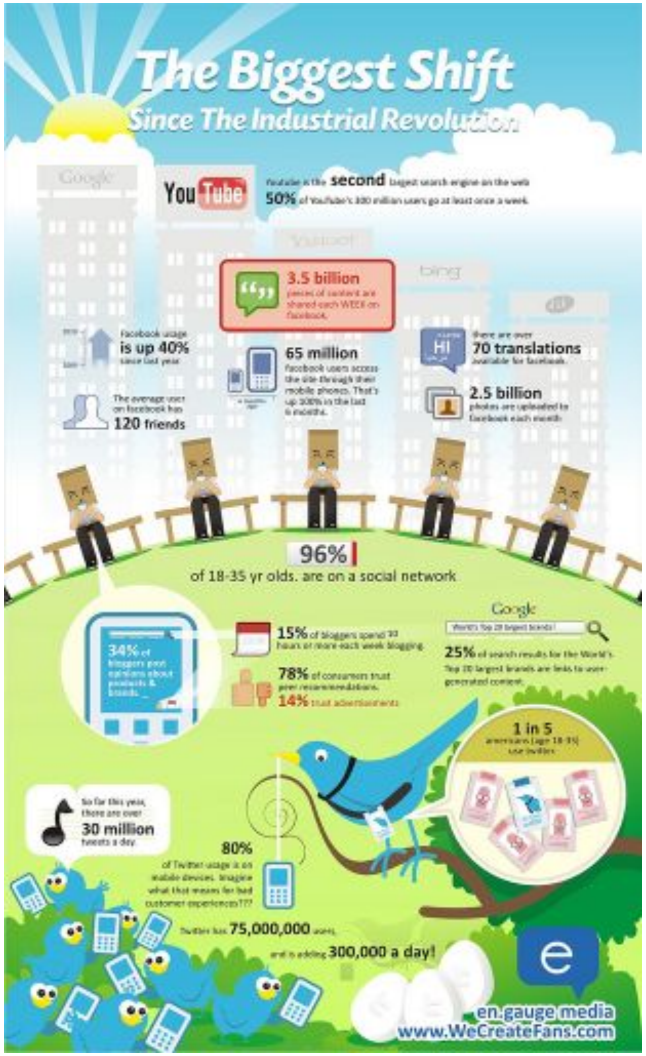


Low-level
Time
Geo
User Profile
Groups
Social network
Caption
Tags
Favs
Comms
flickr
Deutsches Eck from Ehrenbreitstein Fortress, Koblenz, Germany
by schaengel
121 comments
69 faves
Tagged with koblenz, ehrenbreitstein ...
Taken on November 15, 2009, uploaded November 17, 2009
See more of schaengel photos, or visit his profile.

Two main directions

- 1. Improve access to social media
  - Tag refinement, suggestion, propagation
- 2. Extract **implicit** information, capture emergent semantics
  - Not explicitly identifiable by users
  - Data mining
  - Collective Intelligence

Tags everywhere

Describe content and Search

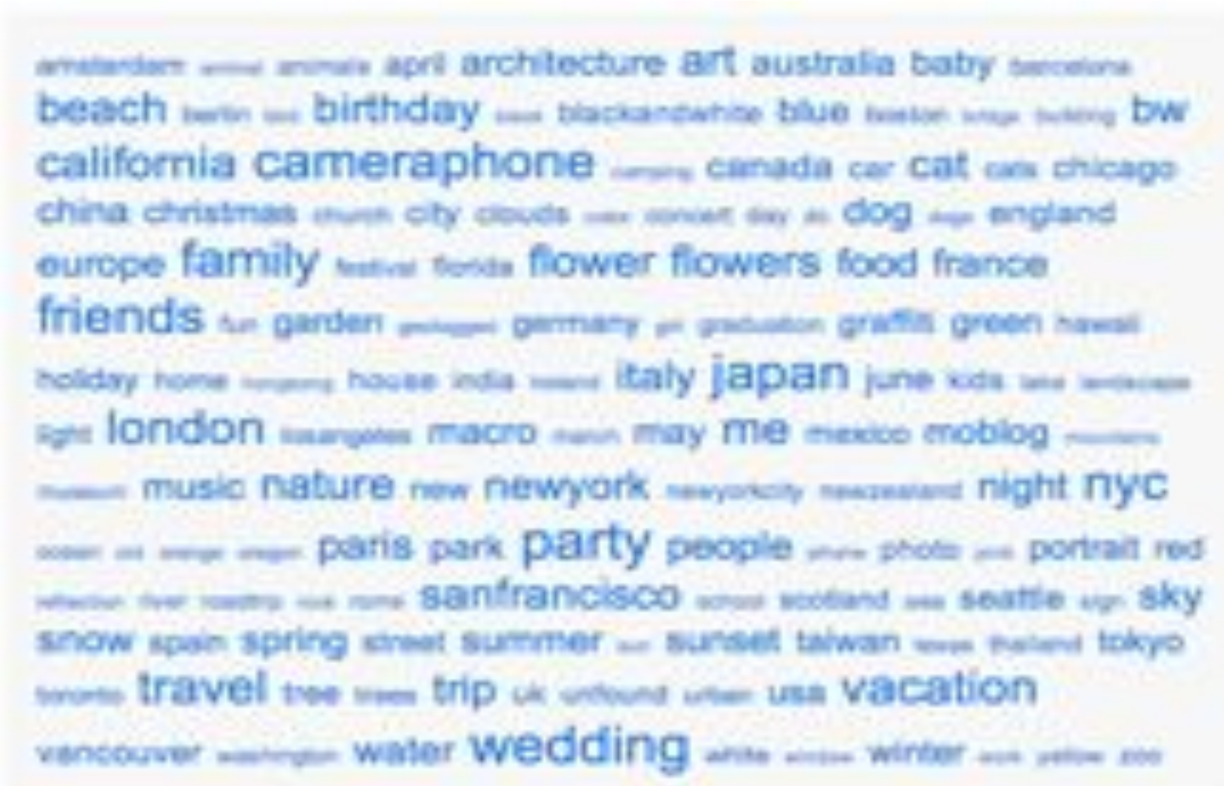

Very low precision

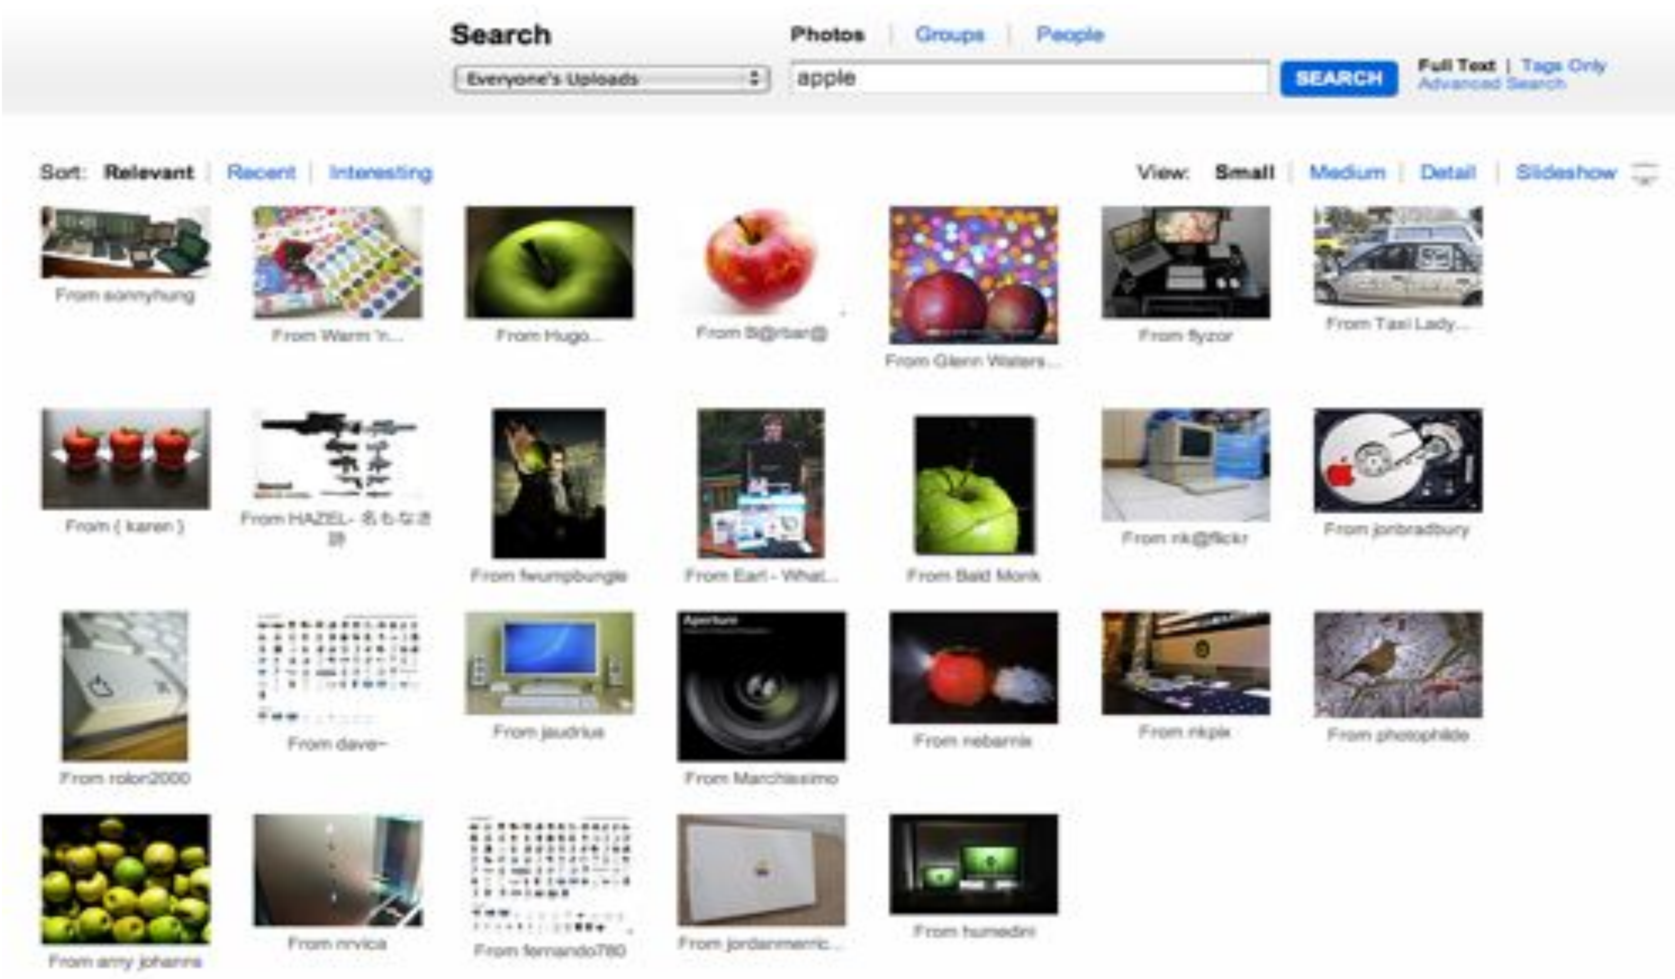

Very low recall

Tags

- Property#1
- Canada
- photo
- image
- digital
- urban
- Halifax
- park
- morning
- afternoon
- night
- Pentax K20D
- Sigma 70-300
- early
- Sackville

Can we improve things?

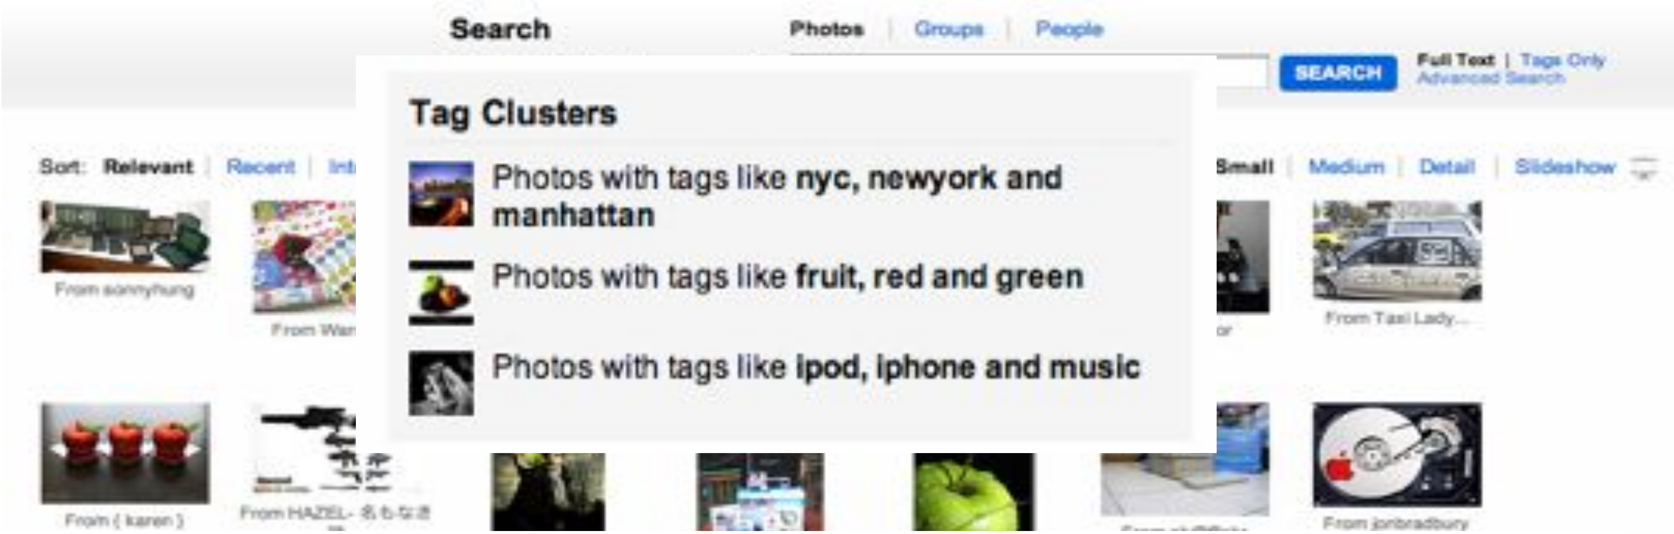

By combining information from many photos - tags, it seems that we can

Stable patterns

in tagging systems over time

Tag refinement, suggestion, ranking

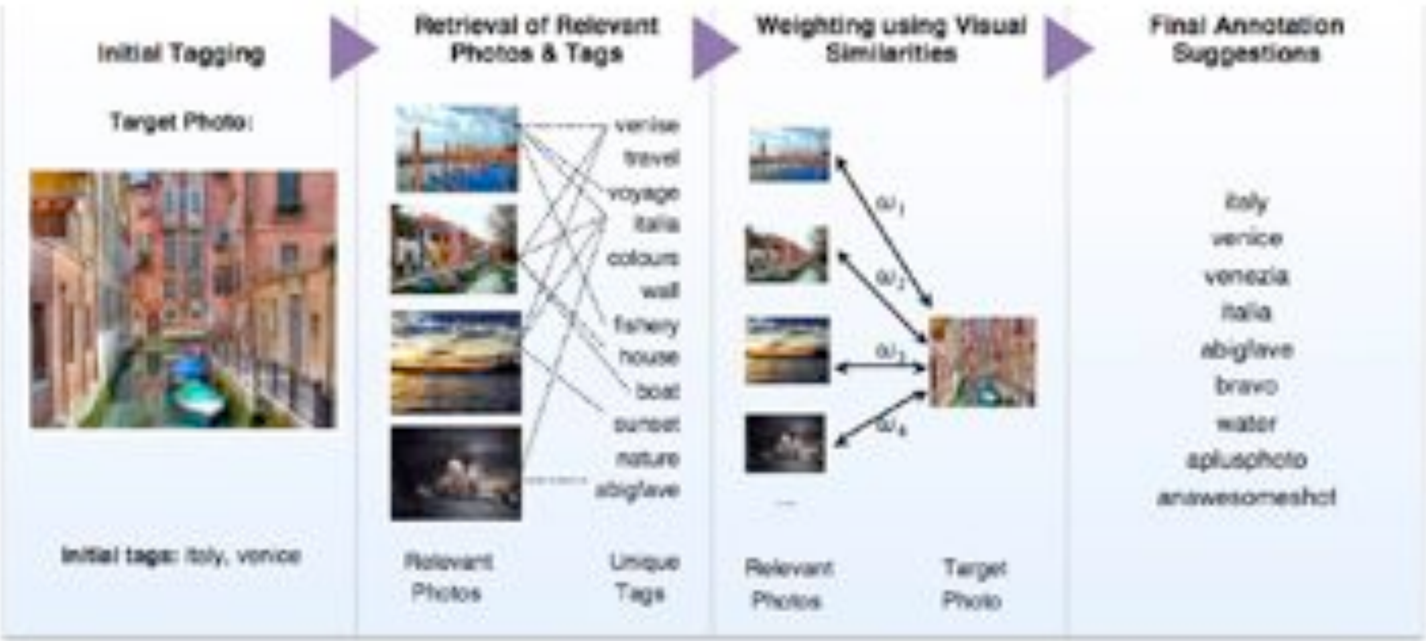

Social Networks and Collective Intelligence

- Social Networks is a data source with an extremely dynamic nature that reflects events and the evolution of community focus (user's interests)
- Potential for much more if we mine the data and their relations and exploit them in the right context
  - Scalable approaches taking into account the content and social context of social networks
- Search and Discovery of meaningful topics, entities, points of interest, social connections and events
- Rather than search for **isolated** or directly connected **social media items**

Extraction of implicit information

trace Flickr users from a chronologically ordered set of geographically referenced photos

Who are the Italians and who are the Americans?

MIT SENSEABLE CITY LAB, "The World's eyes"

What else we can do?

Tags that are “representative” for a geographical area

- 1. Clustering of photos
  - K-means, based on their location [Kennedy07]
- 2. Rank each cluster’s tags
- 3. Get tags above a certain threshold

Contribute to our understanding of the world

Representative tags for San Francisco [Kennedy07]

Sensors and automatically user generated content

Uses the GPS in cellular phones to gather traffic information, process it, and distribute it back to the phones in real time

- online, real-time data processing
- privacy-preservation
- data efficiency, i.e. not requiring excessive cellular network

Mobile Century Project: http://traffic.berkeley.edu/mobilecentury.html

Social Media as real-time Sensors

"…if you're more than 100 km away from the epicenter [of an earthquake] you can read about the quake on twitter before it hits you…"

Web 2.0 Content and Challenges

- **Multi-modality**: e.g. image + tags, image + video
- **Rich Social Context**: spatio-temporal, social connections, relations and social graph
- **Inconsistent quality**: noise, spam, ambiguity
- **Huge volume**: Massively produced and disseminated
- **Multi-source**: may be generated by different applications, user communities, e.g. delicious, StumbleUpon and reddit are all social bookmarking sites
  - Also connected to other sources (e.g. LOD, web)
- **Dynamic:** Fast updates, real-time

collective intelligence

…a form of intelligence emerging from online user activities

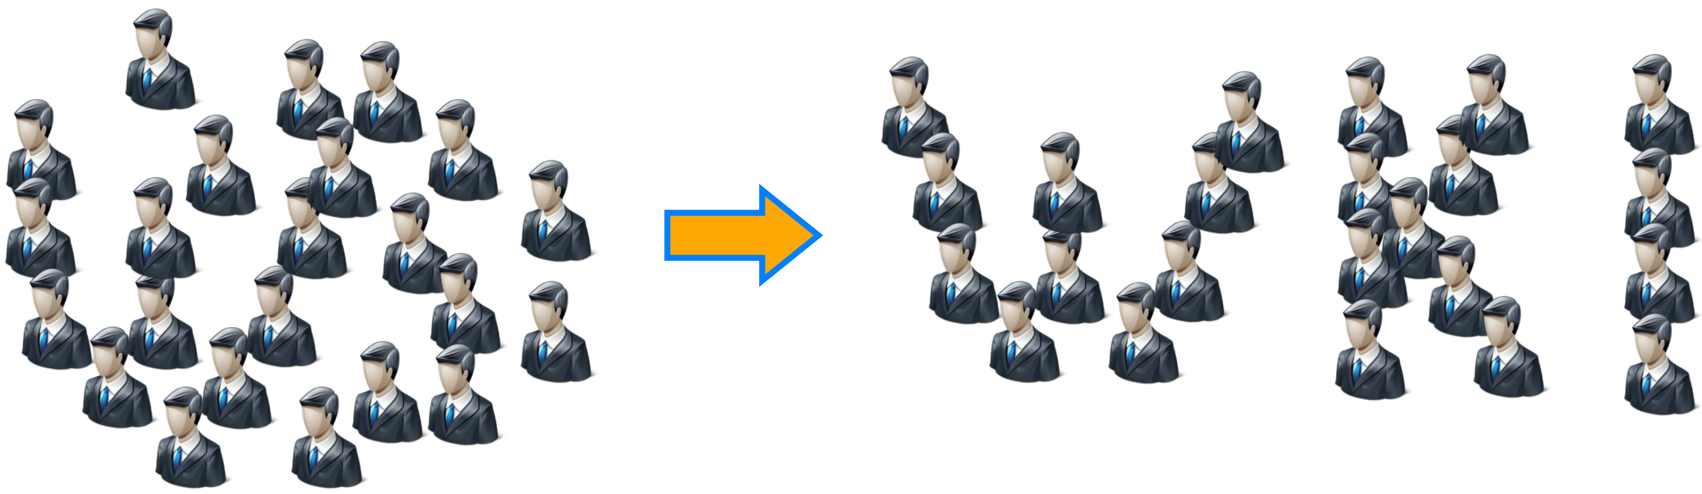

Collective Intelligence >> sum of individuals’ intelligences

Research Fields and Issues

- Statistical analysis, machine learning, data mining, pattern recognition, social network analysis
  - Clustering
- Representation, modeling, data reduction, graph theory
- Image, text, video analysis
- Information extraction
- Fusion techniques
- Stream processing and real-time architecture
- Trust, security, privacy
- Performance, scalability
  - speed, storage, power, grids, clouds

Applications

Xin Jin, Andrew Gallagher, Liangliang Cao, Jiebo Luo, and Jiawei Han. ***The wisdom of social multimedia: using flickr for prediction and forecast,*** International conference on Multimedia (MM '10). ACM.

Figure 7: Reuters/Zogby Poll v.s. Flickr. Y-axis denotes the percentage of popularity for candidate Edwards.

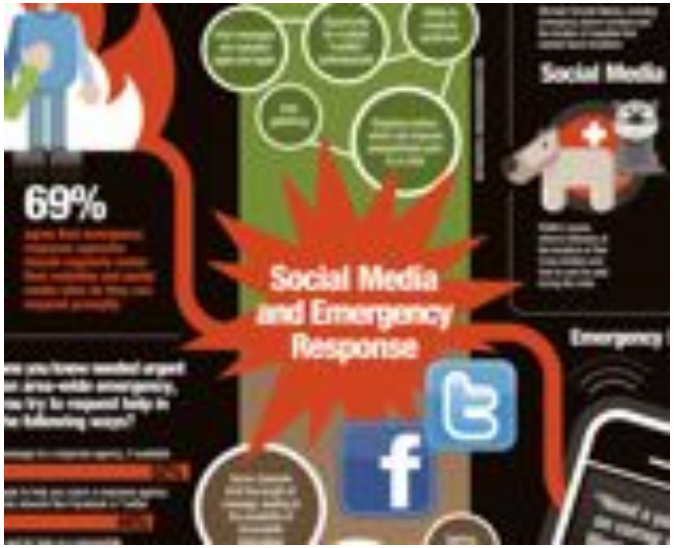

Federal Emergency Management Agency ***plans to engage the public*** more in disaster response by sharing data and leveraging reports ***from mobile phones and social media***

Gogobot: Travel Discovery Goes Social And Visual "The service raised $4 million in funding (Google CEO Eric Schmidt is one of the investors)…This is a $100 billion a year industry in the U.S. It's something like $350 billion worldwide."

Applications

- Science
  - Sociology, machine learning (machine as a teacher), computer vision (annotation)
- Tourism – Leisure – Culture
  - Off-the-beaten path POI extraction
- Marketing
  - Brand monitoring, personalised ads
- Prediction
  - Politics: election resulst
- News
  - Topics, trends event detection
- Others
  - Environment, emergency response, energy saving, etc

Social Media Community Detection

Examples of Social Media networks

Folksonomy (Delicious)

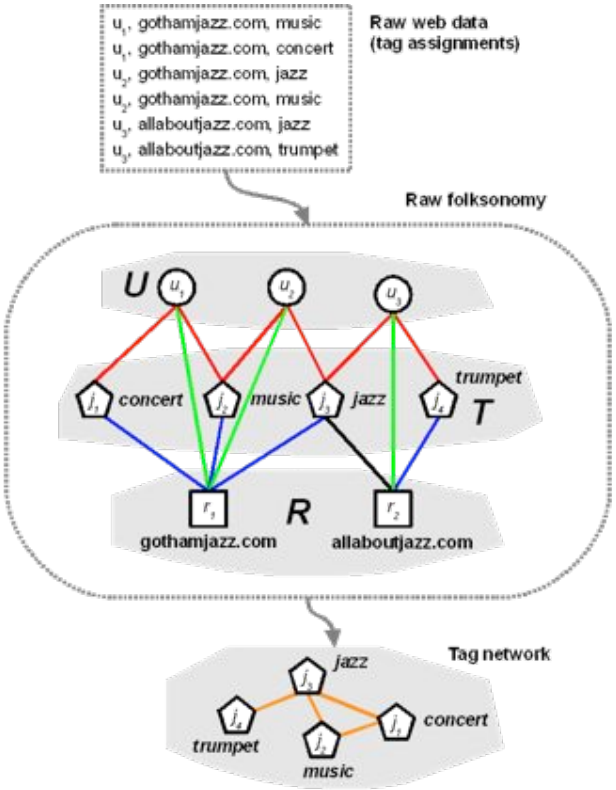

Mika, P. (2005) Ontologies Are Us: A Unified Model of Social Networks and Semantics. Proceedings of the 4th International Semantic Web Conference (ISWC 2005), Springer Berlin / Heidelberg, pp. 522-536

MetaGraph (Digg)

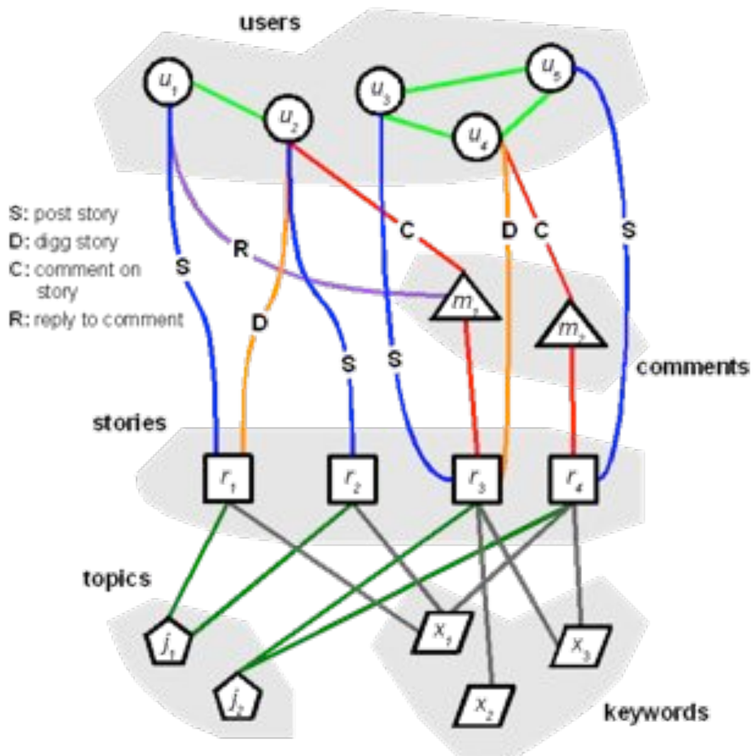

Lin, Y., Sun, J., Castro, P., Konuru, R., Sundaram, H., and Kelliher, A. (2009) MetaFac: community discovery via relational hypergraph factorization. Proceedings of KDD '09, ACM, pp. 527-536

What is a community in a network?

Group of vertices that are more densely connected to each other than to the rest of the network.

Multiple definitions to quantify communities:

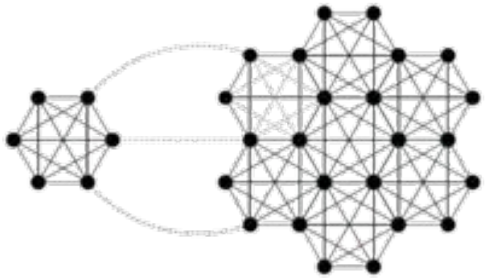

Fortunato S. (2010) Community detection in graphs. Physics Reports486: 75-174

- Global: N-cut, conductance, modularity
- Local: Local modularity, (μ,ε)-cores
- Ad hoc: Label propagation, dynamic synchronization

Related to clustering, but: (a) not necessary to know number of communities, (b) computationally more efficient

In Social Media, we focus on local definitions, because of the properties of Social Media networks: efficiency-scalability and noise resilience.

Challenges in Social Media network mining

No prior assumptions about structure:

- Complex & evolving structure
- No possibility for knowing structural features (e.g. number of clusters on a graph) in advance

→ Unsupervised

Scale

- Tens of millions of active users frequently contributing loads of content links + metadata (tags, comments, ratings)

→ Efficient - scalable

Quality

- Spam is very common. Only a portion of user contributions is worth further analysis.

→ Noise resilient

Approach illustration (1/2)

Two-step process:

- 1st step:

  (μ, ε) – core detection

- 2nd step:
  Local expansion

- 3rd step:
  Characterization of remaining vertices as *hubs* or *outliers*

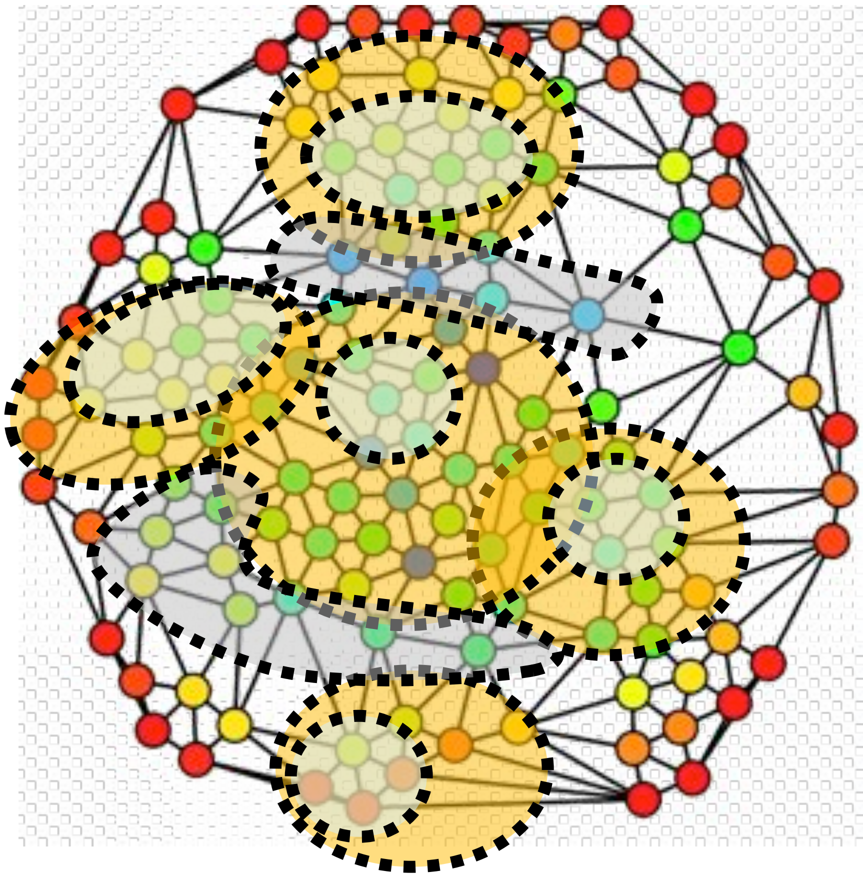

Approach illustration (2/2)

- Structural similarity + Local expansion
  (highly efficient and scalable approach)
- Not necessary to know the number of clusters
- Noise resilient
  (not all nodes need to be part of a community)
- Generic approach adaptable to many applications
  (depending on node – edge representation)

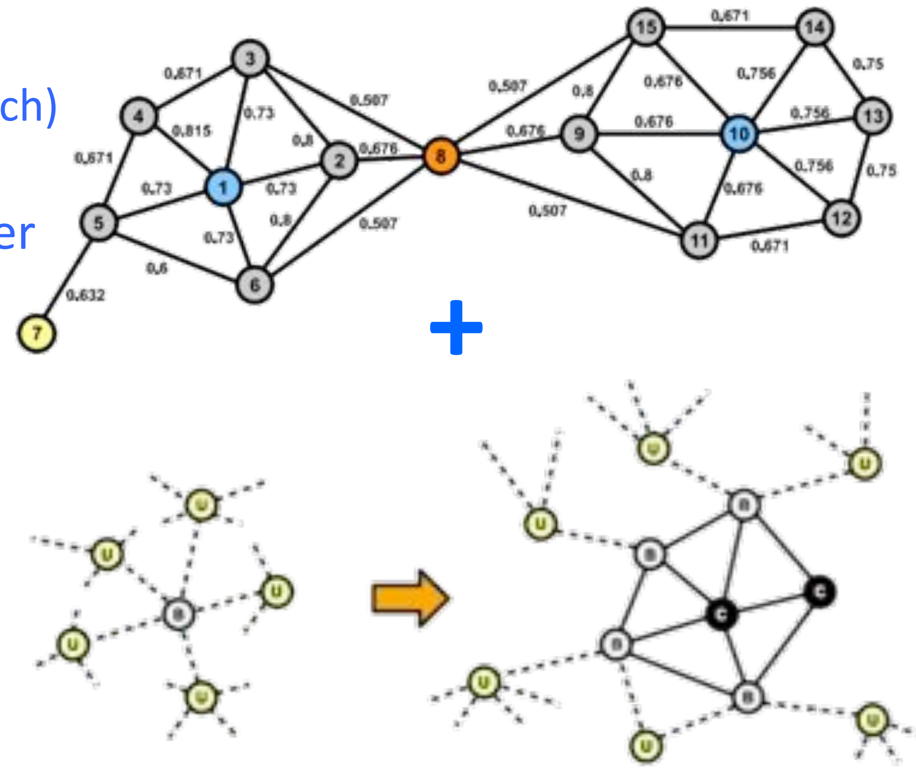

S. Papadopoulos, Y. Kompatsiaris, A. Vakali. "A Graph-based Clustering Scheme for Identifying Related Tags in Folksonomies". In Proceedings of DaWaK'10, Springer-Verlag, 65-76

LYCOS iQ Tag Network

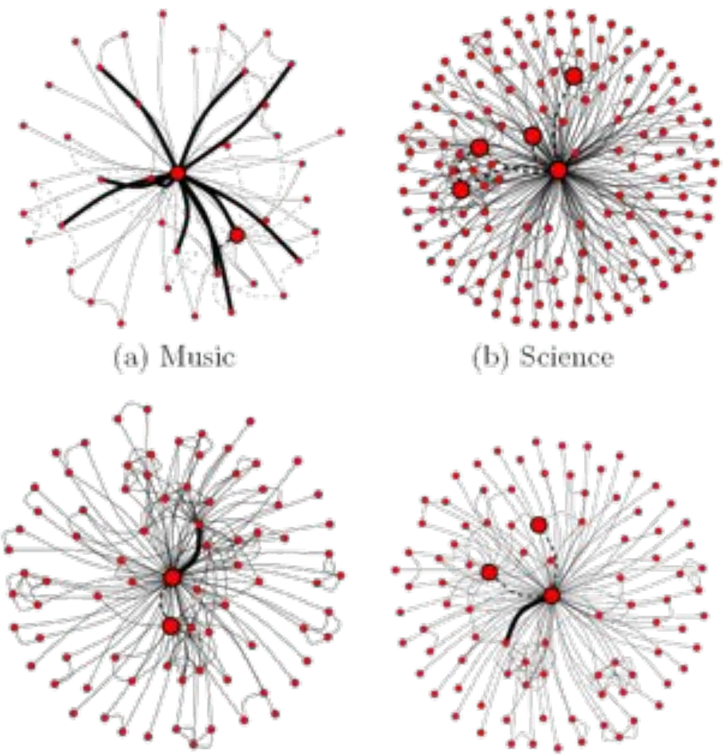

(a) Music (b) Science

(c) Film (d) Animals

Computers:
A densely interconnected community

History:
A star-shaped community

Hybrid photo Clustering

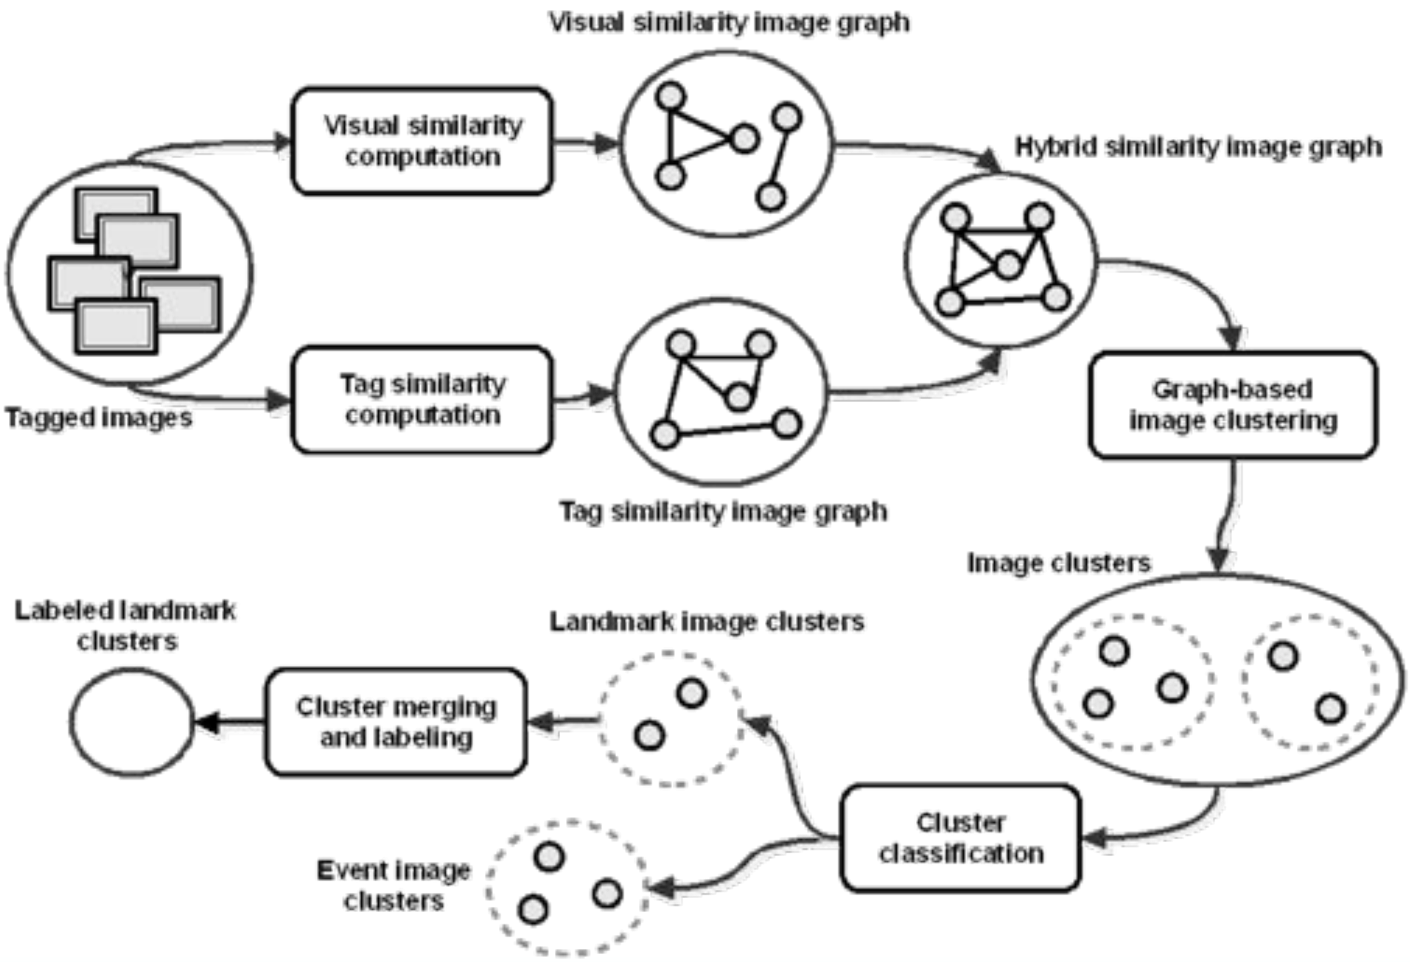

Photo clustering results

Geographic localization of results was also found to be very high.
Most clusters correspond to landmarks or events.

EVENTS

Sample results: [Visual] vs. [Tag] vs. [Visual + Tag]

VISUAL

TAG

HYBRID

clusttour.gr application

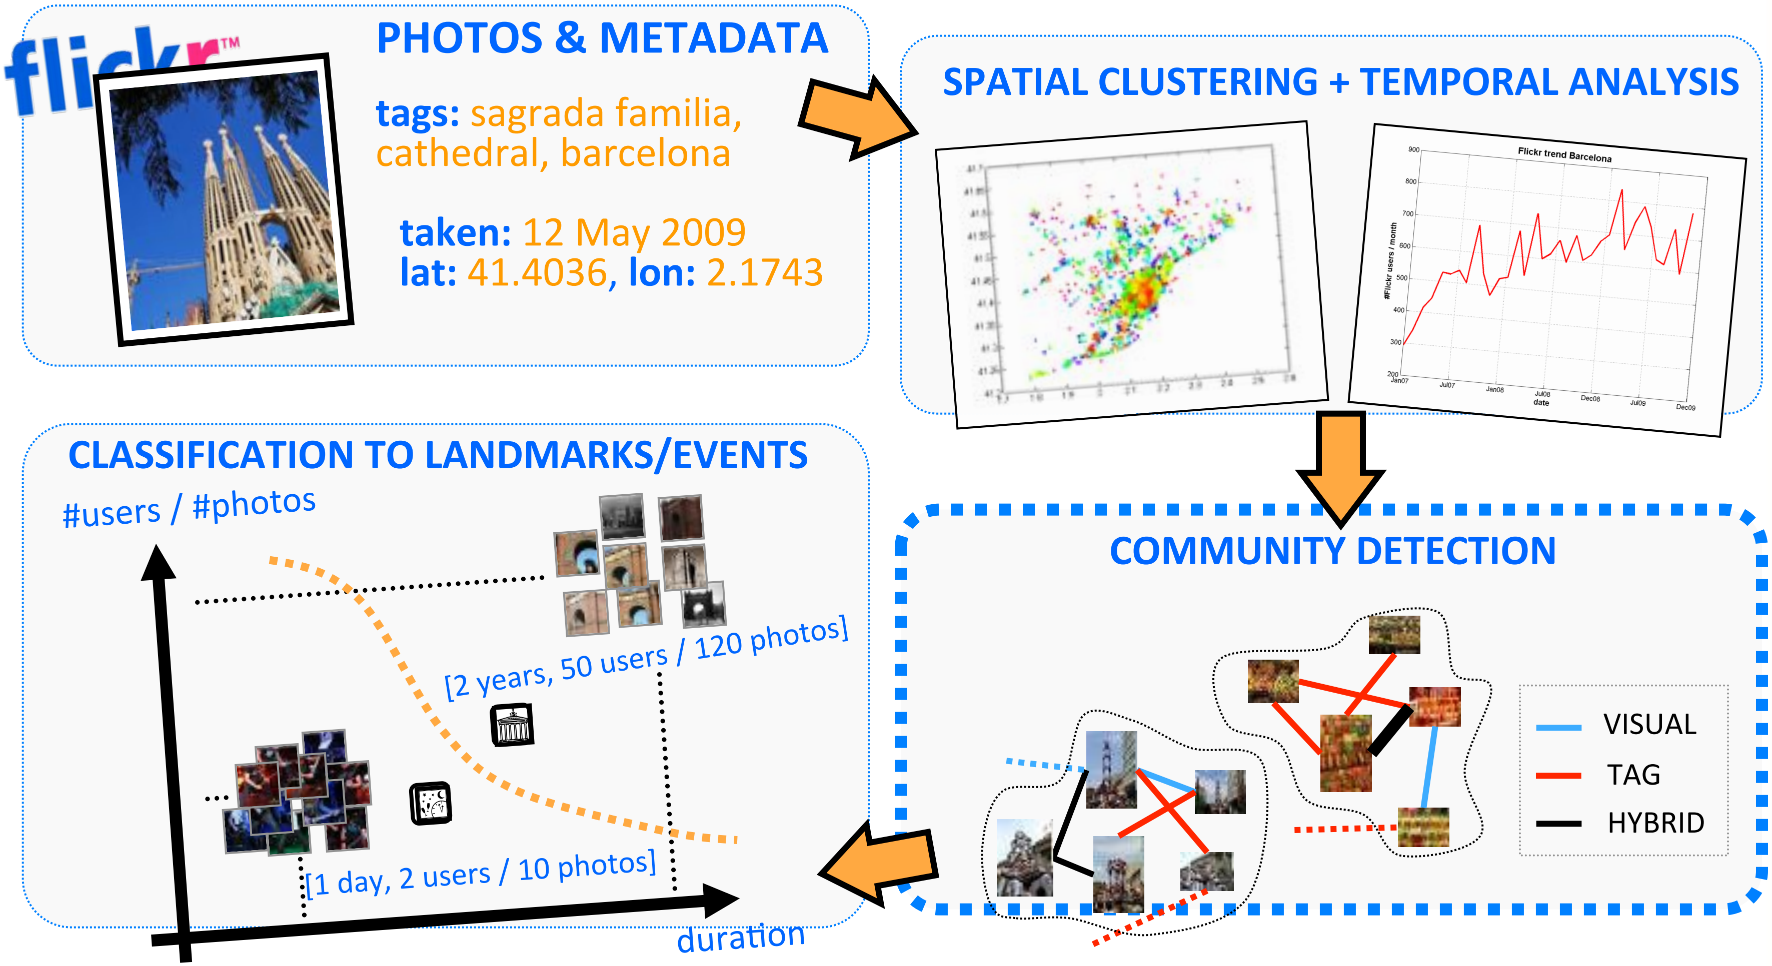

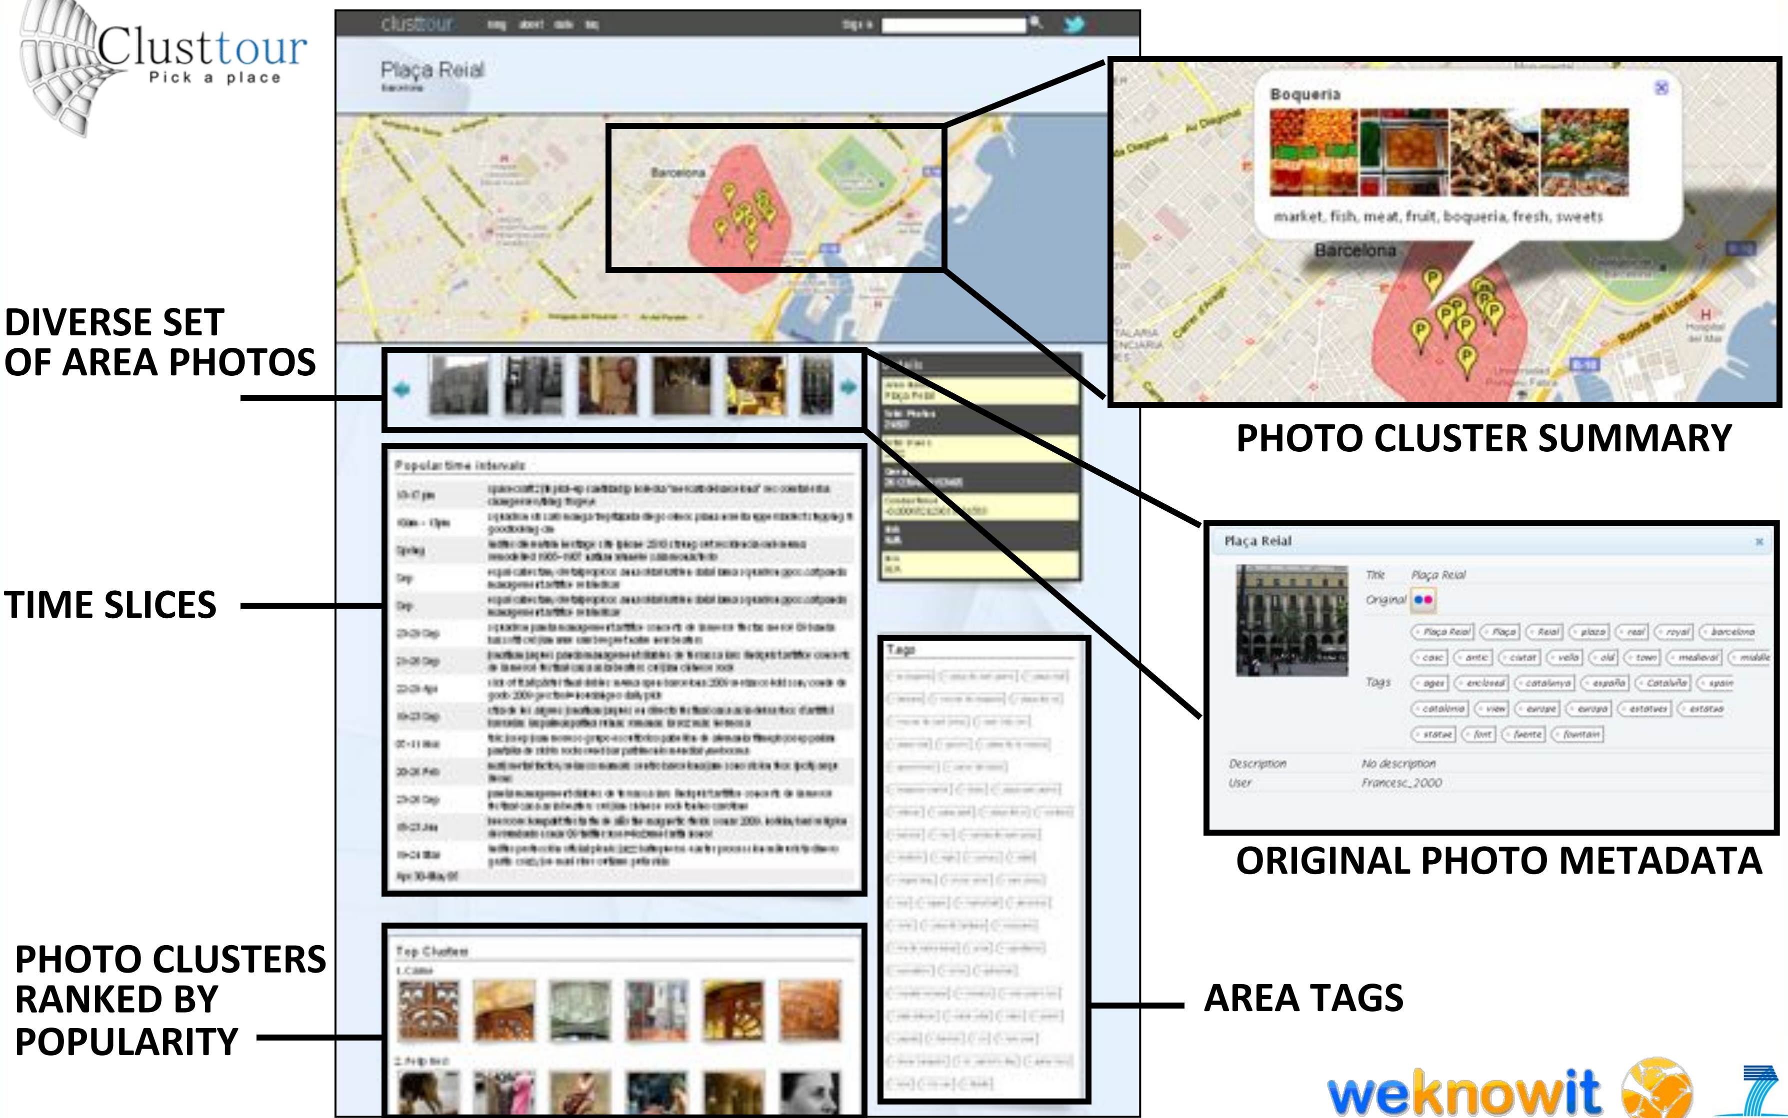
Plaça Reial
Boqueria
market, fish, meat, fruit, boqueria, fresh, sweets
Barcelona
DIVERSE SET OF AREA PHOTOS
PHOTO CLUSTER SUMMARY
TIME SLICES
Popular time intervals
Plaça Reial
Title Plaça Reial
Original
Tags
Description No description
User Francesc_2000
ORIGINAL PHOTO METADATA
PHOTO CLUSTERS RANKED BY POPULARITY
Top Clusters
Tags
AREA TAGS

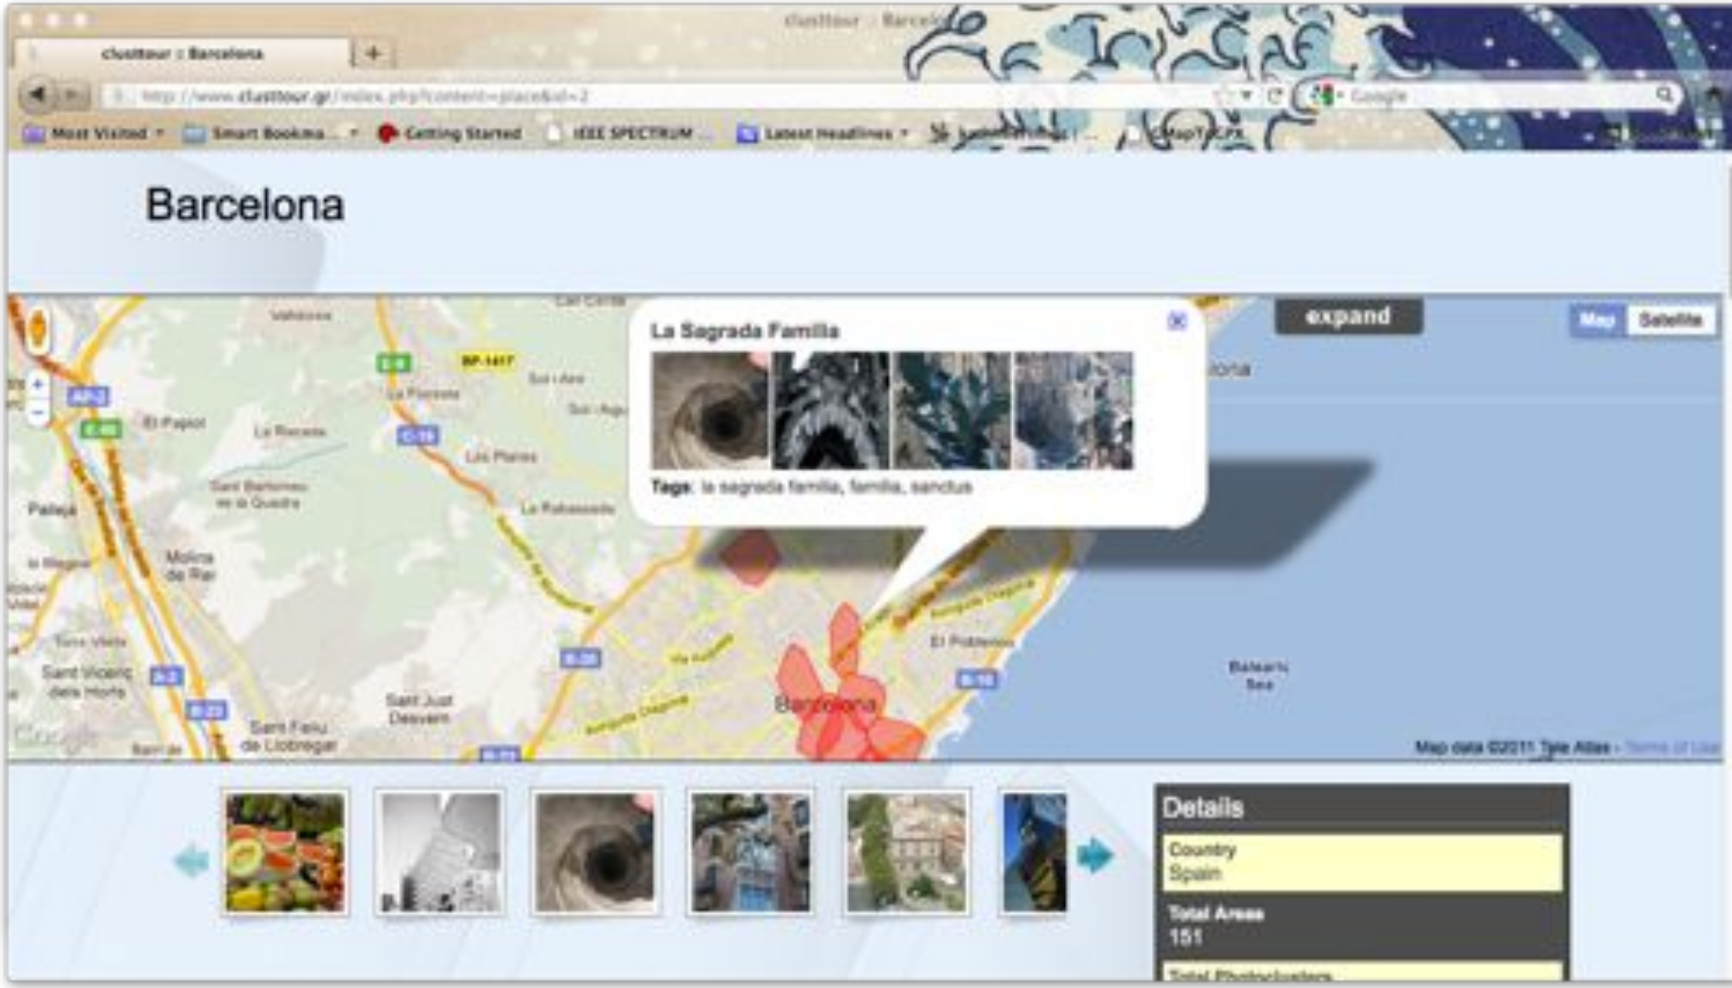
clusttour :: Barcelona
Barcelona
La Sagrada Familia
expand
Map
Satellite
Barcelona
Details
Country
Spain
Total Areas
151

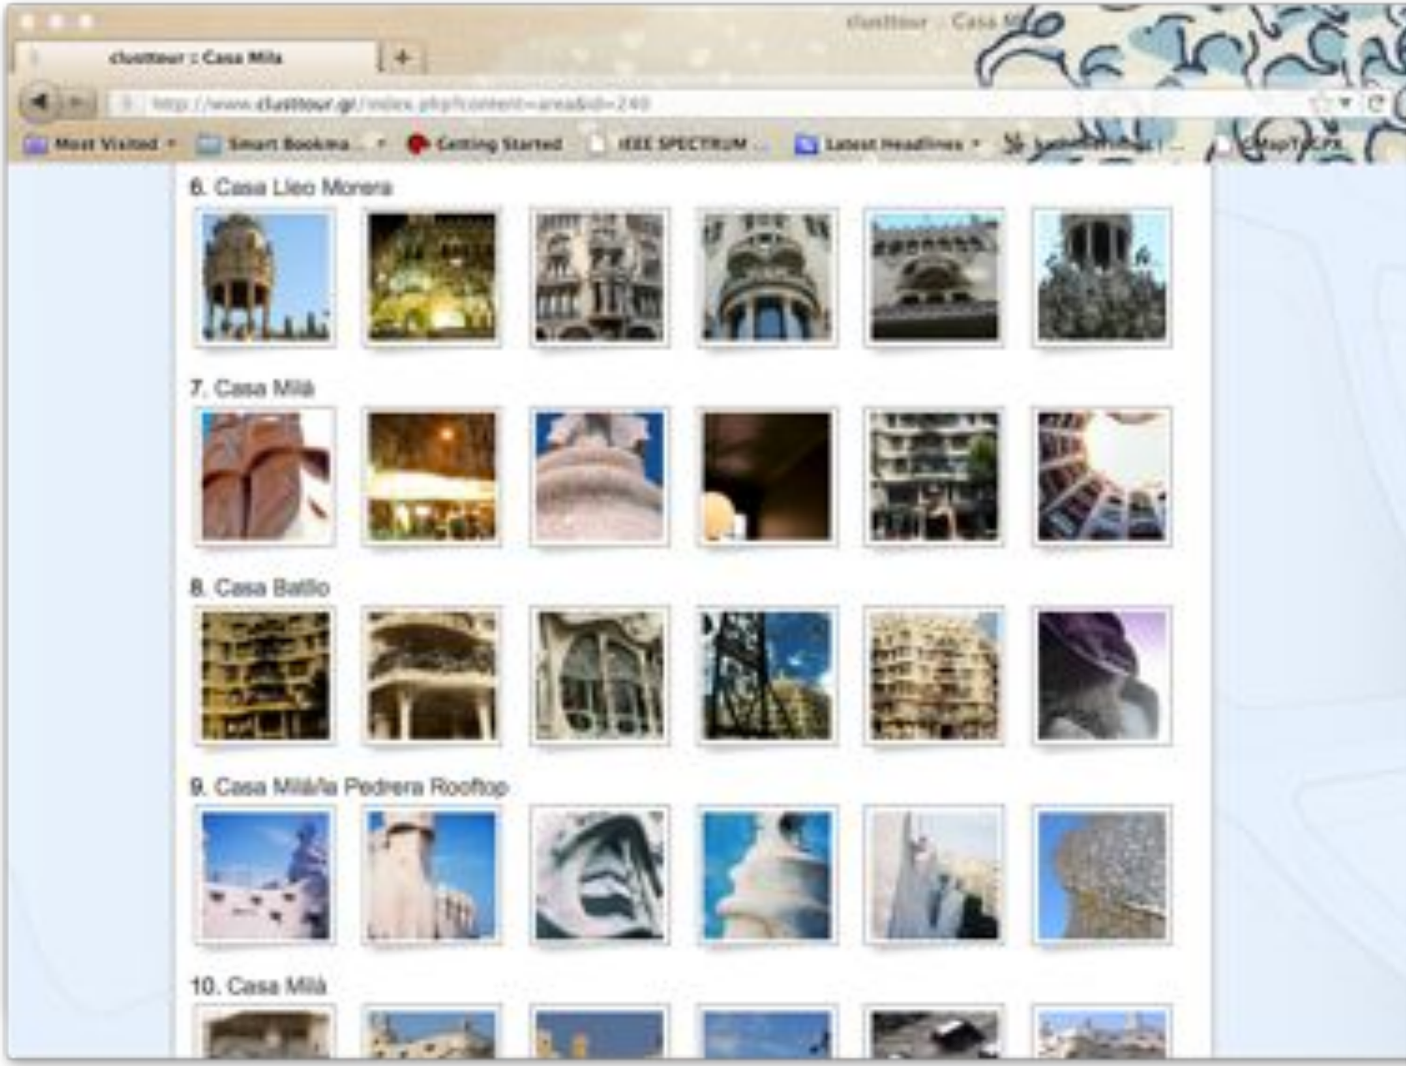
6. Casa Lleo Morera
7. Casa Milà
8. Casa Batllo
9. Casa Milà/la Pedrera Rooftop
10. Casa Milà


clusttour : Homepage
Casa Milà/la Pedrera Rooftop
Casa Mila / Barcelona
expand
Map
Satellite
Avinguda Diagonal
Passeig de Gràcia
Jardins del Palau Robert
Details
Title
Casa Milà/la Pedrera Rooftop
Area
Casa Mila

Social Media “teacher” of the machine

Exploiting clustering for machine learning

Objective: Develop a framework able to create strongly annotated training samples from weakly annotated images

Problems:
- Object detection schemes require region-detail annotations
- Manual annotation is laborious and time consuming

Tagged images

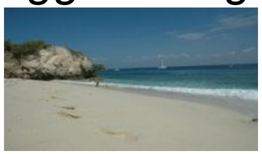

sand, wave, rock, sky

sea, sand

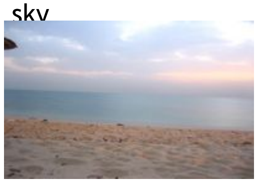

sand, sky

person, sand, wave, see

Social information + Image analysis

Region-detail annotated i

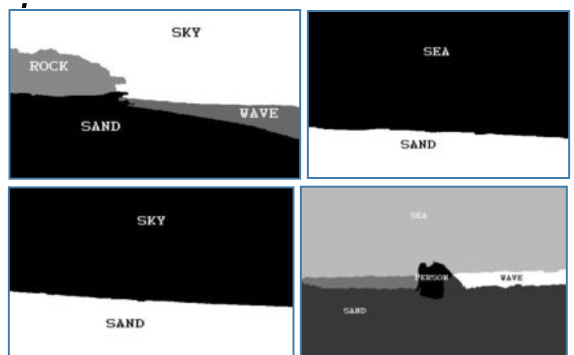

[Chatzilari09]

Machine Learning

Object Detectors

Solutions:
- Exploit user tagged images from social sites like flickr
- Combine techniques operating on tag and visual information space

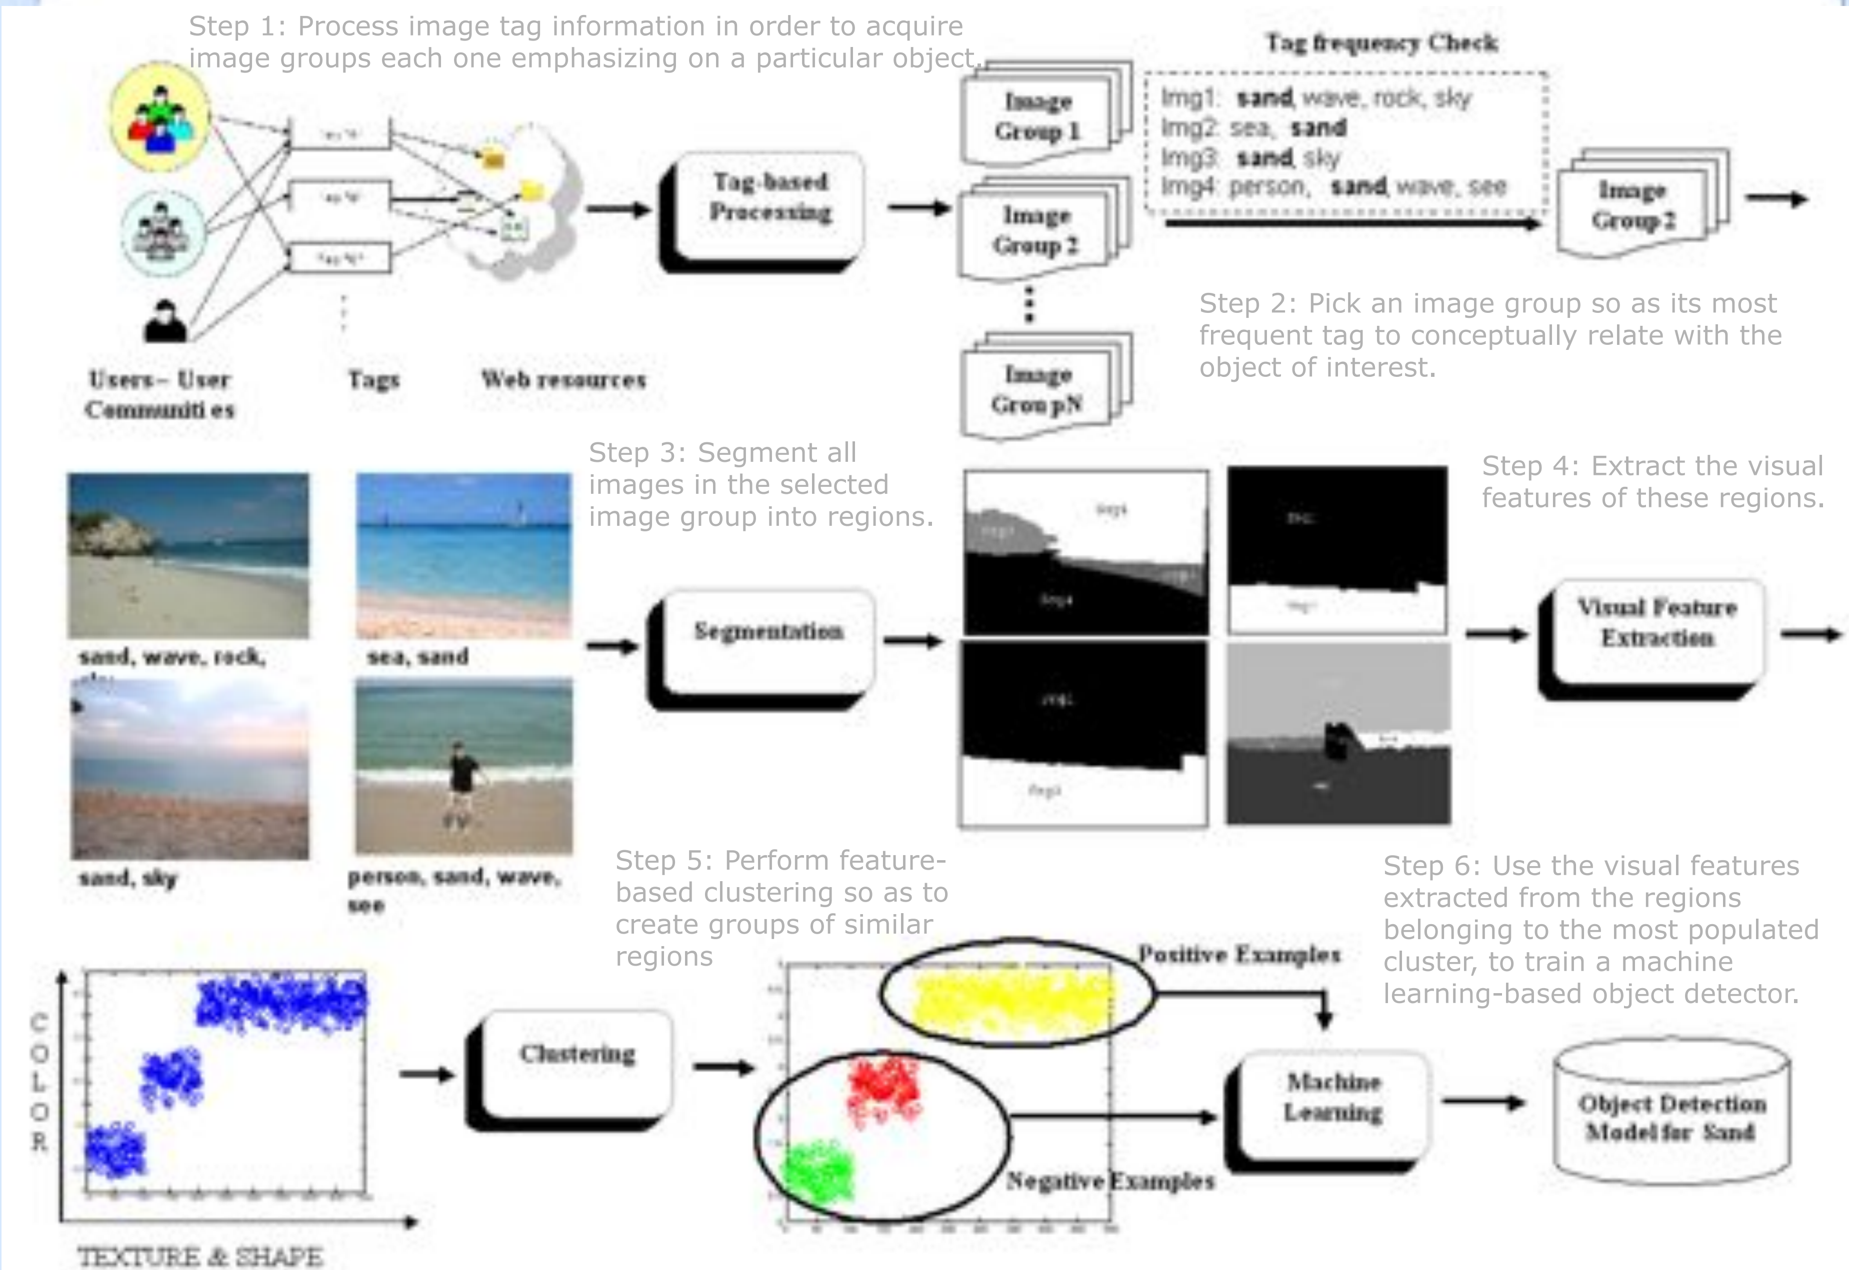
Step 1: Process image tag information in order to acquire image groups each one emphasizing on a particular object.
Users – User Communities
Tags
Web resources
Tag-based Processing
Image Group 1
Image Group 2
Image Group N
Tag frequency Check
Img1: sand, wave, rock, sky
Img2: sea, sand
Img3: sand, sky
Img4: person, sand, wave, see
Image Group 2
Step 2: Pick an image group so as its most frequent tag to conceptually relate with the object of interest.
Step 3: Segment all images in the selected image group into regions.
sand, wave, rock,
sea, sand
sand, sky
person, sand, wave, see
Segmentation
Step 4: Extract the visual features of these regions.
Visual Feature Extraction
Step 5: Perform feature-based clustering so as to create groups of similar regions
COLOR
TEXTURE & SHAPE
Clustering
Positive Examples
Negative Examples
Step 6: Use the visual features extracted from the regions belonging to the most populated cluster, to train a machine learning-based object detector.
Machine Learning
Object Detection Model for Sand

Tag-based processing

[Giannakidou08]

SEMSOC, vector space model where each image is projected onto a space defined by the most prominent tags

SEMSOC output example

Distribution of objects based on their frequency rank

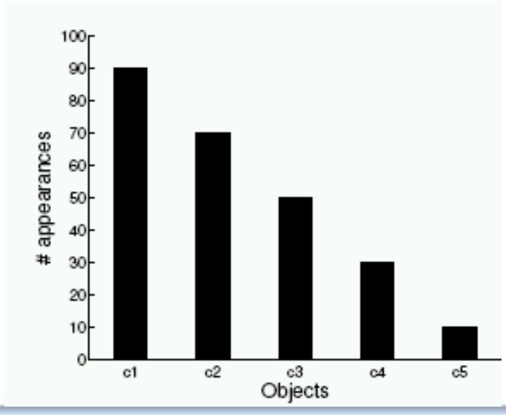

Absolute difference between 1st and 2nd most highly ranked objects increases as n increases

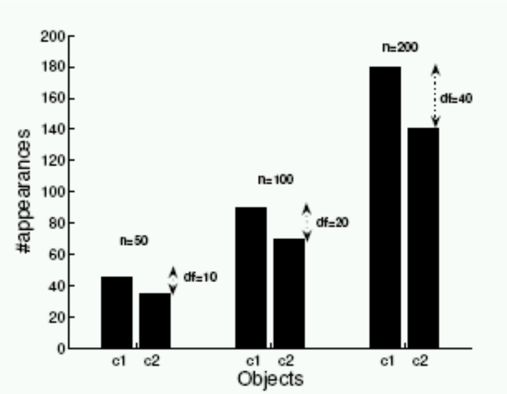

Segmentation & Visual Descriptors

- Segmentation
  - K-means with connectivity constraint (KMCC)

[Mezaris et al., 2004]

- Visual Descriptors
  - MPEG-7 standard
    - *Dominant Color , Color Layout, Color Structure, Scalable Color, Edge Histogram, Homogeneous Texture, Region Shape.*

[Bober et al., 2001], [Manjunath et al., 2001].

Region-based Clustering & Cluster Selection Region clustering

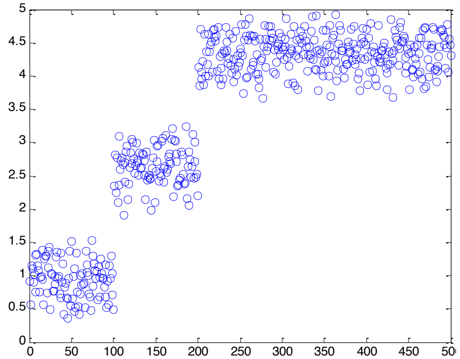

- Perform segmentation and visual feature extraction from all images in an image group (Identified by SEMSOC)

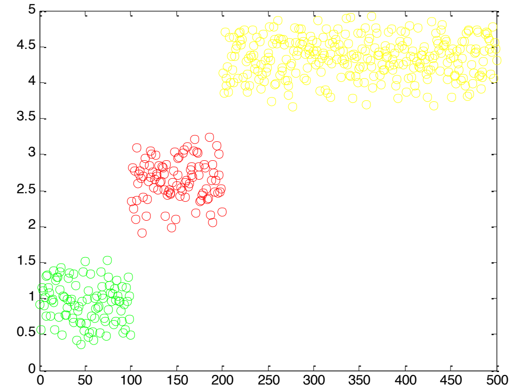

- Perform clustering based on visual features to gather together regions depicting the same object

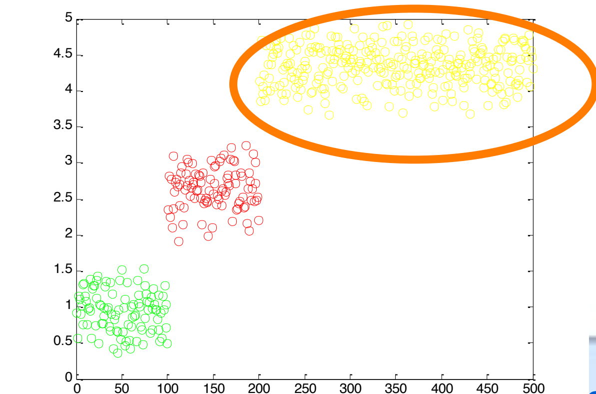

- Pick the most populated cluster as the one representing the most frequently appearing tag of the group

Experimental Results – Cluster Selection

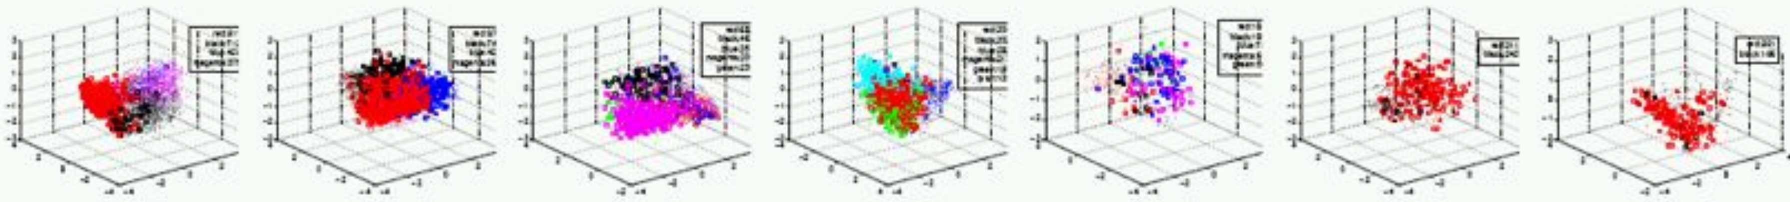

(a) Sky (b) Sea (c) sand (d) Person (e) Boat (f) Vegetation (g) Rock

Setting:

- Visualise the way regions are distributed among clusters
- Use shape-code (squares) to indicate the regions of interest and color-code to indicate a cluster's rank (largest cluster: red)
- Ideally all squares should be painted red and all dots should be painted differently

Goal:

- Validate our theoretical claim that the most populated cluster contains the majority of regions depicting the object of interest

Conclusions:

- Our claim is valid in 5 (i.e., sky, sea, person, vegetation, rock) and not valid in 2 (i.e., boat, sand) cases

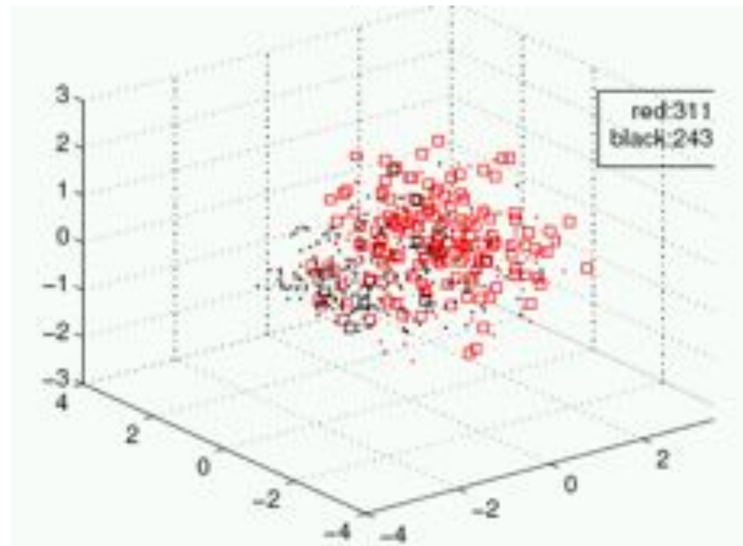

Vegetation in magnification

Experimental Results - Man. vs Autom. trained object detectors

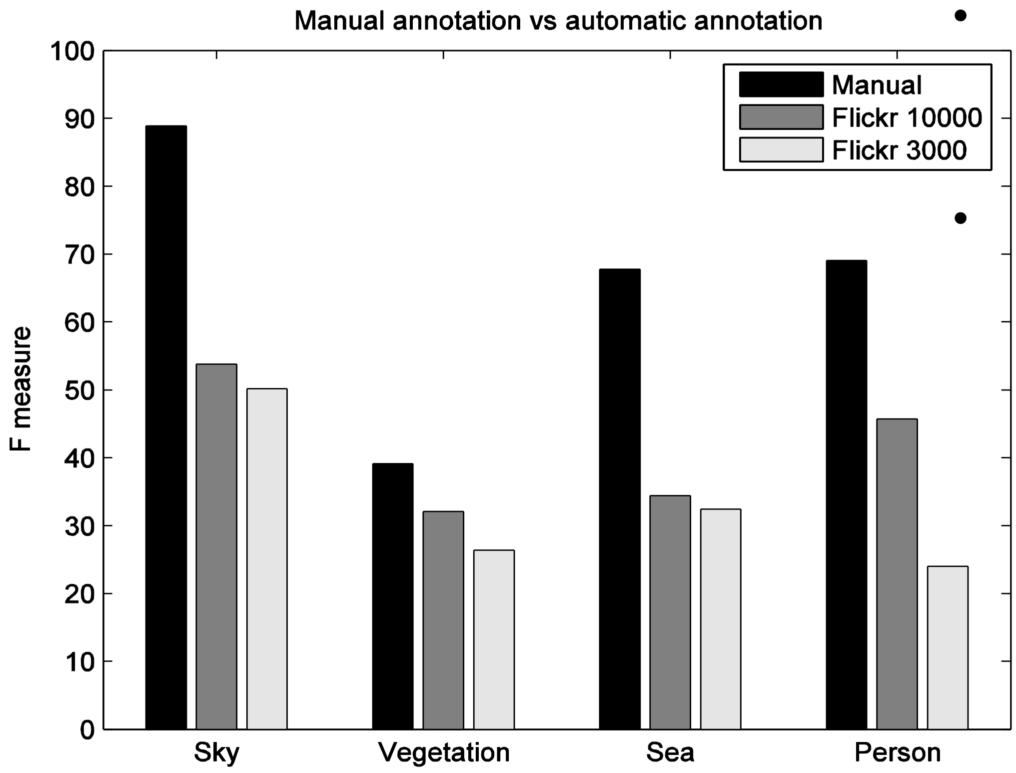

Observations:

- Performance lower than manually trained detectors
- Consistent performance improvement as the dataset size increases

Experimental Results – MSRC Dataset (21 objects)

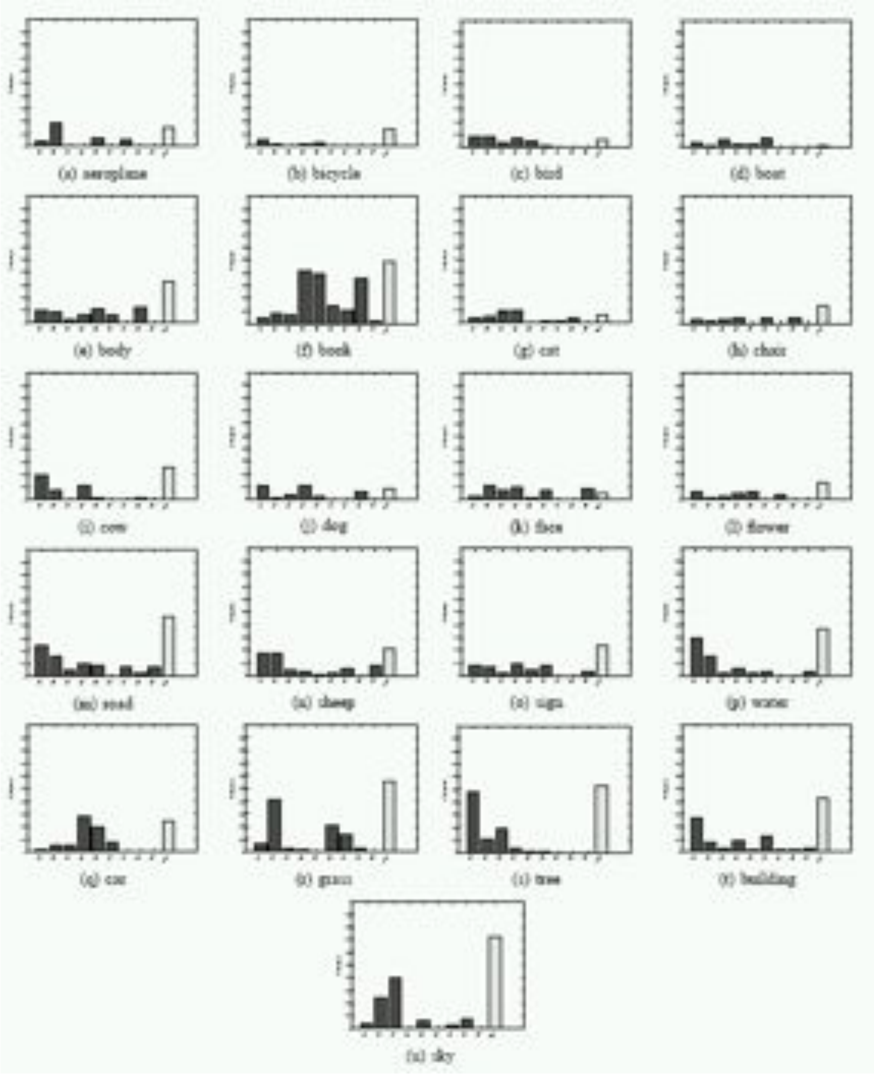

Observations:

- In 5 cases the objects were too diversiform to be described by the employed feature space (not even the manual annotations performed well)
- In 5 cases the annotation we got from Flickr groups were not appropriate
- In 6 cases, our method has failed to select the appropriate cluster
- In 5 cases our method worked well

Experimental Results - MSRC vs Flickr groups Target object: Tree

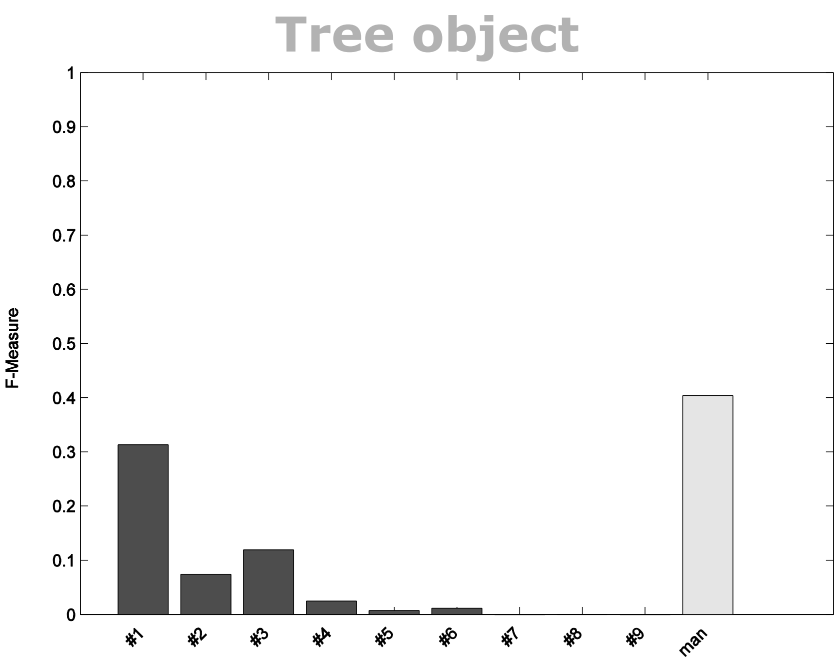

Good example: Semantic objects are correctly assigned to clusters and the most-populated cluster corresponds to the target object)

Experimental Results - MSRC vs Flickr groups Target Object: Sky

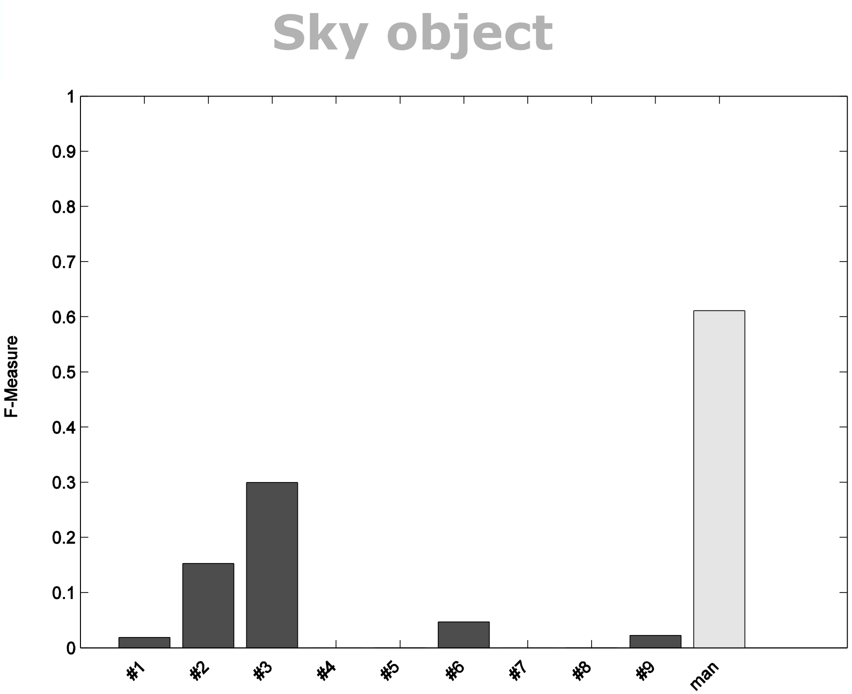

Bad example: Sky regions are split in many clusters and the most populated cluster contains noise regions

WeKnowIt and CI

http://www.weknowit.eu

use case: emergency response

Personal Intelligence

>> Login, Upload

>> Spam detection

>> Personalized Access

Media Intelligence

Photo arrives at ER control centre

>> Automatic localisation of photo

>> Photo & speech auto-tagging

Mass Intelligence

>> Clustering

>> Enrichment from additional sources

Social Intelligence

>> ER Alert Service

>> Reputation Service

Organisational Intelligence

>> Log Merging & Viewing

>> Incident Information Access

use case: travel

Mass Intelligence

>> Landmark & Event detection

>> Ranked facet lists of POIs

>> Hybrid Image Clustering

Media Intelligence

>> Image Localisation

>> Tag suggestions

Travel Preparation

Mobile Guidance

Personal Intelligence

>> Personal Recommendations

Social Intelligence

>> Group profiling & recommendations

>> Friends position, alert

Post Travel

results: research

- User modeling & interaction (CURIO, attention streams)
- Media annotation
  (photo/text localization, photo/speech auto-tagging)
- Media organization
  (graph-based clustering, faceted search, event detection)
- Community analysis & management
  (administration, browsing, reputation, notification)
- Knowledge representation & management
  (Event Model F, dgFOAF)

results: applications

http://www.weknowit.eu/tr

Integrated Prototypes

- ER (desktop & mobile)
- Travel (trip planning, mobile guidance, post-travel photo management)

Stand-alone applications

- WKI image recognizer
- VIRAL (visual search and automatic localization)
- ClustTour (city exploration by use of photo clusters)
- Semaplorer++
- STEVIE (mobile POI management)

results: public APIs

http://mklab.iti.gr/wki-apps

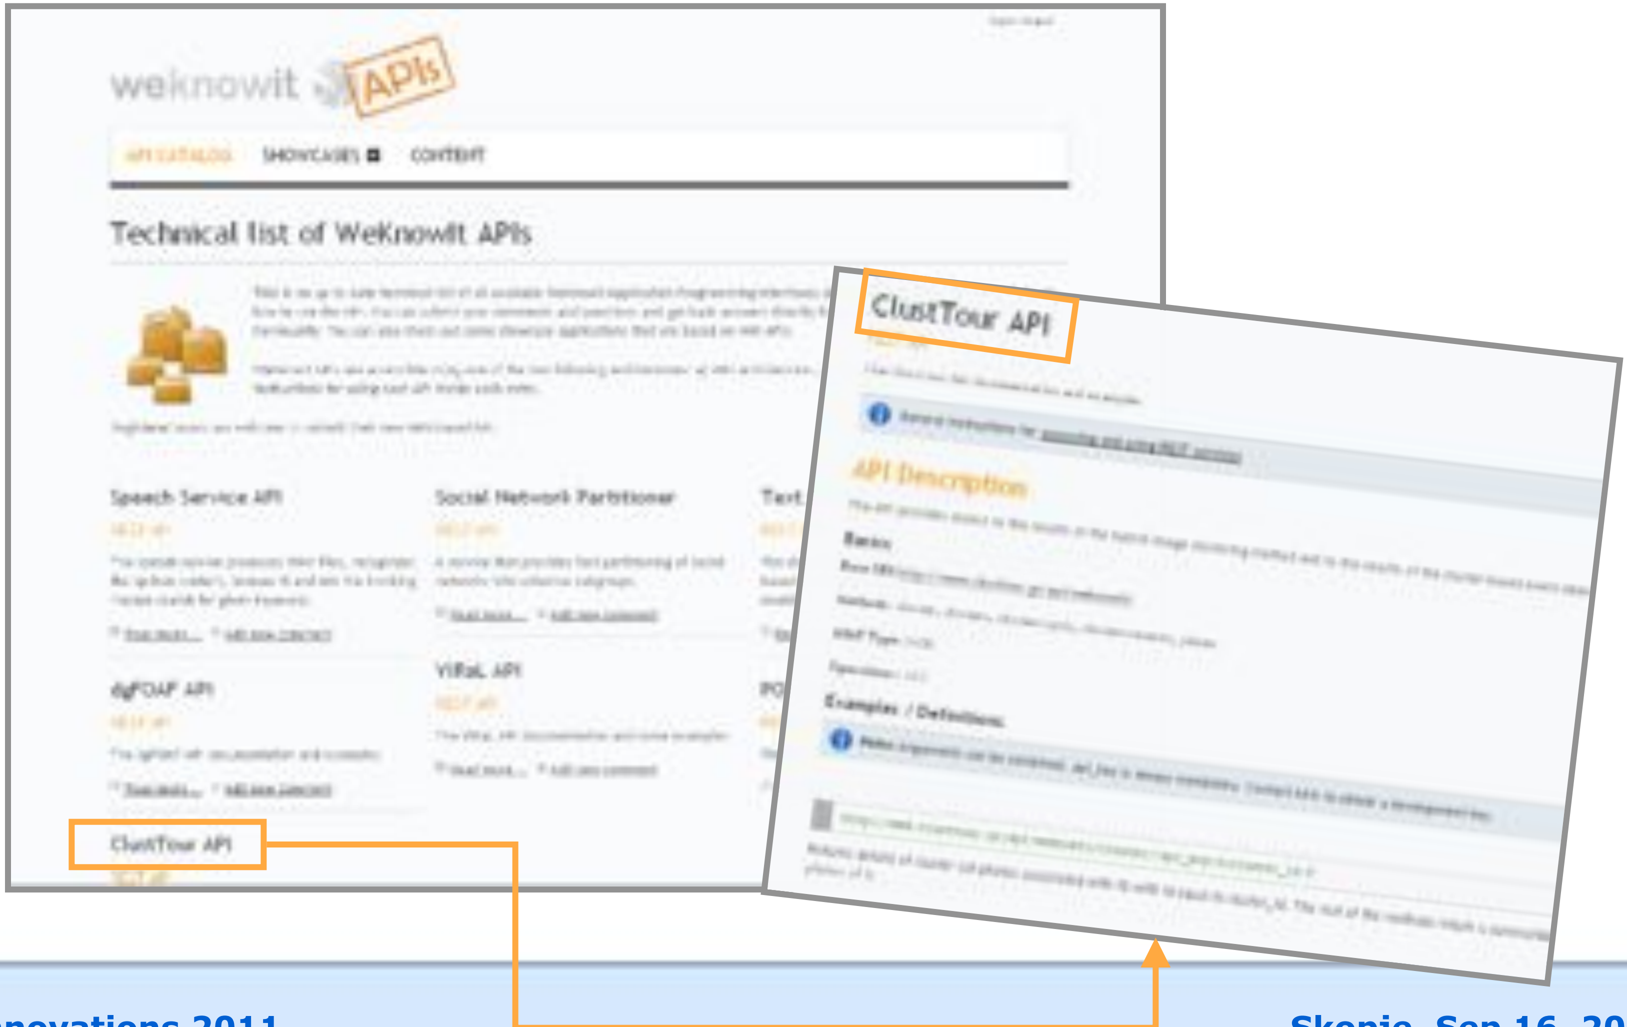

weknowit

http://www.weknowit.eu

Conclusions and Issues

- Social media data mining provides interesting results in many applications
- Not all data always available (e.g. User queries, fb)
- Real-time approaches
  - Efficiency of semantics and analysis
- Real fusion of information
  - not just sum of different analysis
  - formal framework and approach
  - representation
- Linking other sources (web, Linked Open Data)
- Applications and commercialization

Thank you!
http://mklab.iti.gr

References

[Chatzilari09] Elisavet Chatzilari, Spiros Nikolopoulos, Eirini Giannakidou, Athena Vakali and Ioannis Kompatsiaris, "Leveraging Social Media For Training Object Detectors", 16th International Conference on Digital Signal Processing (DSP'09), Special Session on Social Media, 5-7 July 2009, Santorini, Greece.

[Fortunato07a] Santo Fortunato and C. Castellano, “Community structure in graphs”, arXiv:0712.2716v1, Dec 2007.

[Freeman77] L. C. Freeman : A set of measures for centrality . Resolution limit in community detection. PNAS, 104(1). pp. 36-41

[Giannakidou08] E. Giannakidou, I. Kompatsiaris, A. Vakali, "SEMSOC: SEMantic, SOcial and Content-based Clustering in Multimedia Collaborative Tagging Systems", In Proc. 2nd IEEE International Conference on Semantic Computing (ICSC' 2008), IEEE Computer Society, August 4-7, 2008 Santa Clara, CA, USA

[Kennedy07] Lyndon S. Kennedy, Mor Naaman, Shane Ahern, Rahul Nair, Tye Rattenbury: How flickr helps us make sense of the world: context and content in community-contributed media collections. ACM Multimedia 2007: 63

[Kumar99] R. Kumar, P. Raghavan, S. Rajagopalan, and A. Tomkins. Trawling the web for emerging cyber-communities. Computer Networks, 31(11-16):1481-1493, 1999.

[Lindstaedt09] S. Lindstaedt, R. Mörzinger, R. Sorschag, V. Pammer, G. Thallinger. Multimed Tools Appl (2009) 42:97–113, DOI 10.1007/s11042-008-0247-7

[Quack08] Till Quack, Bastian Leibe, Luc Van Gool. World-scale mining of objects and events from community photo collections, In Proceedings of the 2008 international conference on Content-based image and video retrieval, Jul-08

[Zhang06] Y. Zhang, J. Xu Yu, J. Hou : Web Communities: Analysis and Construction, Springer 2006.

Telematics and Informatics

[Crespoa09]Angel García-Crespoa, Javier Chamizoa, Ismael Riverab, Myriam Menckea, Ricardo Colomo-Palaciosa and Juan Miguel Gómez-Berbísa, “SPETA: Social pervasive e-Tourism advisor”, Volume 26, Issue 3, August 2009, Pages 306-315, Mobile and wireless communications: Technologies, applications, business models and diffusion.