Big Social Media

Organizing, Retrieving and Trusting Contents

Francesco De Natale

MMLab@DISI – University of Trento

CNIT – National Consortium of Telecommunications

Acknowledgements

The Multimedia Signal Processing and Understanding Lab (MMLab) deals with relevant aspects of multimedia data processing, focusing on both theoretical and application-driven research issues.

Since media can be considered as the technological extensions of human capabilities, particular attention is paid to smart multimedia data management, analysis, transmission and protection, and to those applications where ambient intelligence can provide advanced services beyond human limitations.

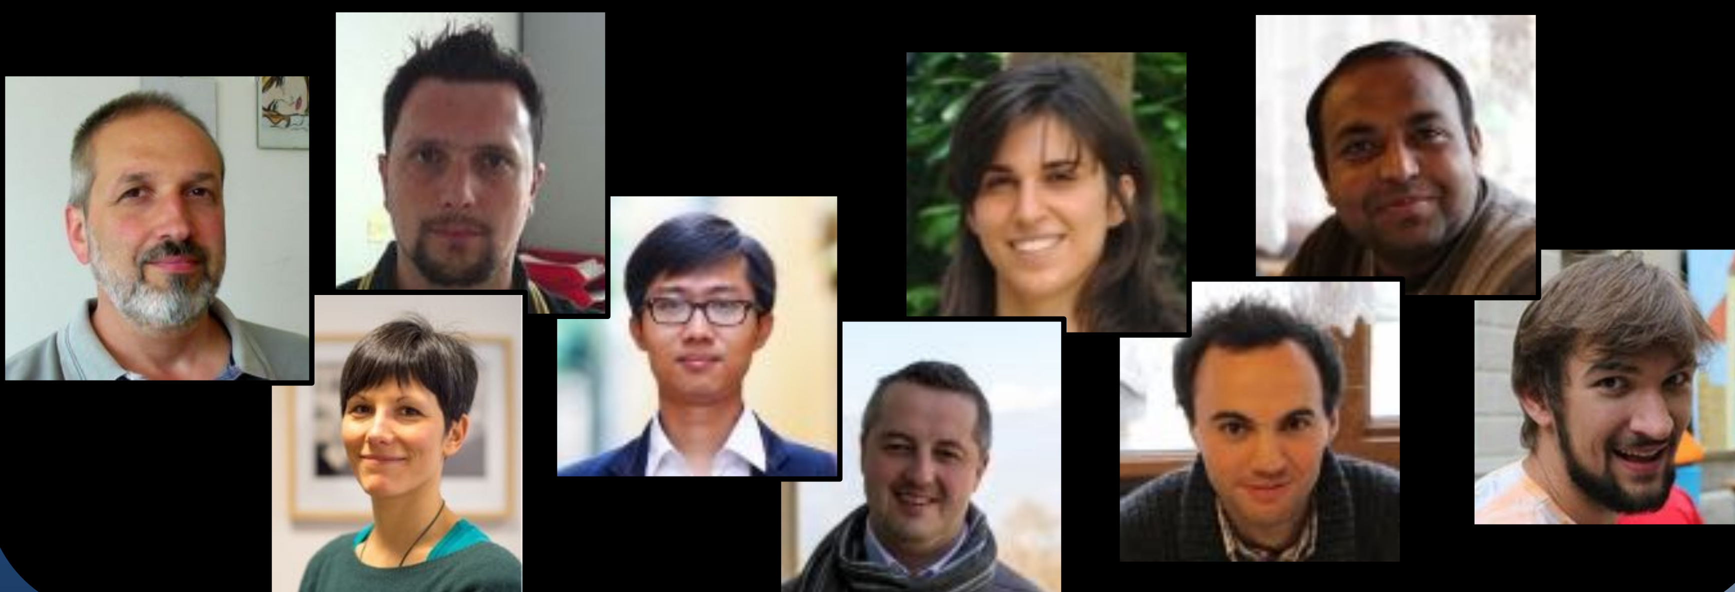

Outline

- Why social media are a good example of ‘big data’
- Challenges of social media:
  - Making structured the unstructured
    - Event-based media structuring
    - Media Synchronization
    - Media Diversity
  - Making trustworthy the untrustworthy
    - Source and provenance detection
    - Manipulation detection and biasing
- Open problems in big social media research

Outline

- Why social media are a good example of 'big data'
- Challenges of social media:
  - Making structured the unstructured
    - Event-based media structuring
    - Media Synchronization
    - Media Diversity
  - Making trustworthy the untrustworthy
    - Source and provenance detection
    - Manipulation detection and biasing
- Open problems in big social media research

What is “big”

- “Big” is a relative concept:
  - a data system is ‘big’ if the hw/sw used to capture and process it, is unable to do that in reasonable time/space [1]
  - what is “big” today can become treatable tomorrow…

    … unless it continues to grow: in that case we have a problem!

[1] *Teradata 2011, McKinsey Global Institute*

Is that big enough?...

Source: George Carey-Simos, wersm.com

‘Big’ does not only refer to size:

- Huge (and increasing) in volume
- Largely unstructured
- Largely diverse in content and format
- Partly untrustworthy
- Originating from different sources
- Used for indirect scopes (different from the application they have been collected for)

‘Big’ does not only refer to size: the 3V model

- According to Doug Laney (Gartner’s Inc, 2001):
  - Big data is high-volume, -velocity and –variety information assets that demand cost-effective, innovative forms of information processing for enhanced insight and decision making.

    also known as the “3-Vs model”
  - Later extended to many more Vs (2011), including:
    - Variability, Value, Veracity, Volatility, Vulnerability, Verification, …

    from 3-Vs to 12-Vs and more

The 3V model

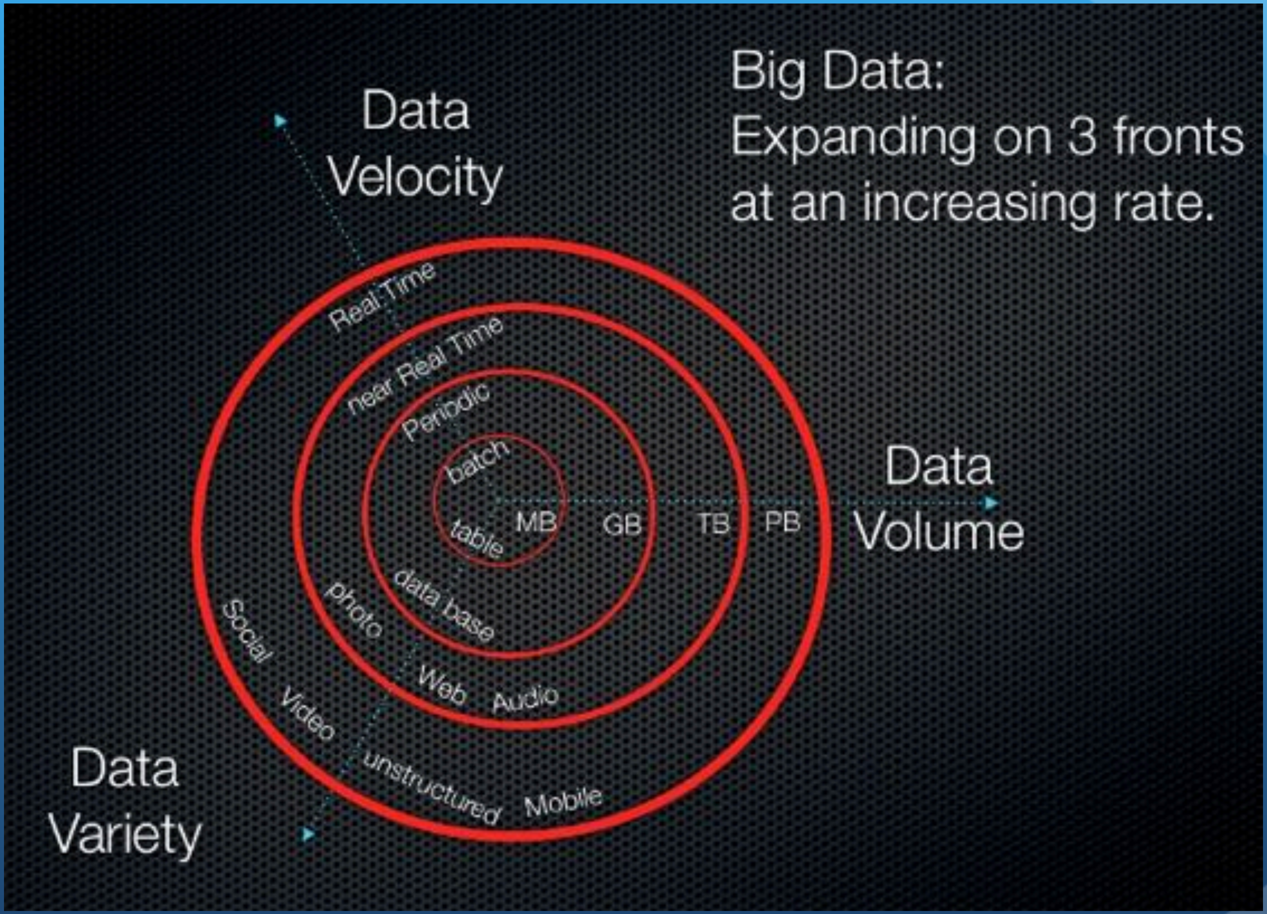

Source: Diya Soubra's, The 3Vs that define Big Data, on Data Science Central

The 3 Vs of Social media

	Monthly active users	Average user activity	Data volumes and generated traffic	Data types
Facebook	2.1B	20 min/day	350 Mphoto/day 9M messages/h 100Mh video watched/day	Images, video, audio, text, graphics, links, …
Youtube	1,57B	Avg. session 40’	5+ B shared videos today 5 B watched videos/day 300 h upload videos/min	Video and audio
Instagram	800M	8 likes/day 1 photo upld/week	40B photo shared 95M photo upld/day	Photos, stories
Twitter	330M	5 tweets/day	100M tweets+img/year	Mostly text but also images and video

+ instant messaging, webcams, professional media, ...

Source: Salman Aslam, Omnicore Agency

The other ‘Vs’…

- Value:
    - Facebook stock market valutation tops $245 billion
    - Annual revenue of Google for Youtube is around $13 billion
    - Instagram is expected to generate about $5 billion in mobile advertising sales in 2018
    - Twitter is currently valued at $16 billion
  - Value connected to content, users, ads, influencing
- Variability and volatility:
  - Contents can be modified, removed, linked, re-used, etc.
  - No guarantee to find the same information the next day

The other ‘Vs’…

- Veracity and Verification
  - Fake news are becoming a big problem in Social networks
  - Upload rate is overwhelming: impossible to check contents

and the trends…

- Also trends are typical of big data
  - continuous growth (although rate tends to saturate)

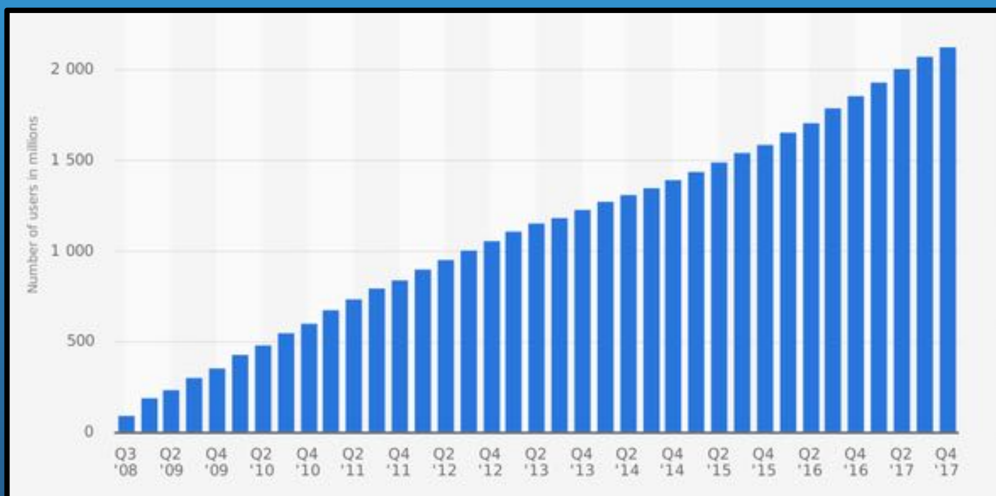

Monthly active Facebook users '10-'17

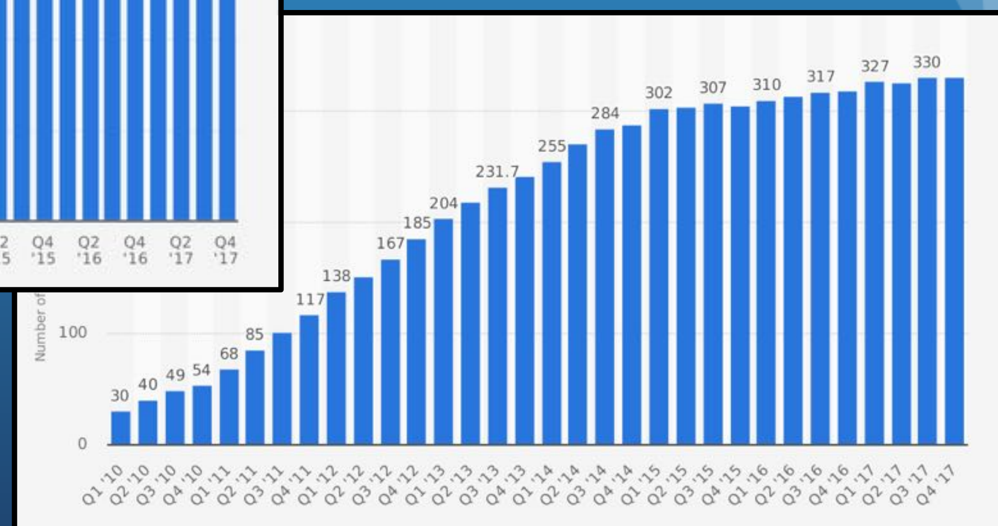

Monthly active Twitter users '08-'17

Source: statista.com

The challenges of Big Social Media

- Social Media are indeed big data
- Multimedia in general is today the killer application for systems and networks
  - Video is expected to cover 80% of total Internet traffic in 2019
- There is a dramatic direct and indirect value connected to media contents, BUT…
  - to make it (re-)usable and profitable, there is a need of structure, accessibility, trust, cross-linking, etc.
  - we will focus on two major challenges:
    - *Making structured the unstructured*
    - *Making trustworthy the untrustworthy*

A content-based perspective

Outline

- Why social media are a good example of 'big data'
- Challenges of social media:
  - Making structured the unstructured
    - Event-based media structuring
    - Media Synchronization
    - Media Diversity
  - Making trustworthy the untrustworthy
    - Source and provenance detection
    - Manipulation detection and biasing
- Open problems in big social media research

Making structured the unstructured

- Structure is needed to make data easy to access, (re)use, retrieve, organize, summarize, present, ...
- Structure can be achieved through indexing, annotation, linking, time-stamping and any other instrument that allows capturing the relationships among pieces of content along some dimension (facet)
- We will focus on three concepts:
  - Event-based media structuring
  - Media synchronization
  - Media diversity

Media vs. Semantics

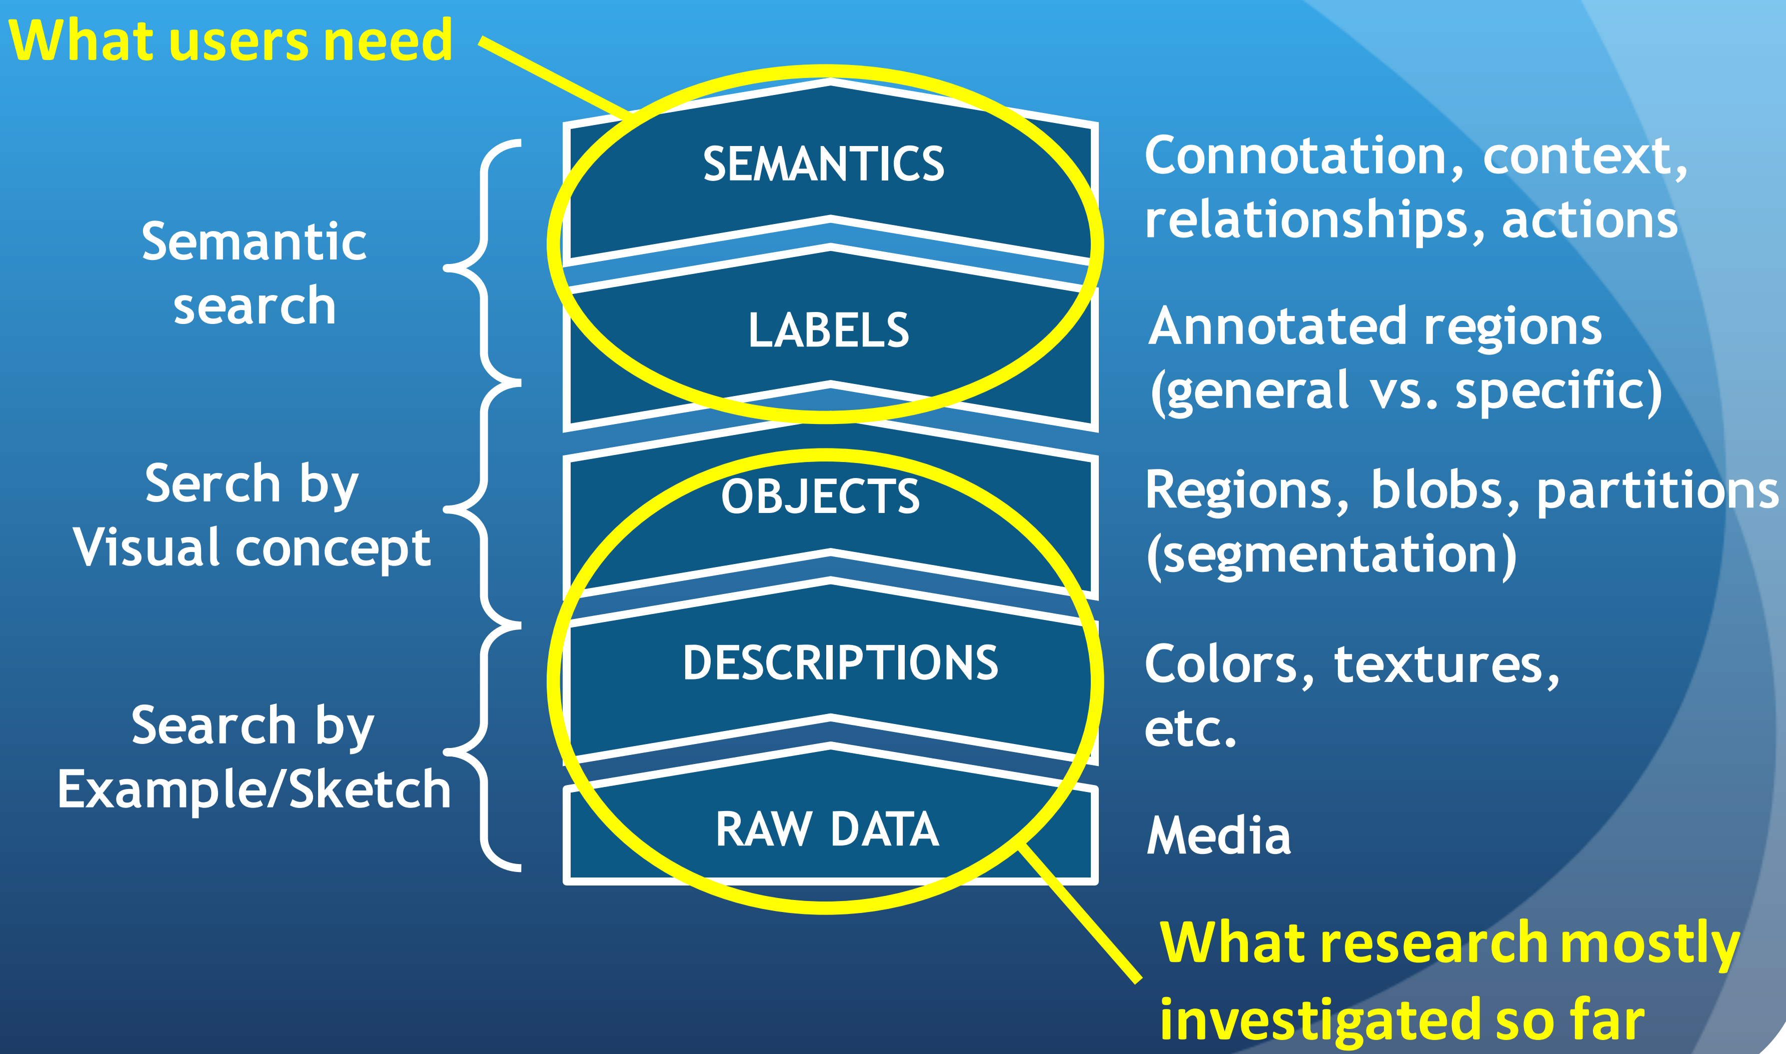

Media vs. Semantics

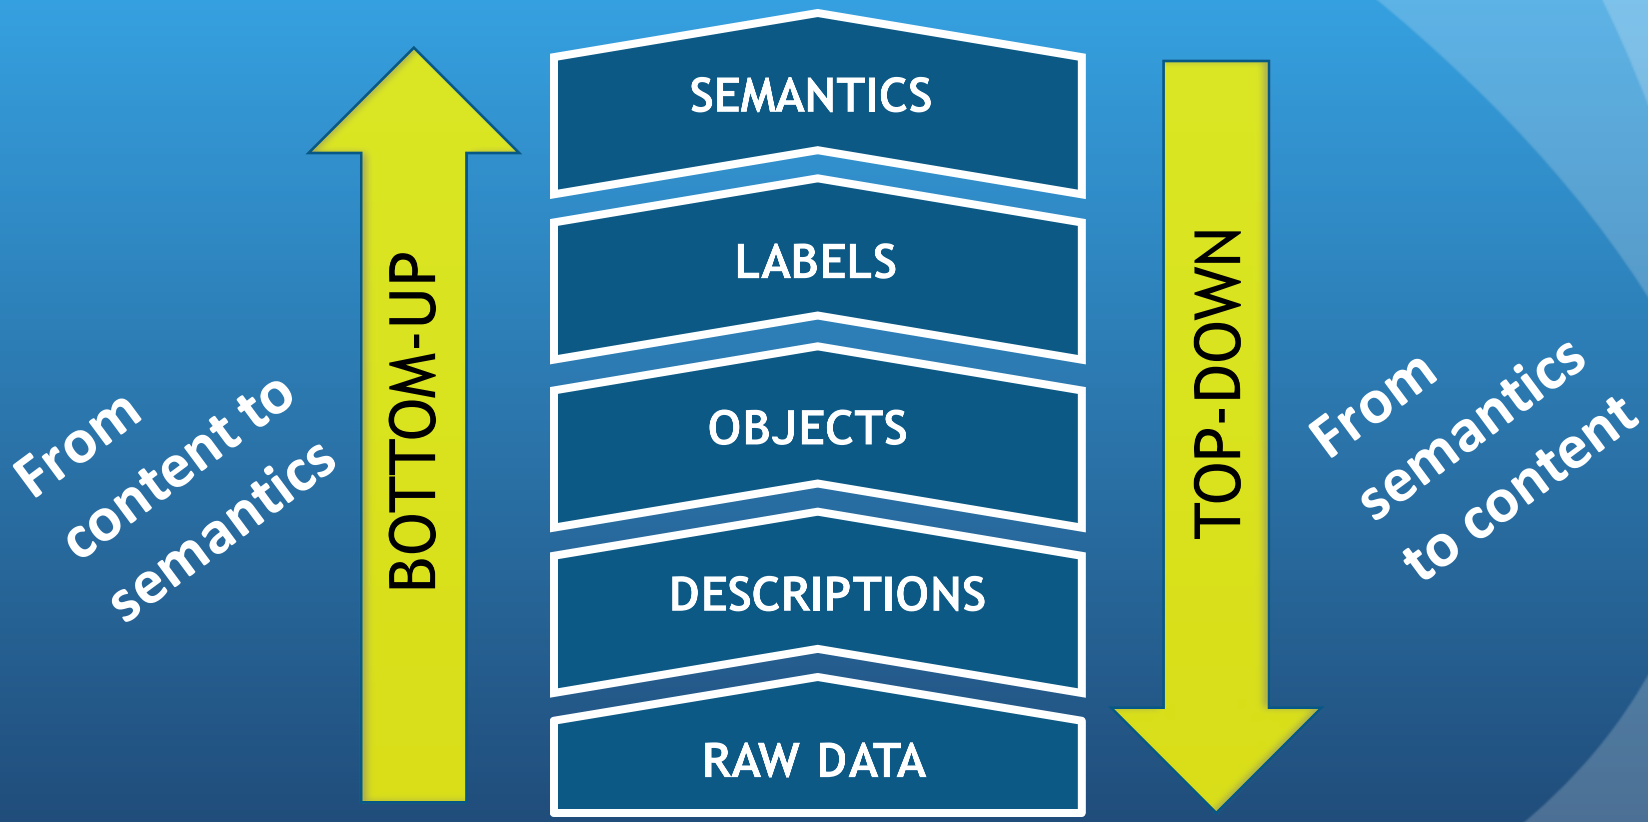

How to model semantics in media

- Explicit Knowledge representation
  - Taxonomies (description and classification of things)
  - Ontologies (formal conceptualizations of domains)
  - Events (formal conceptualizations of whatever is happening)
- Human-based computation
  - Relevance feedback (capture human semantics via interaction)
  - Crowdsourcing (use collective intelligence)
  - Gamification (use human intelligence ‘concealed’ by games)

Media and Events

- An interesting idea to put associate a semantic to media is using the concept of events
- Events have been widely used in information retrieval and successively applied to media retrieval
- Events are the way we organize our own memories, then they are a very natural way to organize media
- Events facets typically answer to questions such as:
  - What, Where, When, How, Who, …

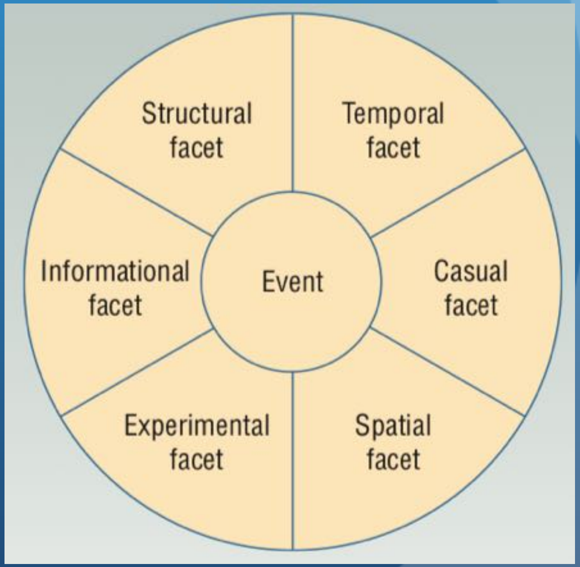

R. Jain, "EventWeb: Developing a human-centered computing system, IEEE Computer, 2008

Event webs

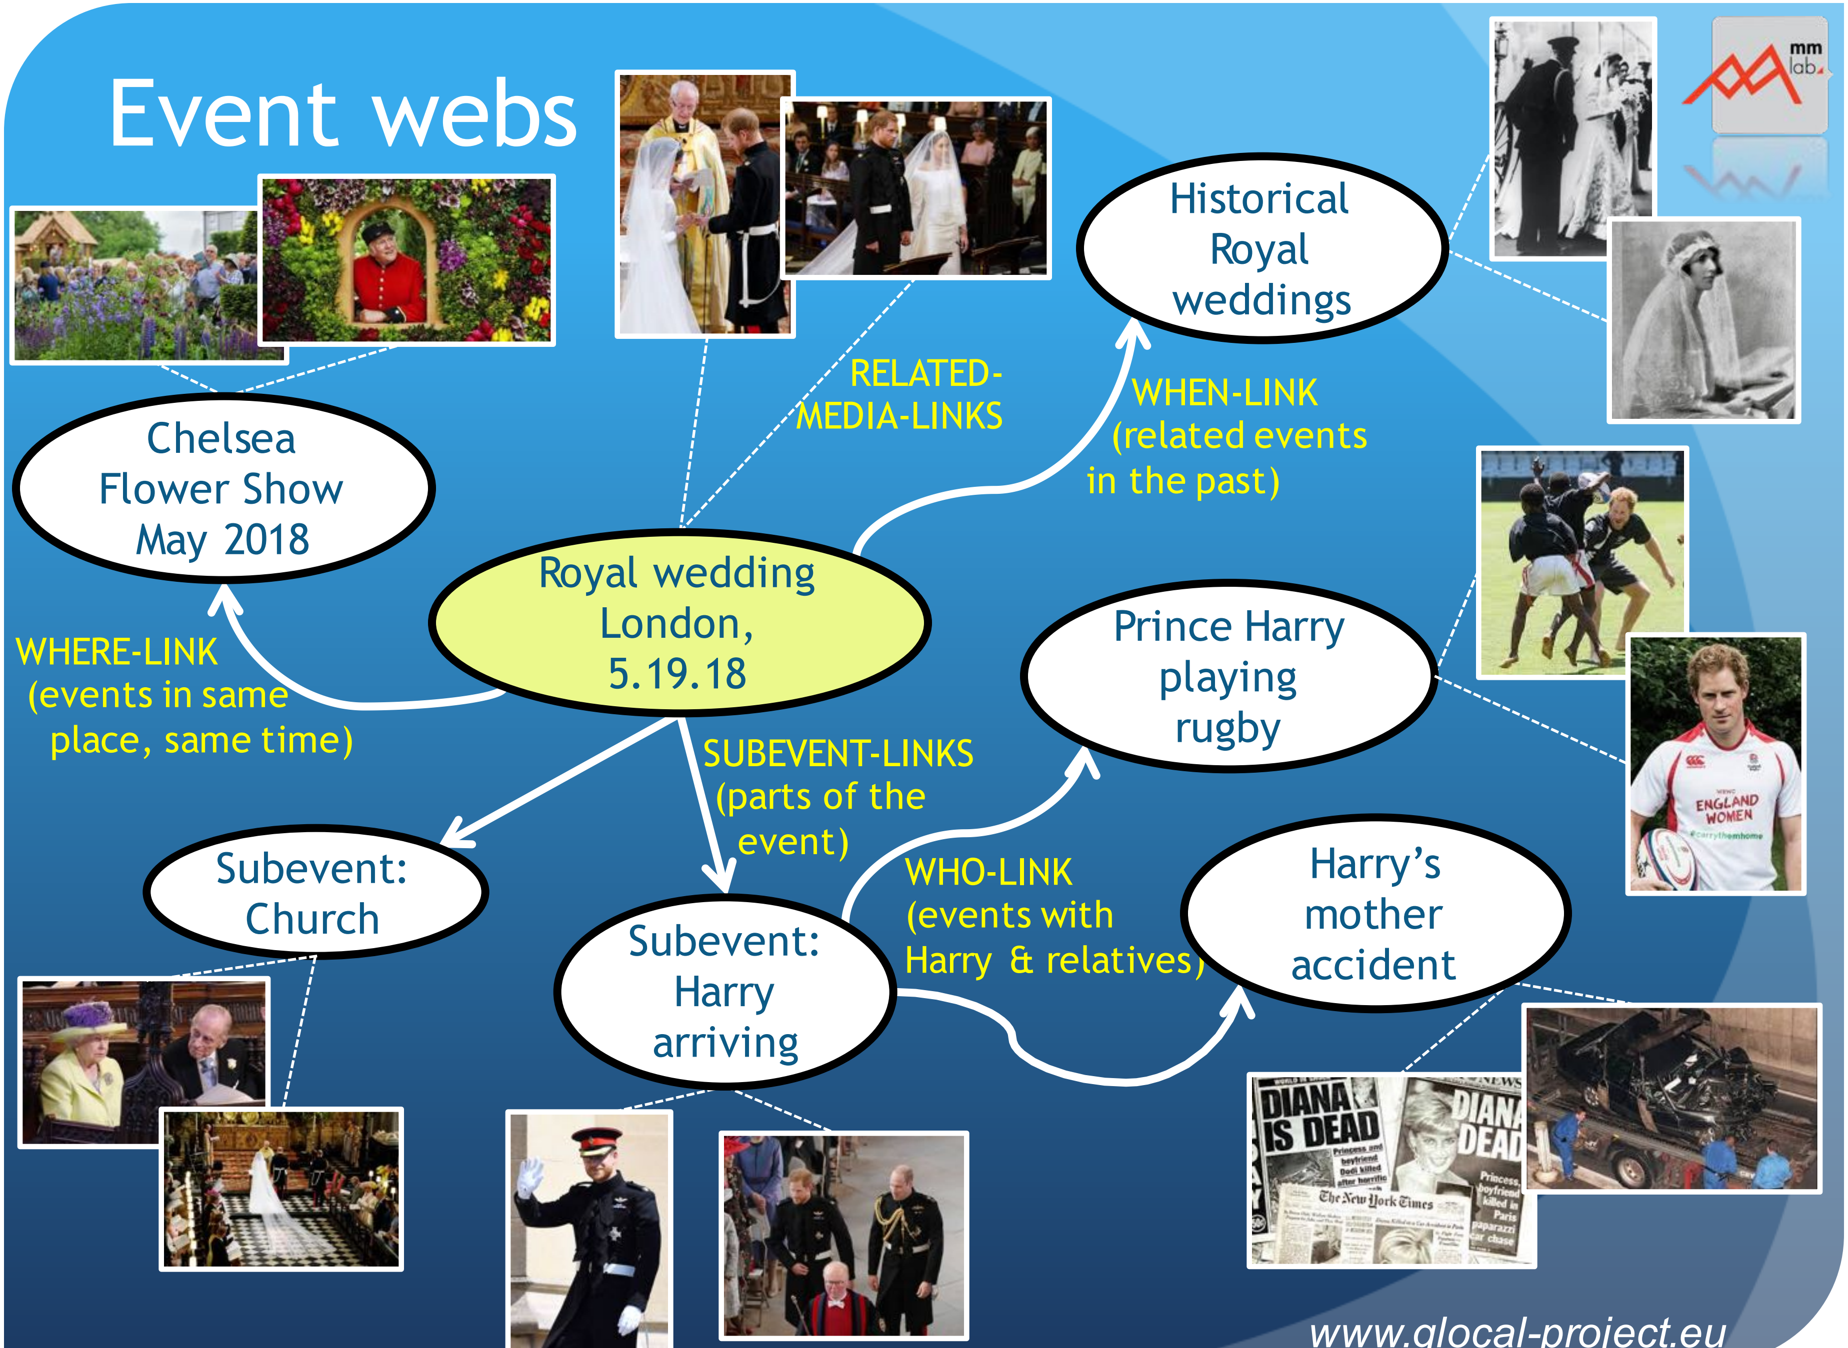

Historical Royal weddings
RELATED-MEDIA-LINKS
WHEN-LINK (related events in the past)
Chelsea Flower Show May 2018
Royal wedding London, 5.19.18
WHERE-LINK (events in same place, same time)
Prince Harry playing rugby
ENGLAND WOMEN
SUBEVENT-LINKS (parts of the event)
Subevent: Church
WHO-LINK (events with Harry & relatives)
Harry's mother accident
Subevent: Harry arriving
DIANA IS DEAD
DIANA DEAD
The New York Times

Discovering events from media

- All of this is possible if there is a way to automatically associate media and events
  - This rises many research questions:
    - How is it possible to discover the event associated to a media item?
    - How is it possible to discover the event associated to a collection of media (e.g., a photo album)?
    - What is important in a media item to understand the associated event?
    - How is it possible to link media that share some facets?

The concept of Event Saliency

- Recently, we introduced the concept of **event saliency**
  - Identifying the parts of the image that reveal the event
  - Not as simple as **visual saliency** (typically based on appearance, color, contrast, position, etc.)
- Two major problems
  - Detecting the event saliency
    - Gamification experiments
    - Crowdsourcing experiments
  - Using the event saliency
    - Introducing saliency in media event detection

Which events are those photos related to?

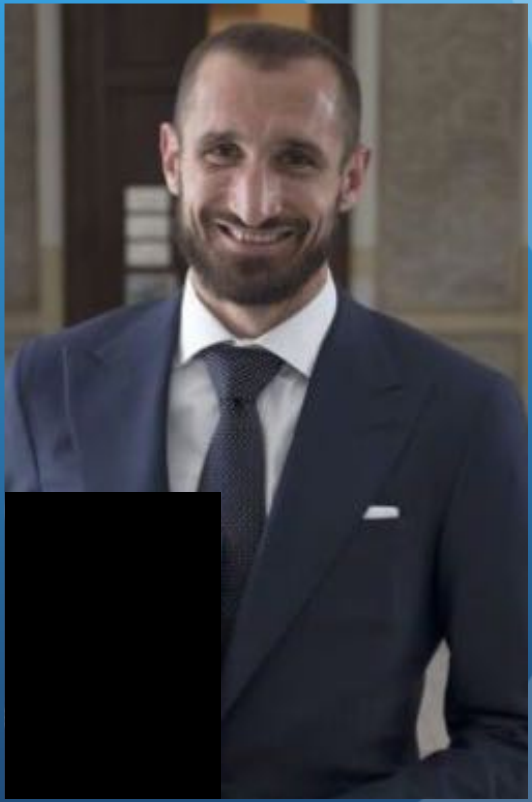

Which events are those photos related to?

A gamified event saliency detector

- An example of human computation using games
- EventMask is an adversarial game formulated as inversion problem:
  - We ask people to hide only what may reveal the event
  - ‘Maskers’ win if ‘discoverers’ cannot recognize the event
  - Earned points are inversely proportional to covered area
  - Players earn point in both hiding and discovering
- Event saliency map is generated combining maskers and discoverers results.

A.Rosani, G.Boato, F.De Natale, “EventMask: A game-based framework for event-saliency identification in images”, IEEE Trans. on Multimedia, 2015

Dataset: http://mmlab.science.unitn.it/EventMaskDataset/

Example of EventMask saliency

Original Image

Event saliency

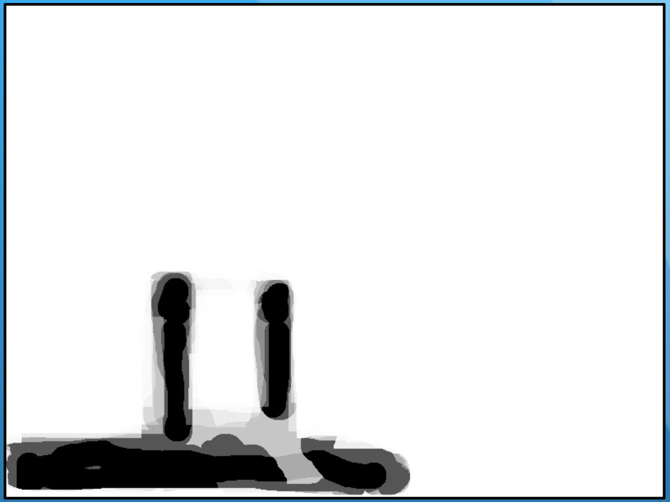

Generated masks

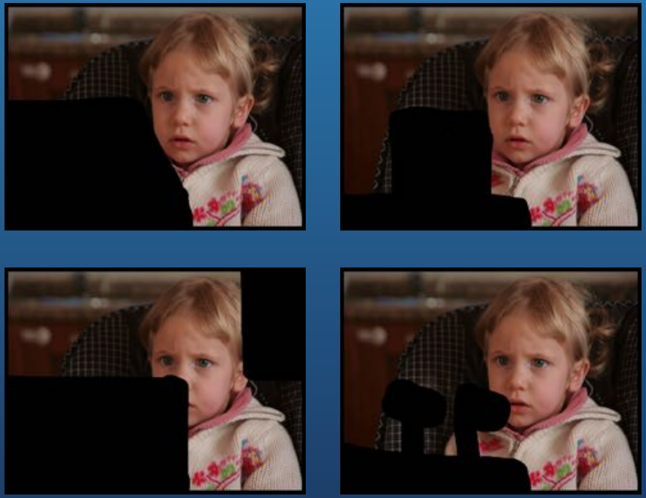

Traditional saliency

Event saliency by crowdsourcing

- As an alternative, we used crowdsourcing for event saliency detection
  - Image fragments are generated by an agent, and micro-tasks consist in discovering the event
  - Reward is in form of micro-payments
  - Mechanisms are used to prevent cheating

Introduction to the Crowdsourcing Task

We are carrying out non-profit research at a university to build an event retrieval system. By accepting this task, you agree that we may publish parts of your answers as part of our research study. We will NOT publish any information that could be linked to you. We do NOT use your worker ID, or any other information that links to you, during data analysis or storage. Your answers are used only by researchers for the purposes of gaining insight into general opinions concerning events related multimedia. Beyond the people who are doing research in this area, no other parties are allowed to use your answers.

Event Representation via a region

Questions

(i) From these 8 possible events, which one do you think has been the one presented to you?

- Option 1
- Option 2
- ⋮
- Option n

(ii) Briefly explain, why did you choose the particular option in question i (open question)

Next

K.Ahmad, N.Conci, F.De Natale, "A Saliency-based Approach to Event Recognition", Signal Processing: Image Communication, 2017

Salient region extraction process

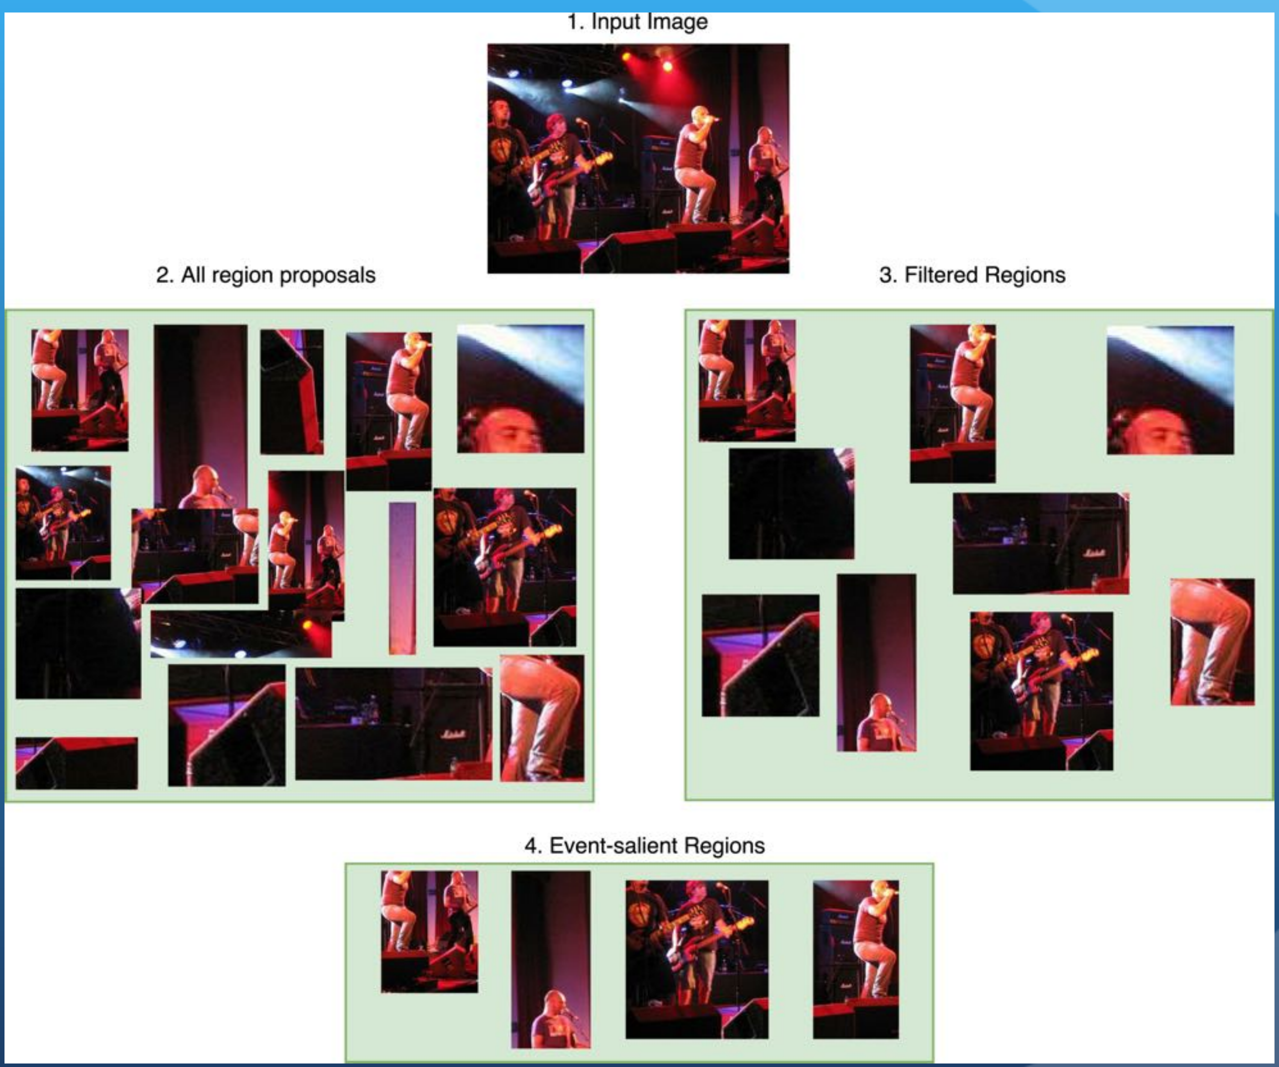

Using saliency for event discovery from single image: MIL

- We adopt a Multiple-Instance-Learner to associate salient visual concepts to the relevant events

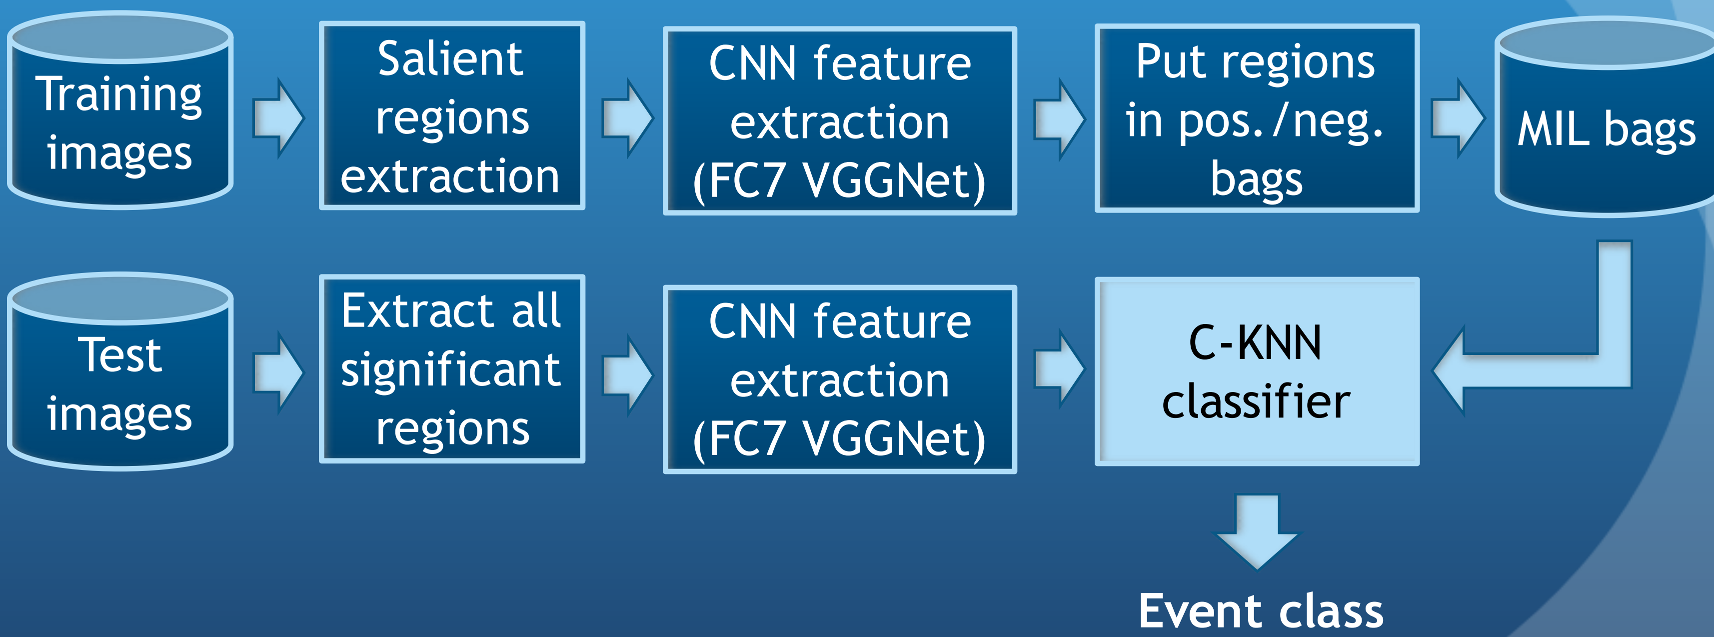

K.Ahmad, N.Conci, F.De Natale, “A saliency-based approach to event recognition”, Signal Processing: Image Communications, 2017

Using saliency for event discovery from single image: MIL

- Results on different event detection datasets:

Method	Avg. Accuracy
Schinas et al 2012	0.334
Rosani et al 2015	0.459
Ahmad et al 2015	0.700
Ahmad et al 2016	0.858
Proposed	0.912

SED 2013

Method	Avg. Accuracy
Ahmad et al 2015	0.700
Rachmadi et al 2016	0.720
Proposed	0.771

USED

Method	Avg. Accuracy
Li et al 2007	0.734
Zhou et al 2014	0.944
Zhou et al 2016	0.950
Wang et al 2016	0.988
Proposed	0.983

UIUC

SED
MediaEval Social Event Detection (SED) Task 2011-14 - www.mediaveal.org
USED
K. Ahma, N.Conci, G.Boato, F.De Natale, USED: A large scale social event detection dataset, MMSys 2016 - http://loki.disi.unitn.it/~used/
UIUC
L-J.Li, L.Fei-Fei, "What, where and who? Classifying event by scene and object recognition, ICCV 2007 - http://vision.stanford.edu/lijiali/event_dataset/

Using saliency for event discovery from photo collections: MIL

- Even more challenging is the case of multiple images (photo collections, albums)
  - Annotation at collection-level only
  - Not every picture is relevant (e.g., close-ups, outliers)
- Again, MIL can provide superior performance

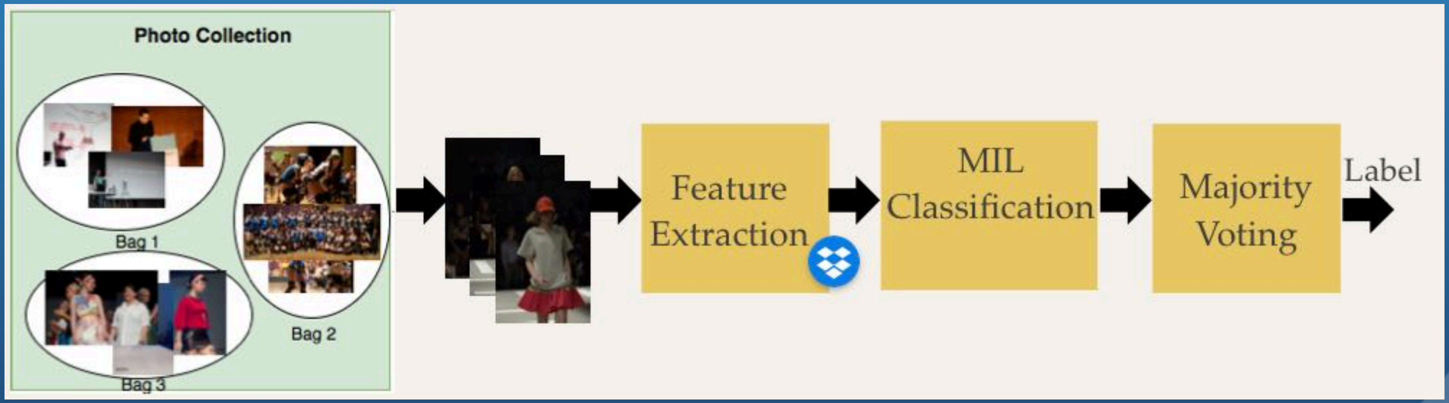

K.Ahmad, N.Conci, G.Boato, F.De Natale, "Event recognition in personal Photo collections via multiple instance learning-based classification of Multiple images", Journal of Electronic Imaging, 2017

Using saliency for event discovery from photo collections: MIL

- Comparative results on PEC dataset[1]:

Method	Avg. Acc.	Method	Avg. Acc.
AgS [1]	0.4143	ShMM [1]	0.5571
Huang et al. [2]	0.7343	HAS [3]	0.8632
R-OS-PGM [4]	0.7428	SVM-CNN (Majority voting)	0.8230
MIL_SMO	0.9153	**Our Approach**	**0.9527**

[1] L.Bossard, M.Guillamin, L.VanGool, "Event recognition in photo collections with a stopwatch HMM, ICCV, 2013

Using saliency for event discovery: ensamble of deep models

- Both object-related and background regions may contain useful information, complementing each-other
- We proposed a method based on the fusion of different classifiers, differently trained for the two areas
  - Object-level info coded by a network pre-trained on ImageNet dataset
  - Scene-level info coded by a network pre-trained on Places dataset
  - Different neural architectures are considered (AlexNet, VGGNet, GoogleNet, ResNet)
  - 3 different fusion models are considered (IOWA, PSO, GA)

*K.Ahmad, M.L.Mekhalfi, N.Conci, F.Melgani, F.De Natale, "*Ensemble of Deep Models for Event Recognition*", ACM* Trans. on Multimedia Computing Communications and Applications*, 2018*

Using saliency for event discovery: ensamble of deep models

- Overview of the proposed scheme:

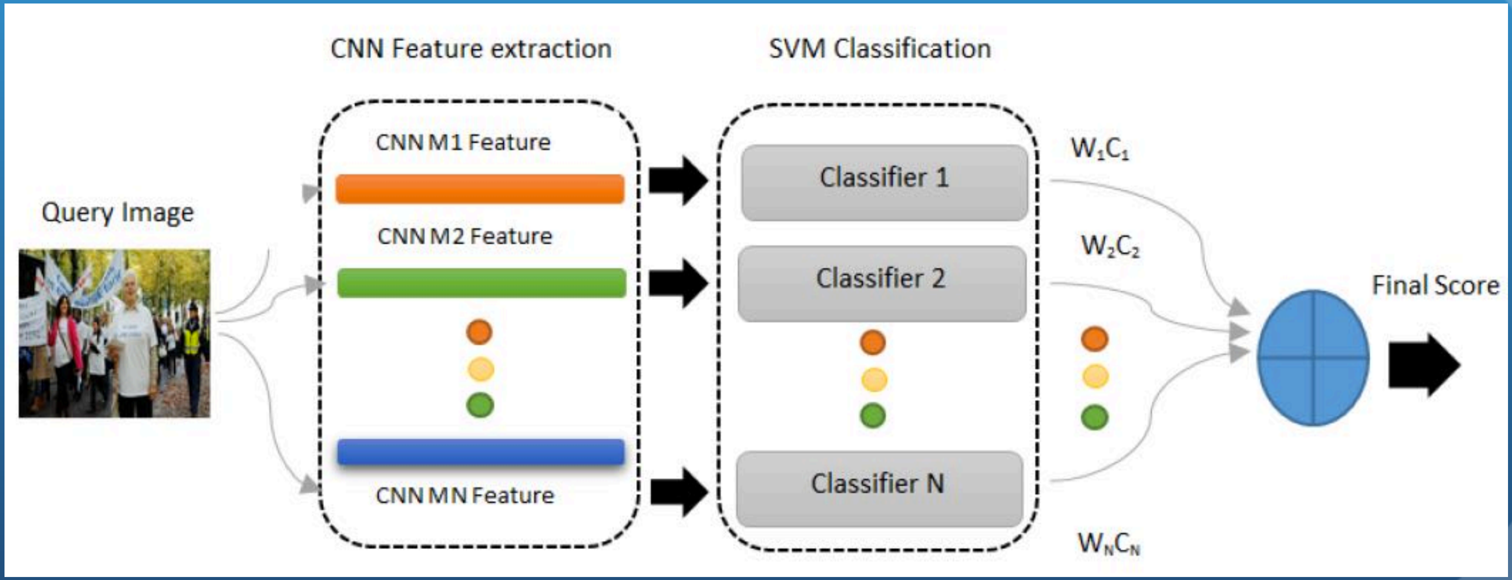

Using saliency for event discovery: ensamble of deep models

Sample images for which training on background fails (theater event class)

Sample images for which training on foreground fails (theater event class)

Using saliency for event discovery: ensamble of deep models

CNN Model	Avg. Acc.	CNN Model	Avg. Acc.
AlexNet (ImageNet)	0.422	AlexNet (Places Dataset)	0.4154
VGGNet16 (ImageNet)	0.477	VGGNet16 (Places Dataset)	0.454
VGGNet19 (ImageNet)	**0.479**	VGGNet19 (Places Dataset)	0.4682
GoogleNet (ImageNet)	0.448	ResNet 50 (ImageNet)	0.3010
ResNet 152 (ImageNet)	0.3008	ResNet 101 (ImageNet)	0.3006

From single network to ensamble

WIDER Dataset		UIUC Dataset	
Method	Avg. Acc.	Method	Avg. Acc.
Baseline [1]	0.397	Baseline [2]	0.734
Deep Channel Fusion [1]	0.424	Placess CNN [3]	0.941
Reza et al. [4]	0.440	GoogleNet GAP [4]	0.950
Zhou et al. [5]	0.530	Method in [5]	0.988
Our Approach (IOWA)	0.5840	Our Approach (IOWA)	0.9854
Our Approach (GA)	0.5826	Our Approach (GA)	0.9870
Our Approach (PSO)	**0.5908**	**Our Approach (PSO)**	**0.9887**
Our Approach (equal weights)	0.5593	Our Approach (equal weights)	0.9731

Structuring vs. Synchronization

- Another interesting problem is synchronization:
  - If there are multiple users taking part to the same event (possibly extended in time and space) how can we place their media into a single time scale?
    - Timestamp could be either unavailable or unreliable
    - Media could be captured from different perspectives, with different purposes, and with different devices
    - We could have images (single instant) and/or video (time span)
  - Also in this case, content can reveal a lot.

Structuring vs. Synchronization

Synchronization of Multi-User Event Media (SEM) Dataset Task in Mediaeval 2015

MediaEval Benchmark

N.Conci, V.Mezaris, F.De Natale, M.Matton, "Synchronization of multi-user event media at MediaEval (SEM): Task description, datasets, and evaluation", Proc. MediaEval, 2014-2015

Dataset: http://mmlab.disi.unitn.it/MediaEvalSEM2014

Automatic synchronization of multi-user photo collections

- Research question:
  - is it possible to find a common timeline for different sets of (possibly etherogeneous) media taken from a variety of users and devices, using content + metadata?

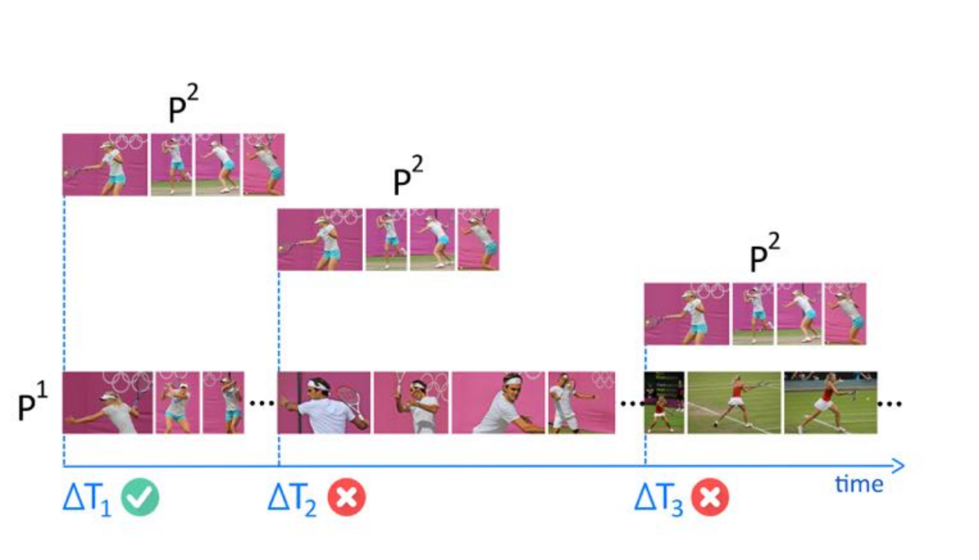

E.Sansone, K.Apostolidis, N.Conci, G.Boato, V.Mezaris, F.De Natale, "Automatic Synchronization of Multi-user Photo Galleries", IEEE Trans. on Multimedia, 2017

Automatic synchronization of multi-user photo collections

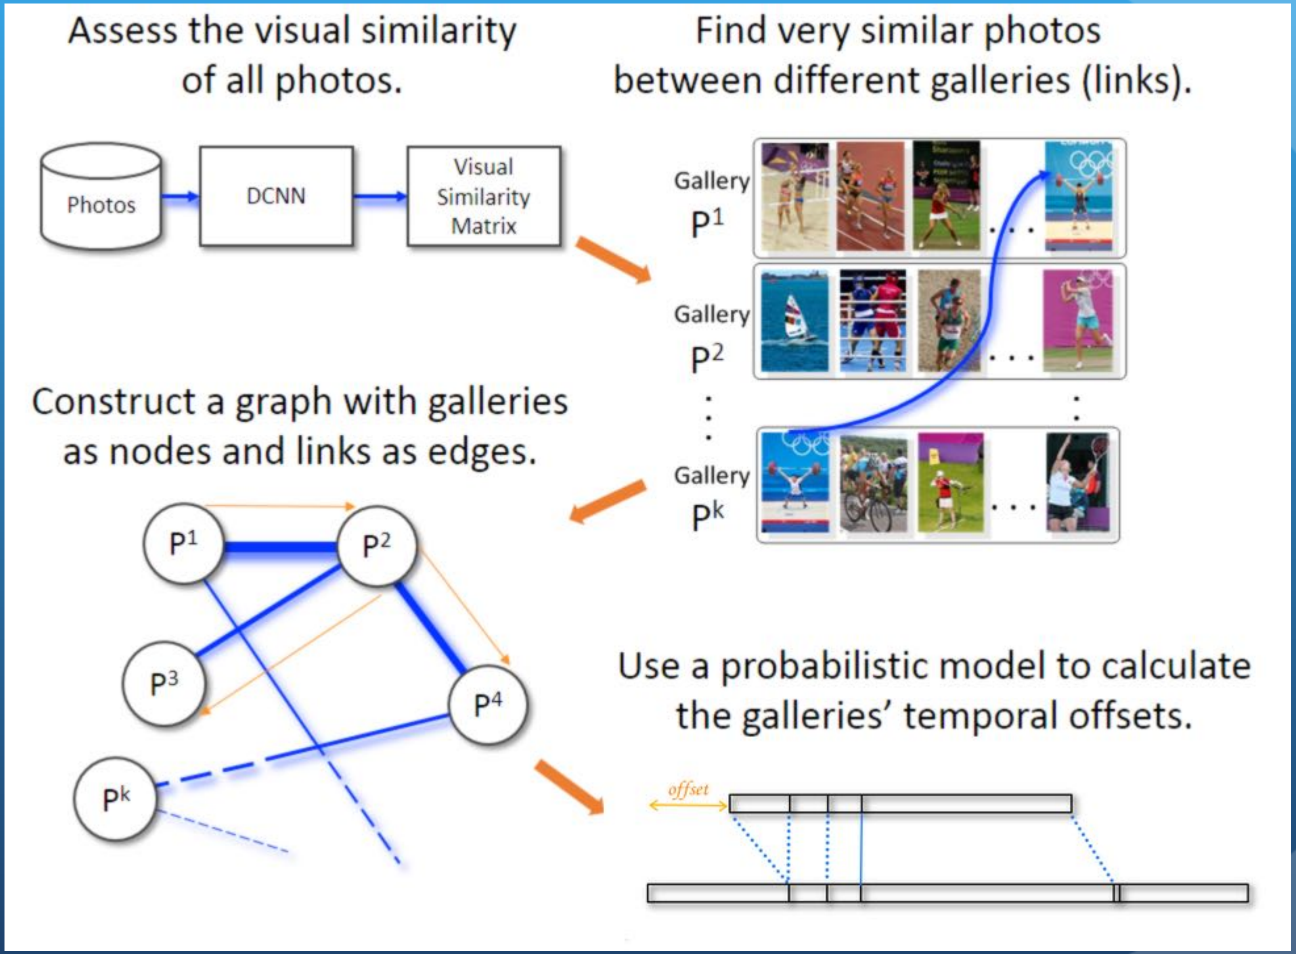

Automatic synchronization of multi-user photo collections

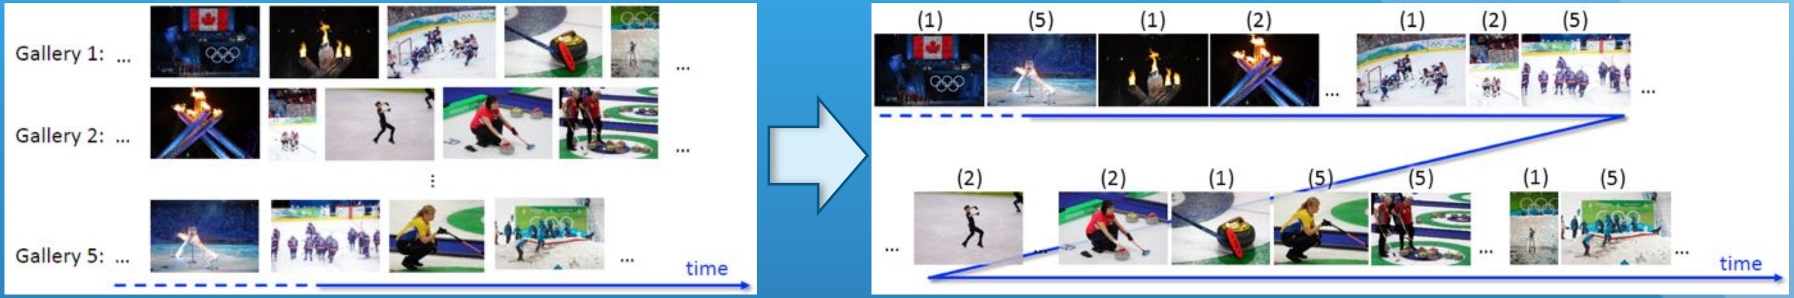

Method	Vancouver		London		NAMM15		TDF14		Average and standard deviation across all datasets		
	P(%)	A(%)	P(%)	A(%)	P(%)	A(%)	P(%)	A(%)	P(%)	A(%)	H(%)
[24]	61.8	81.6	33.3	88.6	50.0*	73.5*	25.0	80.6	43.9 ± 17.6	81.1 ± 6.2	55.5 ± 15.9
[31]	94.1	79.2	47.2	87.5	-	-	-	-	70.7 ± 33.2	83.4 ± 5.9	73.6 ± 17.2
[34]	5.0	65.0	15.0	**92.0**	-	-	-	-	10.0 ± 7.1	78.5 ± 19.1	17.6 ± 11.6
[33]	91.2	72.8	61.1	71.3	77.9*	90.4*	9.4	79.3	61.4 ± 37.0	78.5 ± 8.7	62.7 ± 31.8
[35]	-	-	-	-	40.0	78.0	5.0	**92.0**	22.5 ± 24.5	85.0 ± 9.9	31.2 ± 30.6
[37]	94.1	76.0	61.1	65.6	83.3	90.8	12.5	84.5	62.8 ± 36.2	79.2 ± 10.9	64.0 ± 30.1
[36]	35.0	86.0	25.0	89.0	44.5*	82.4*	15.6	83.2	30.0 ± 12.5	**85.2** ± 3.0	43.2 ± 13.7
[38]	**97.1**	86.0	63.9	75.0	50.0*	71.5*	21.9	90.8	59.6 ± 30.9	80.8 ± 9.1	64.5 ± 23.0
[36] modified (inception3a)	5.9	**93.4**	5.6	73.5	47.1	85.0	9.4	73.2	17.0 ± 20.1	81.3 ± 9.8	24.7 ± 24.1
[38] modified (inception3a,*exact*)	85.3	56.3	47.2	74.6	**88.9***	88.7*	**65.6**	84.6	73.1 ± 20.9	76.1 ± 14.4	72.7 ± 14.1
[38] and [36]	8.8	54.7	19.4	67.4	58.8	66.8	25.0	80.6	28.0 ± 21.6	67.4 ± 10.6	36.5 ± 19.8
Proposed (inception3a,*exact*,MRF)	**97.1**	83.7	**75.0**	84.3	83.3*	**93.4***	**65.6**	73.7	**80.3** ± 13.3	83.8 ± 8.0	**81.7** ± 9.4

Structuring vs. Diversity

B.Ionescu, A.Popescu, M.Lupu, H.Muller, "Retrieving Diverse Social Images", Proc. MediaEval, from 2013 to 2107 (5 editions)

Winner of 2014 Edition

D-T.Dang-Nguyen, G.Boato, F.De Natale, G.Giacinto, L.Piras, "Retrieval of Diverse Images by Pre-filtering and Herarchical Clustering", Proc. MediaEval, 2014

Achieving diversity with RF

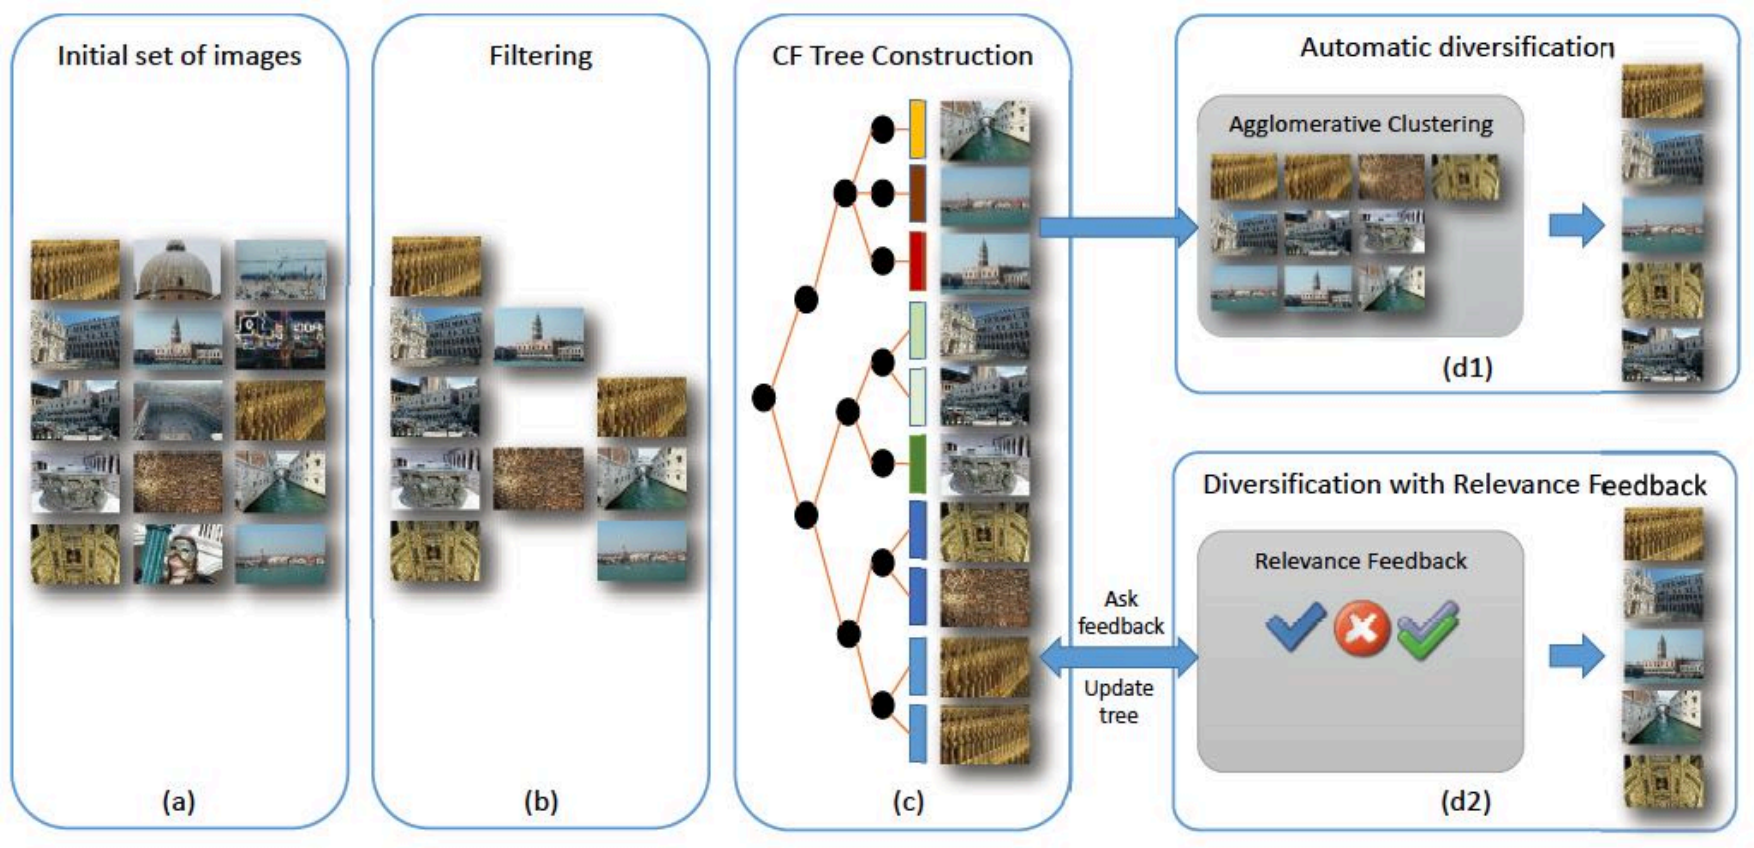

D-T.Dang-Nguyen, L-Piras, G.Giacinto, G.Boato, F.De Natale, "Multimodal retrieval with diversification and relevance feedback for tourist attraction images", ACM TOMM, 2017

Diversity by RF: Some results

(a) Results from Flickr. (b) Results without relevance feedback. (c) Results with relevance feedback.

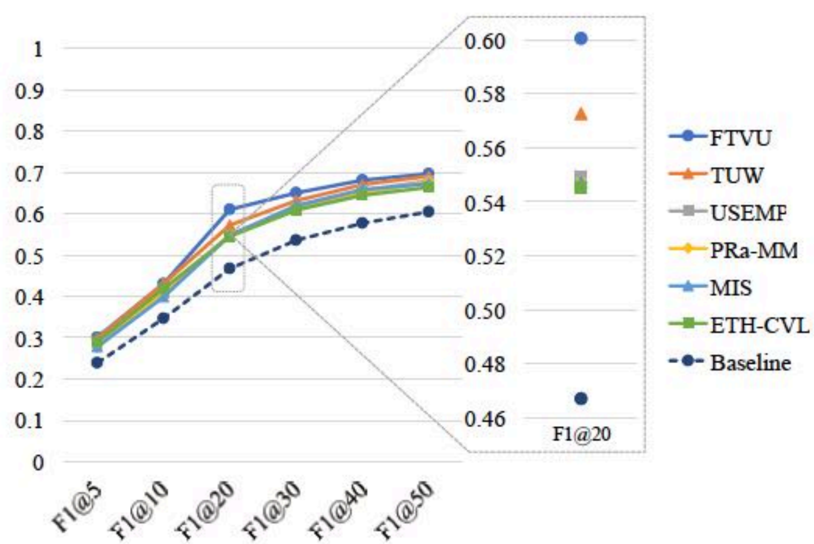

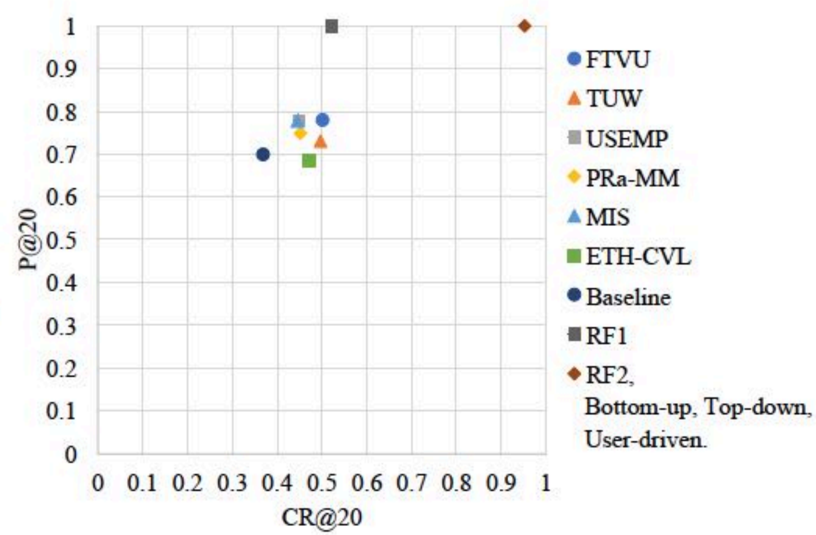

Outline

- Why social media are a good example of ‘big data’
- Challenges of social media:
  - Making structured the unstructured
    - Event-based media structuring
    - Media Synchronization
    - Media Diversity
  - Making trustworthy the untrustworthy
    - Source and provenance detection
    - Manipulation detection and biasing
- Open problems in big social media research

Seeing is not believing

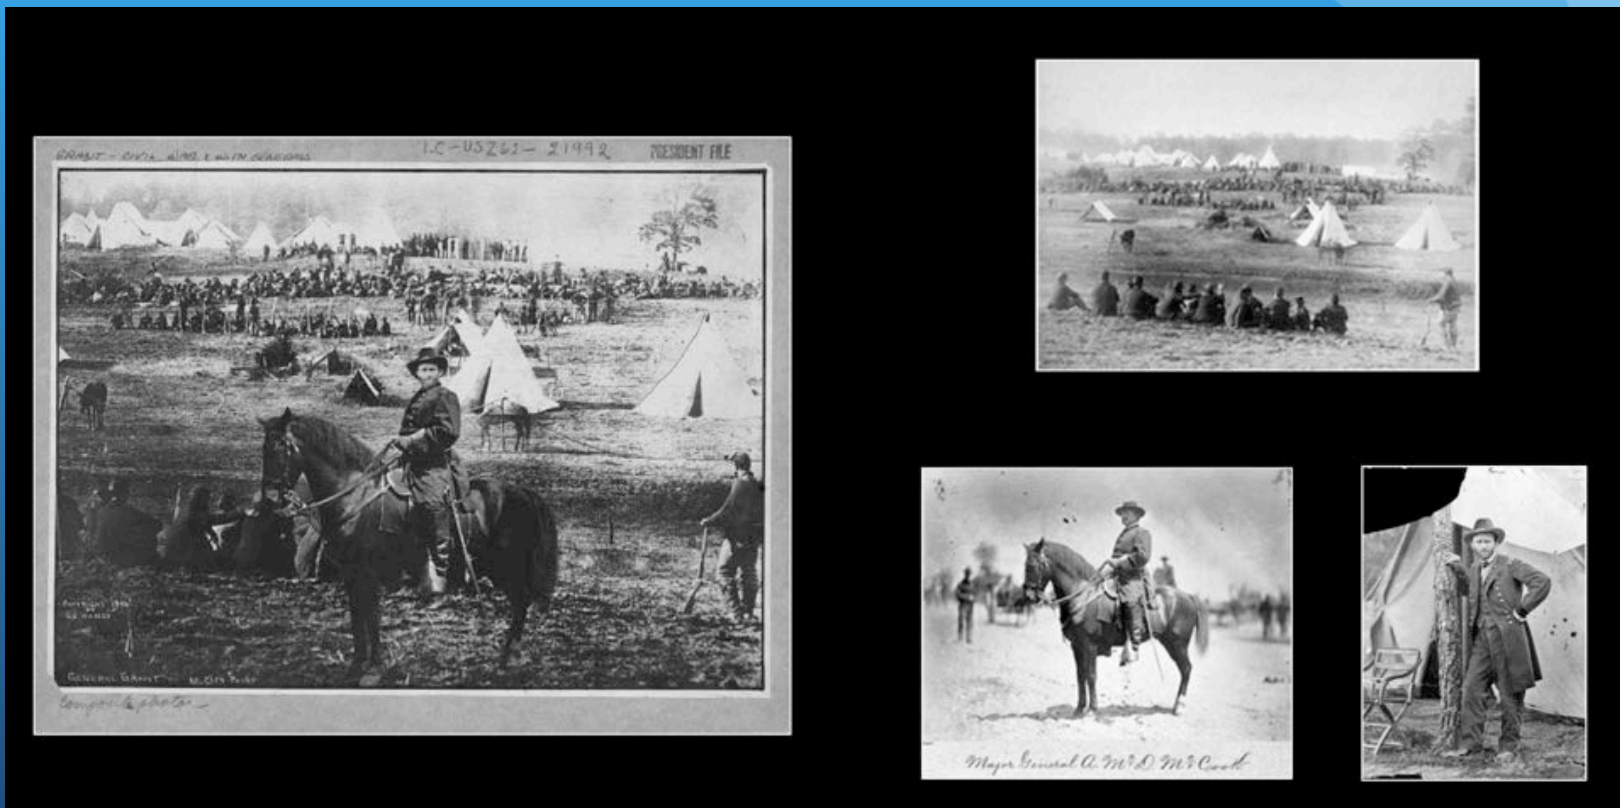

Fake news in the 18th century

Source: ELT-CY4GATE Workshop

Seeing is not believing

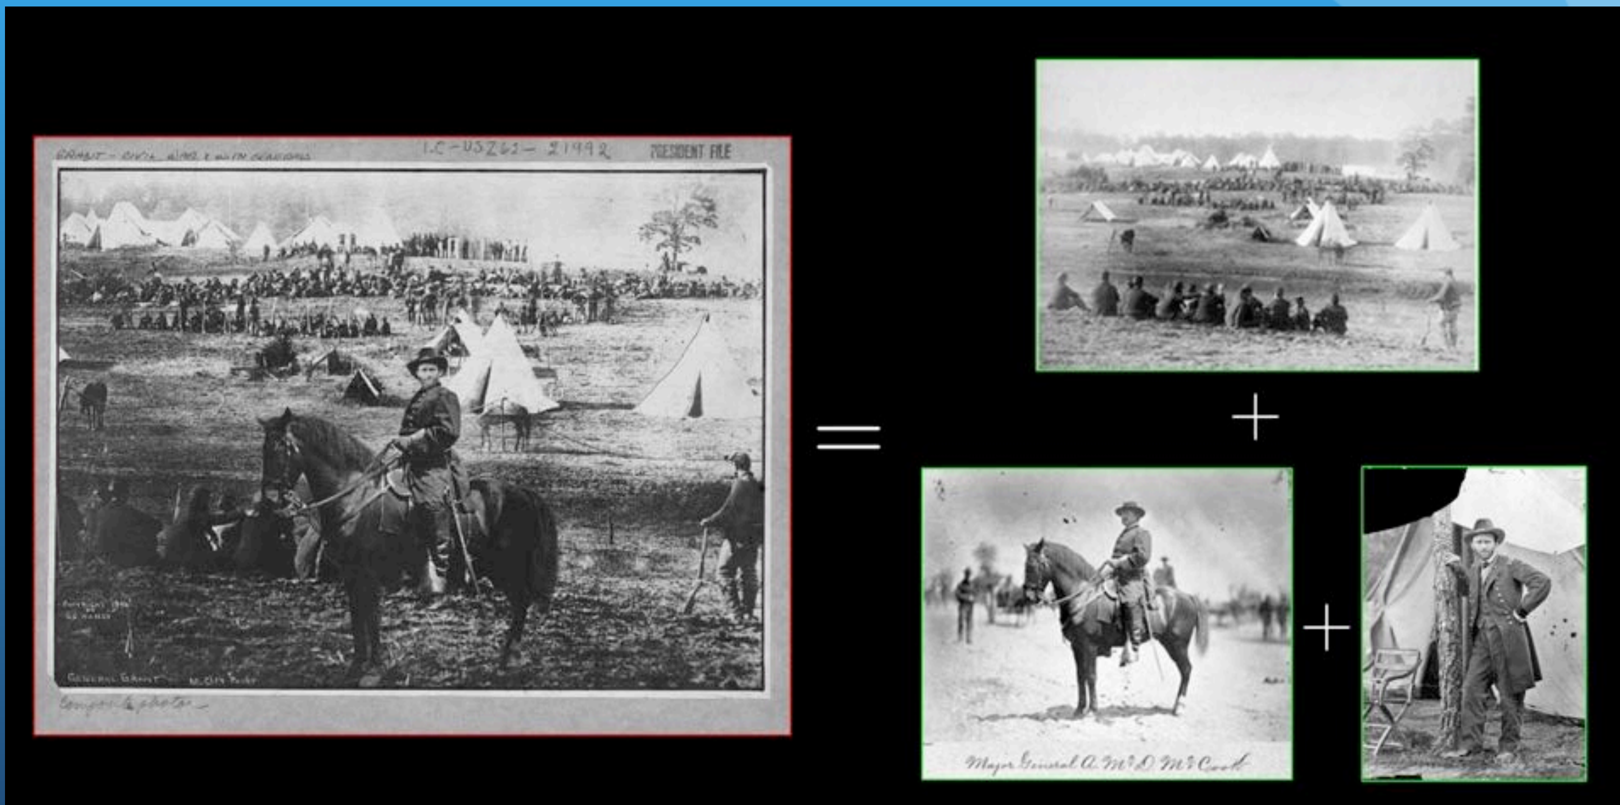

Fake news in the 18th century

Source: ELT-CY4GATE Workshop

Seeing is not believing

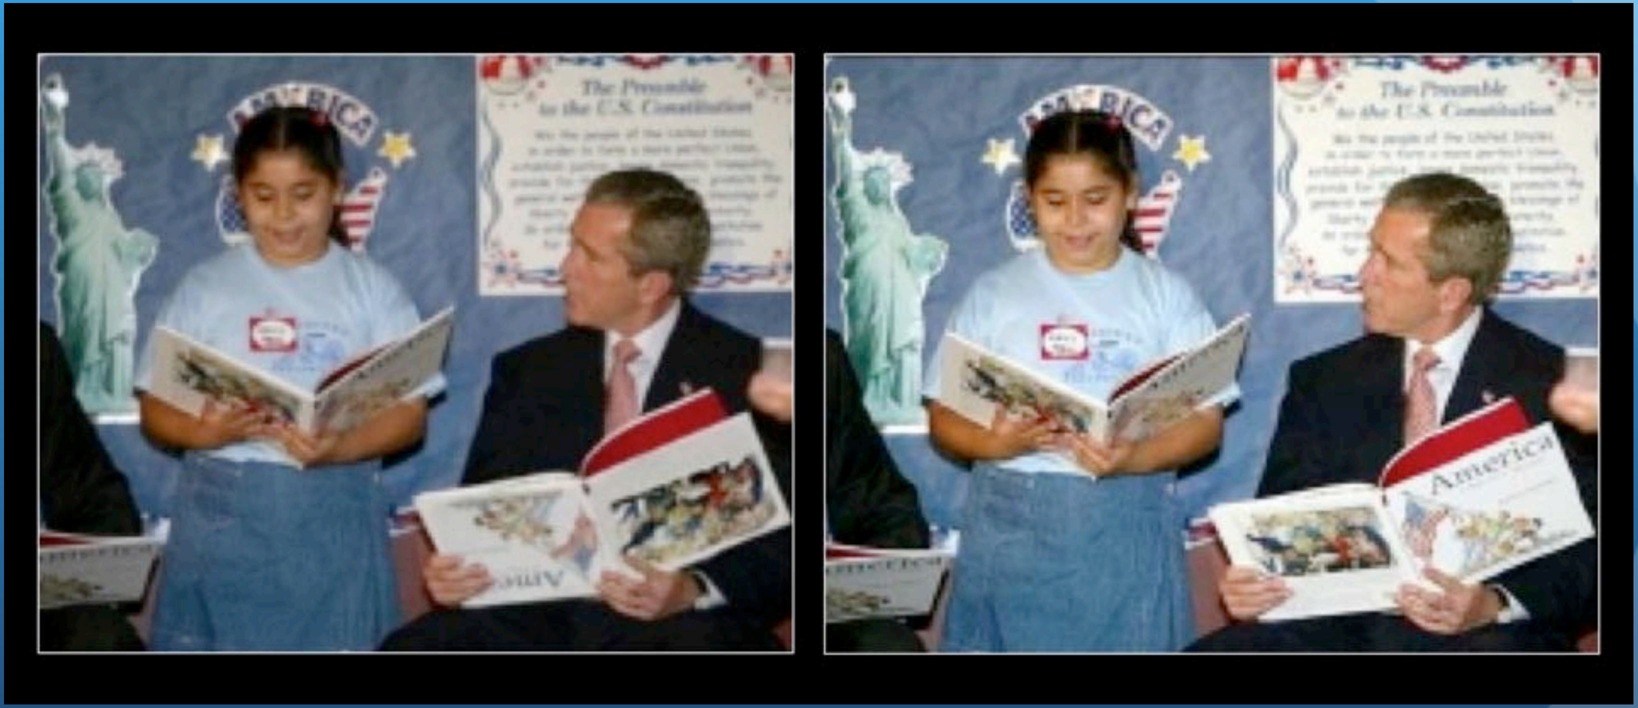

Fake news today

Source: ELT-CY4GATE Workshop

Seeing is not believing

Fake news today

Source: ELT-CY4GATE Workshop

Media impact on perception

- **Modified data may influence people opinions** and even alter their attitudes in response to the represented event

How is your feeling if I say that this image reports the status of a polluted area?

Media impact on perception

- **Modified data may influence people opinions** and even alter their attitudes in response to the represented event

and now?...

- it is important to be able to automatically verify the **authenticity** and the **integrity** of digital images in order to guarantee their trustworthiness.

Media impact on perception

Social media do the rest of the job: fast and large scale propagation (viral media) of modified multimedia content, potentially carrying a distorted semantic message

Multimedia forensics

- We refer to multimedia forensics as the set of methodologies that allow revealing information about the “history” of a media item from the instant of capturing to its current publication
  - Identifying the device (and possibly the person) that captured the media
  - Determining possible tampering (from simple image processing to complex and multiple manipulation)
  - Discriminating between real and computer generated media (or parts of)

Multimedia forensics is not a wizard

Source identification

- Acquisition devices leave unique traces (fingerprints) on the media, such as guns do it on bullets

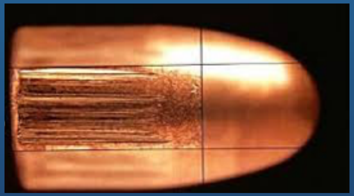

Source identification

- Source identification uses such traces to link media contents to a particular (class of) acquisition device(s)
- The problem can be solved at different levels:
  - Identifying the class of the device (e.g., scanner-vs-camera)
  - Identifying the specific device (e.g., Canon-vs-Nikon)
  - Identifying the brand/model of the device (e.g., SN_X-vs-SN_Y)

Source traces

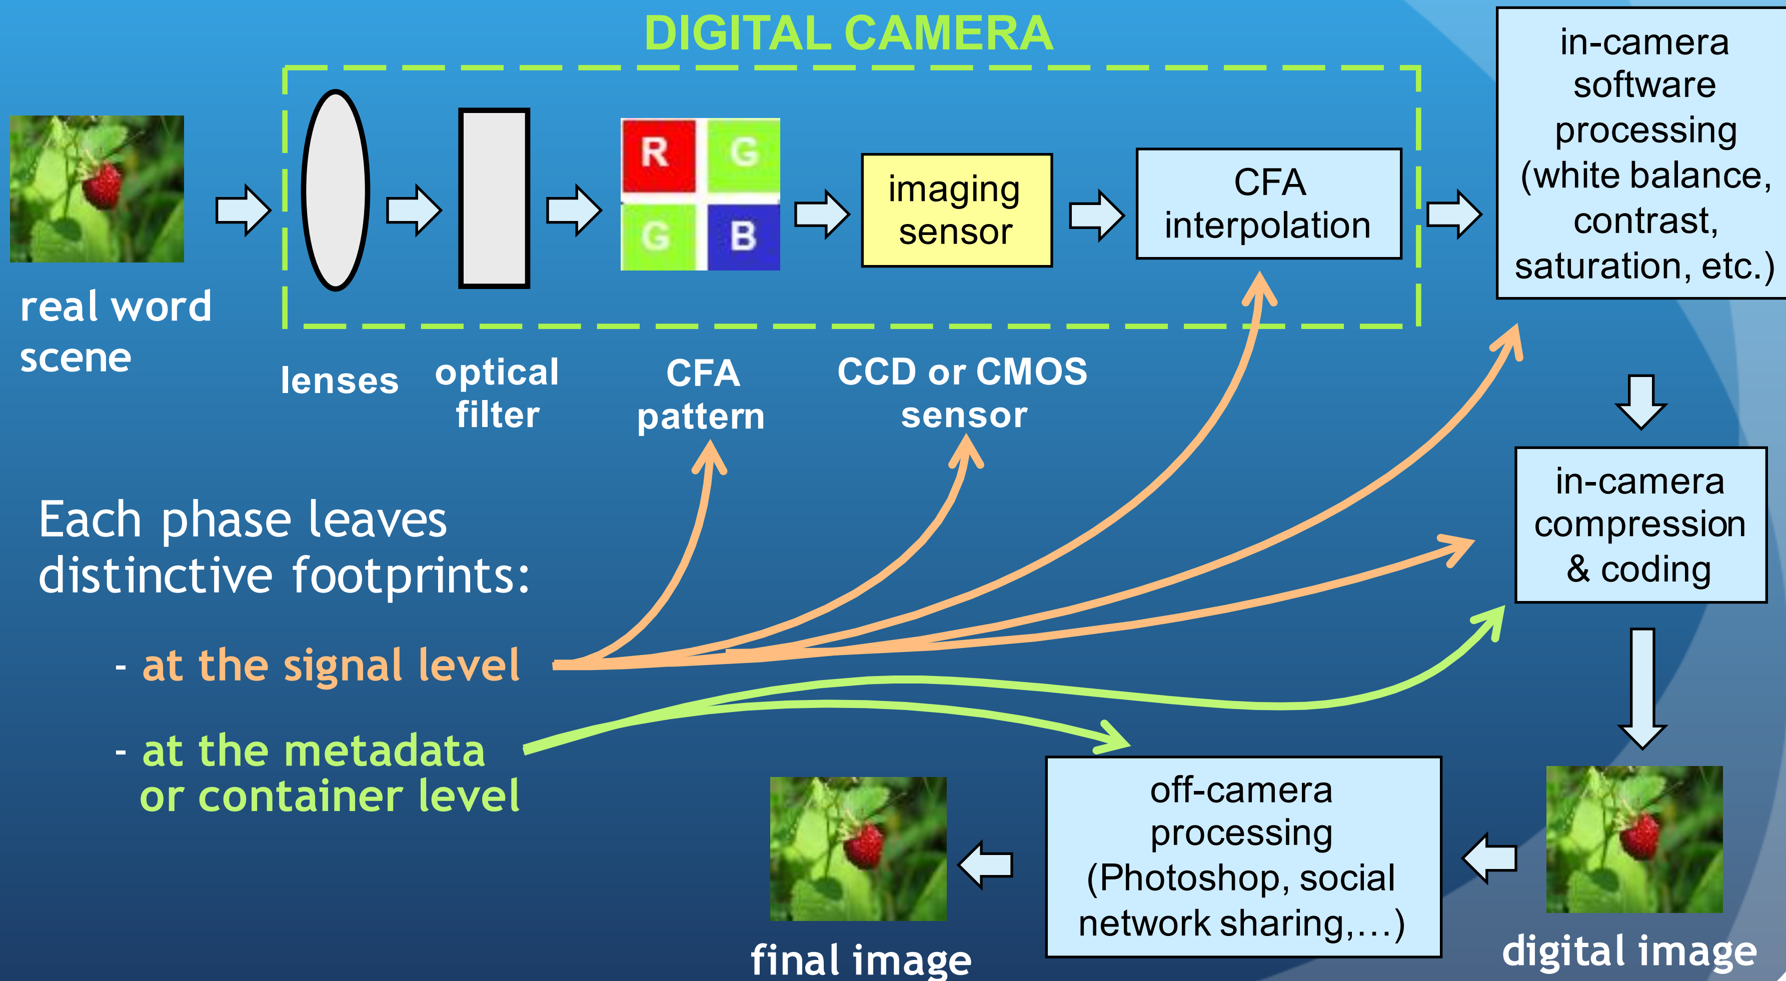

Source traces: PRNU

- **Photo Response Non Uniformity Noise** (PRNU) is typical of a CCD/CMOS sensor

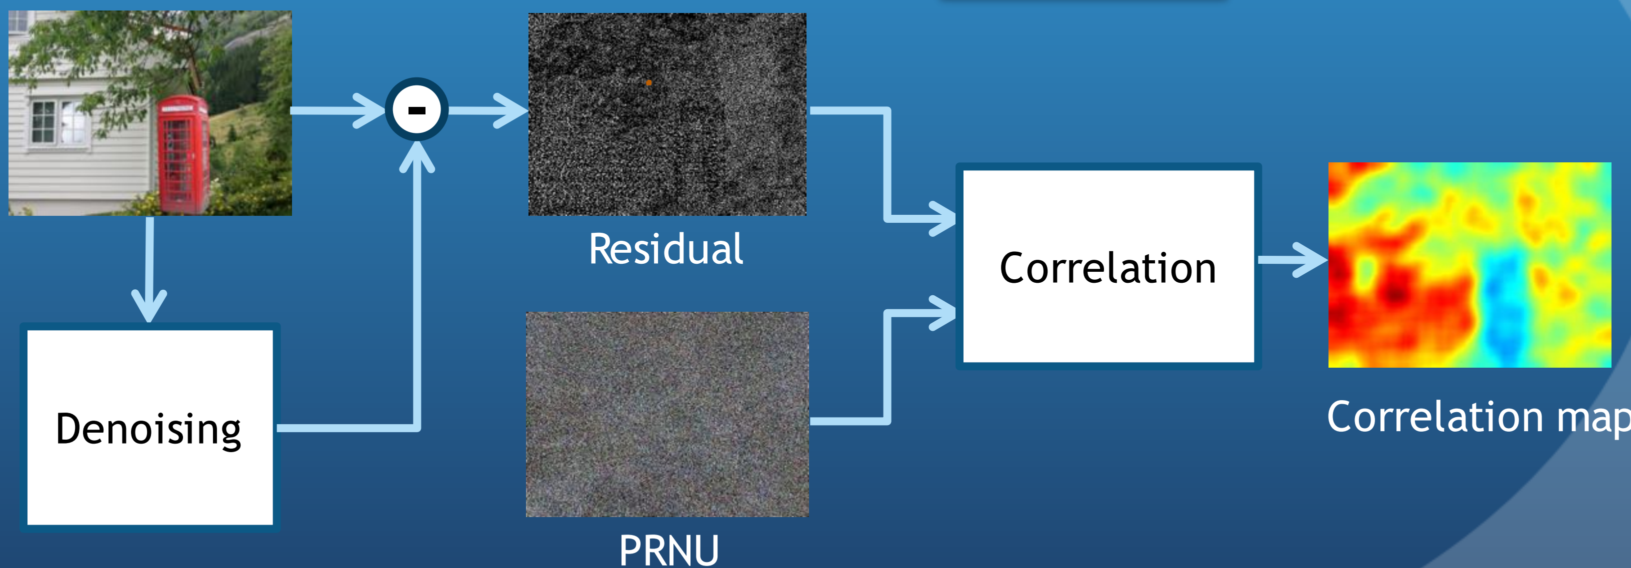

M.Chen, J.Fridrich, M.Goljan, J.Lukas, "Determining Image Origin and Integrity Using Sensor Noise," in IEEE Trans. on Information Forensics and Security, 2008

Example 1: detecting photo owner on Social Networks

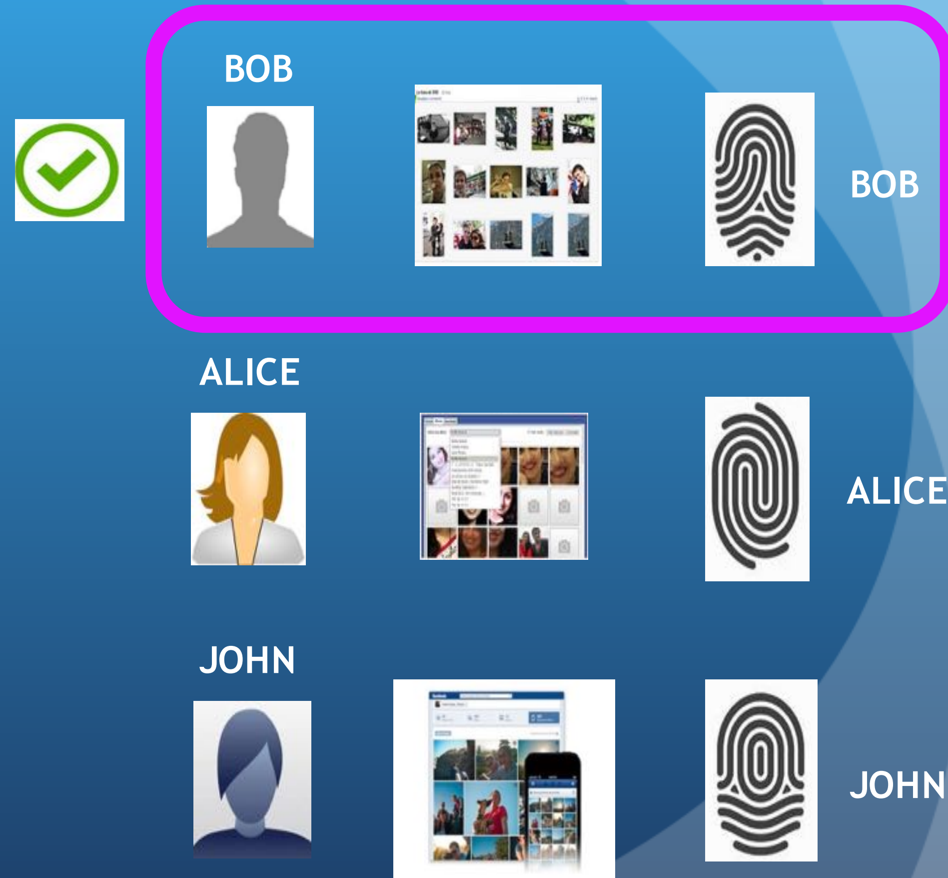

Who is the author of this fake?

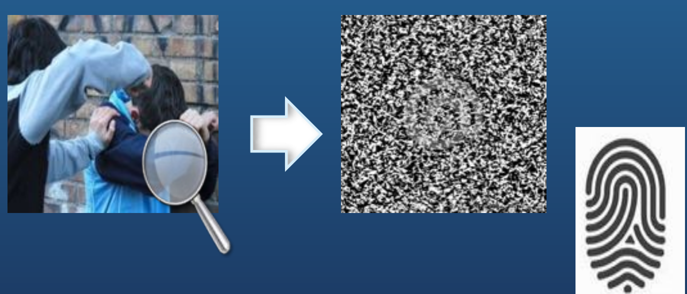

Example 2: Media clustering by source camera

Canon Alpha SX710 HS   Panasonic DMC-LX100S Lumix   Sony Alpha 3000   Kodak PIXPRO AZ251

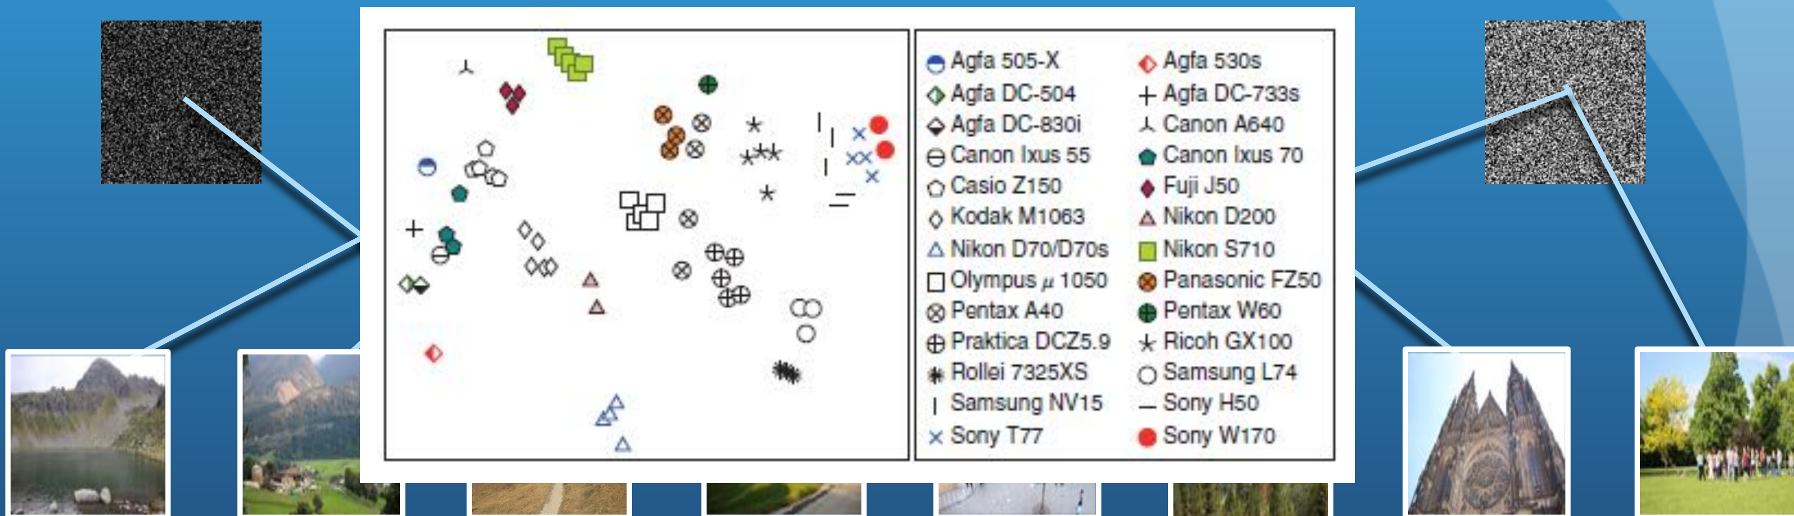

Q-T.Phan, G.Boato, F.De Natale “Image Clustering by Source Camera via Sparse Representation”, MFSec 2017

Q-T.Phan, G.Boato, F.De Natale, “Robust Image Clustering based on Sparse Representation of Camera Fingerprint”, submitted to IEEE TIFS

Also social media leave traces

- Multimedia data phylogeny
  - Trace the history of the data (source, manipulations, sharing)

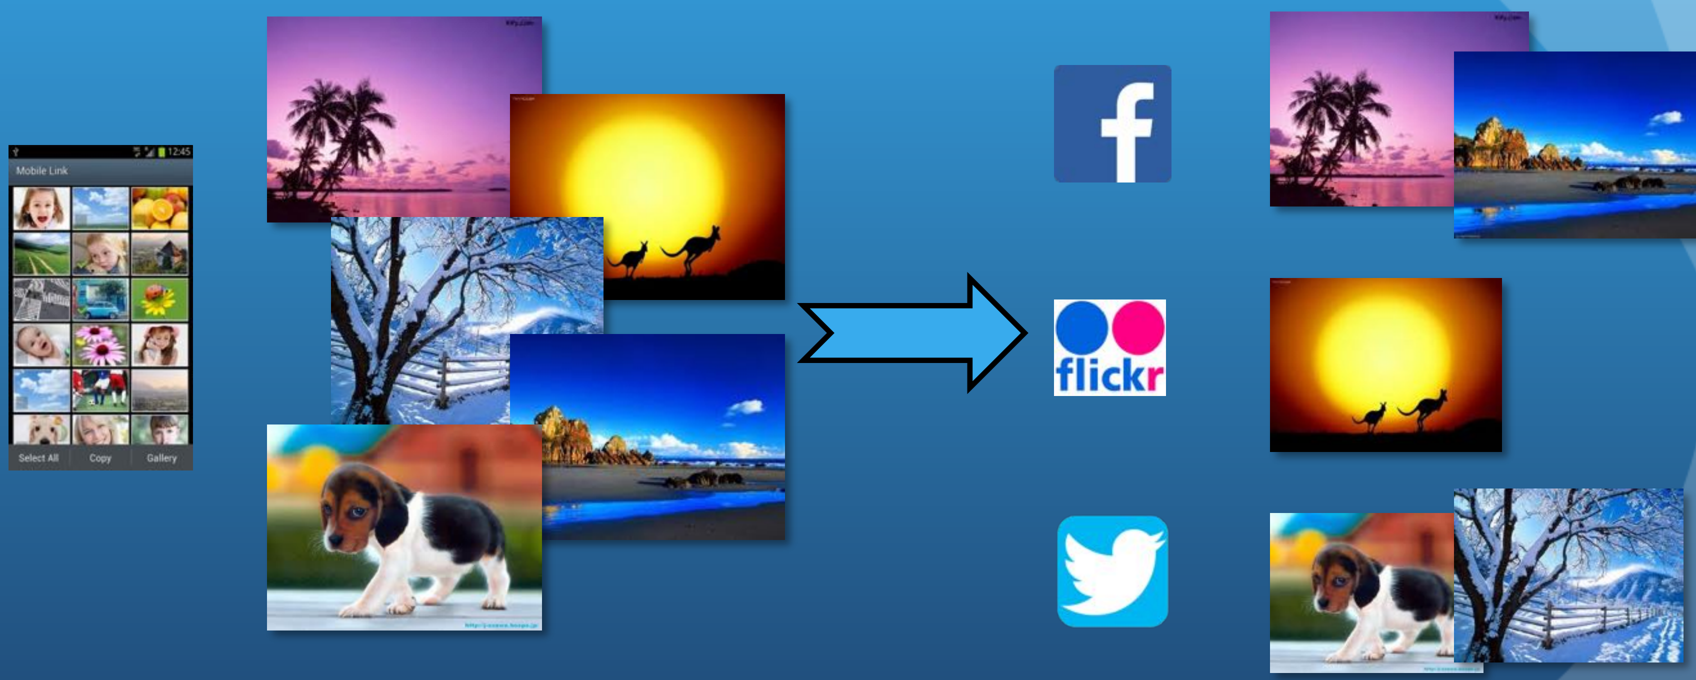

R. Caldelli et al., Image Origin Classification Based on Social Network Provenance, IEEE Trans. on Information Forensics and Security, 2017
I. Amerini et al., Tracing images back to their social network of origin: a CNN-based approach, IEEE WIFS, 2017

Example 1: detecting social media provenance

- Uploading an image on a Social Network: the process alters images
  - Resize
  - Rename
  - Meta-Data deletion/editing
  - Re-Compression
  - NEW JPEG file Structure

E.g., histogram of DCT coefficients allows recognizing different social networks

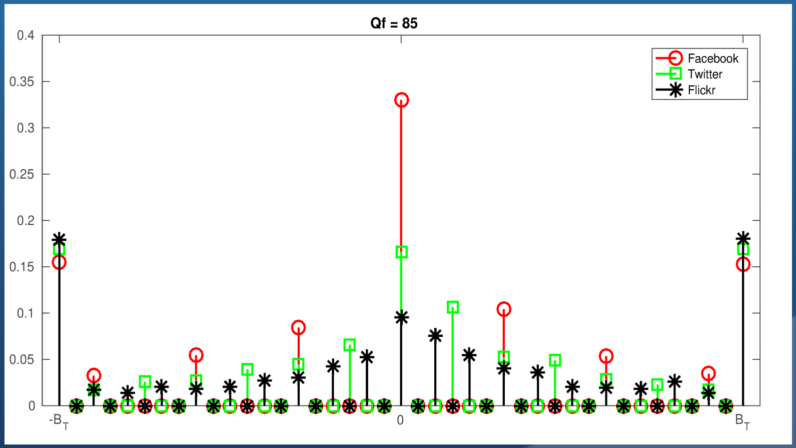

Example 2: detecting instant messaging app provenance

- Considered messaging apps: Messenger, Whatsapp, Telegram
- Two mobile phones of different operating systems: Android and iOS
- Original images are taken from VISION dataset: various resolutions and qualities.
- Features:
  - Histogram of DCT coefficients
  - Additional information from JPEG headers
- Extension to double sharing scenarios.

Q-T.Phan, C.Pasquini, G.Boato, F.De Natale "Identifying image provenance: an analysis of mobile instant messaging apps", IEEE MMSP 2018

Example 2: detecting instant messaging app provenance

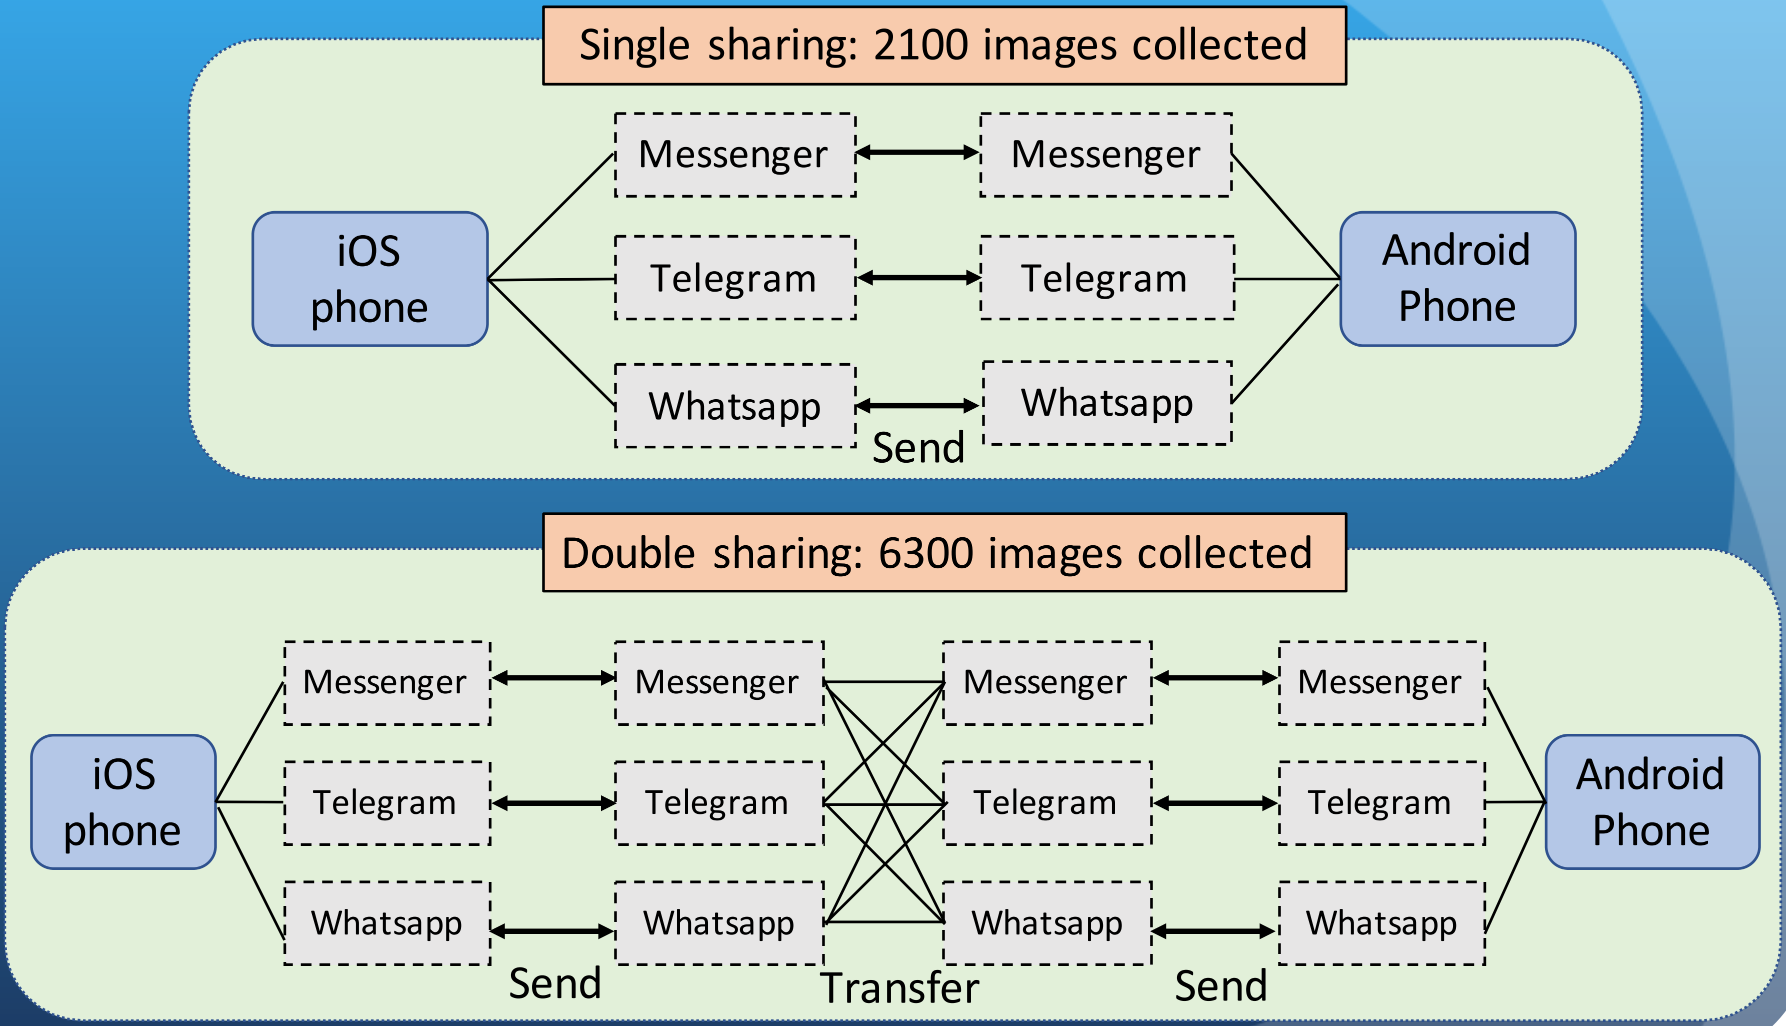

Image provenance detection: mobile instant messaging apps

- APP and OS / LAST APP and LAST OS can be reliably identified in single / double sharing scenarios.
- The PREVIOUS APP can be reliably revealed based on the information of the LAST APP or LAST APP + LAST OS.

	Scenario	Accuracy
APP	Single sharing	100
APP + OS		100
LAST APP	Double sharing	94
LAST APP + LAST OS		89
PREVIOUS APP | LAST APP		77
PREVIOUS APP | LAST APP + LAST OS		80

Example 3: Out of context use of media

- Verifying not only images but also their context
  - Image is original, but its use is not consistent

Real photo
captured April 2011 by WSJ
BUT
heavily tweeted during Hurricane Sandy
(29 Oct 2012)

Verification task @ MediaEval

- Given a post (image+metadata), return a decision (fake, real, unknown) on whether the information presented by the post reflects the reality
  - Photos from past events reposted as being associated to current event
  - Digitally manipulated photos
  - Artworks presented as real imagery

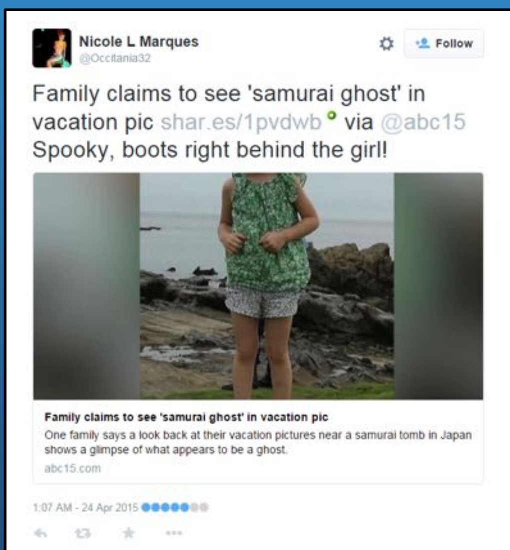

C.Boididu, S.Papadopoulos, S.Middleton, G.Boato, D-T.Dang-Nguyen, M.Riegler, "Task: Verifying Multimedia Use", Proc. MediaEval, 2016
C.Lago, Q-T.Phan, G.Boato, "Image forensics in online news", IEEE MSSP 2018

Tampering detection

Tampering detection

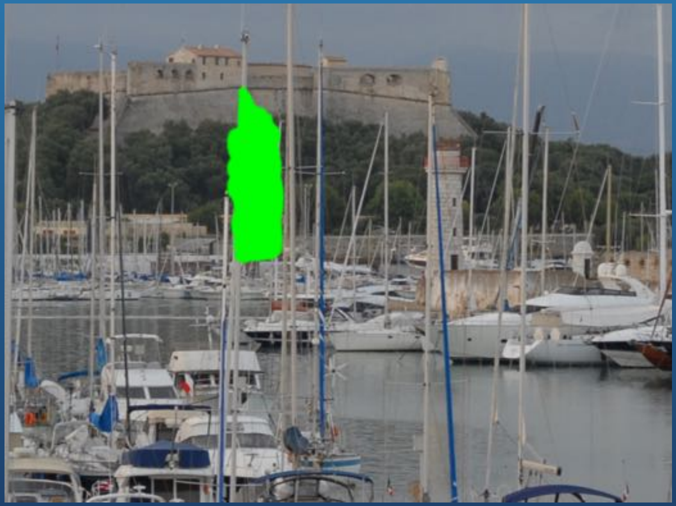

Tampering detection




Tampering detection

- Manipulation **detection**

Forgery Detector

Yes/No

It's forged !

- Manipulation **localization**

Forgery Localizer

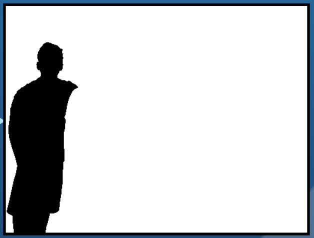

Forgery map

Tampering detection

- Pixel-level traces
  - CFA, PRNU, etc. consistency
  - Filtering detection (linear, non linear, contrast, color, …)
- Scene-level traces
  - Physical and geometric distortion (perspective coherence)
  - Light direction and shadows coherence
- Format-level traces
  - Multiple JPEG compression
  - Other encoding
- Semantic-level traces
  - Date, place, time coherence

Pixel level traces

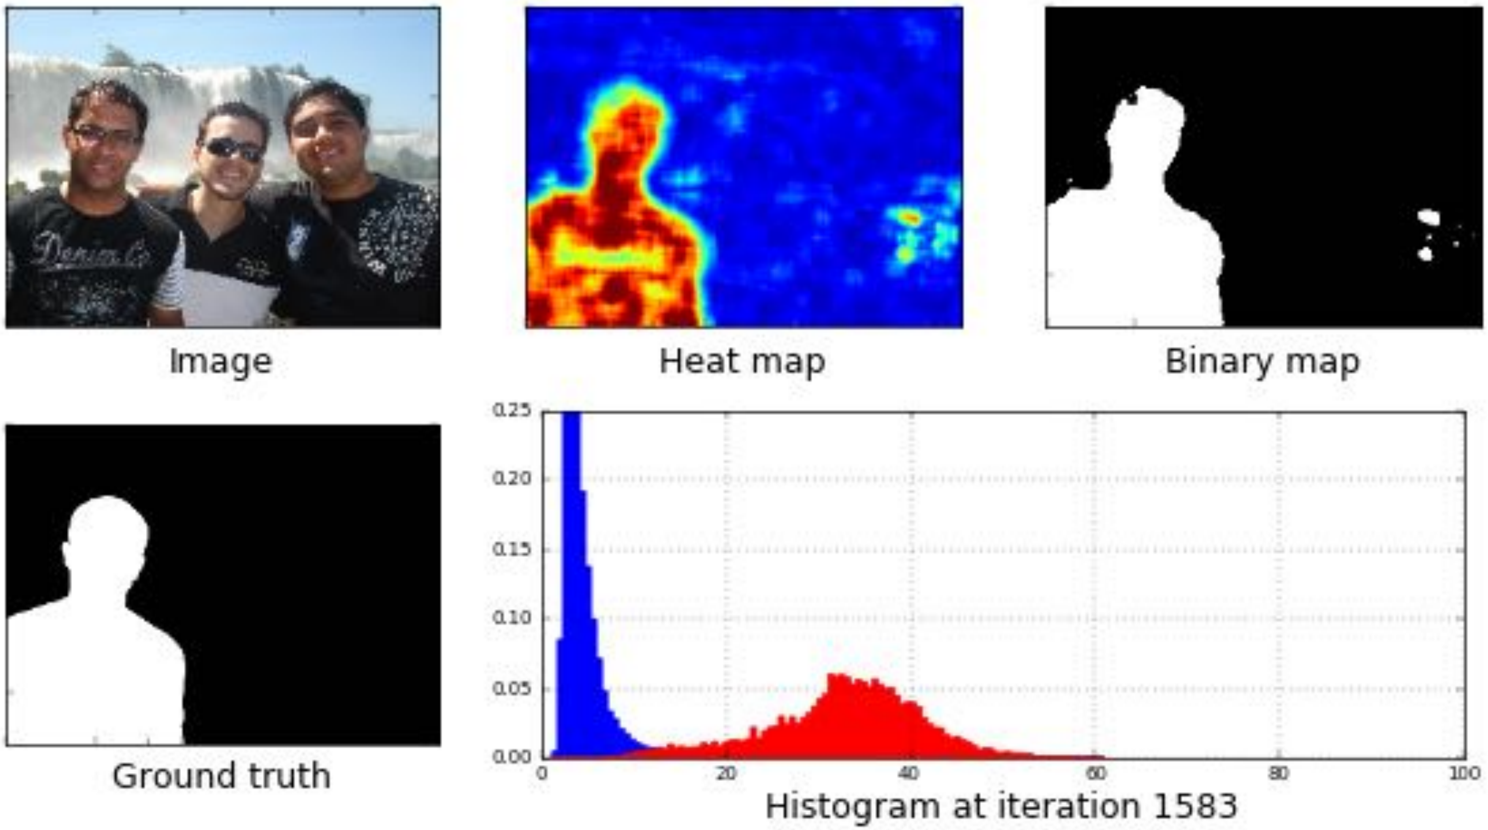

D.Cozzolino, G.Poggi, L.Verdoliva, "Recasting residual-based local descriptors as convolutional neural networks: an application to image forgery detection", IH&MMSec 2017

Filtering traces

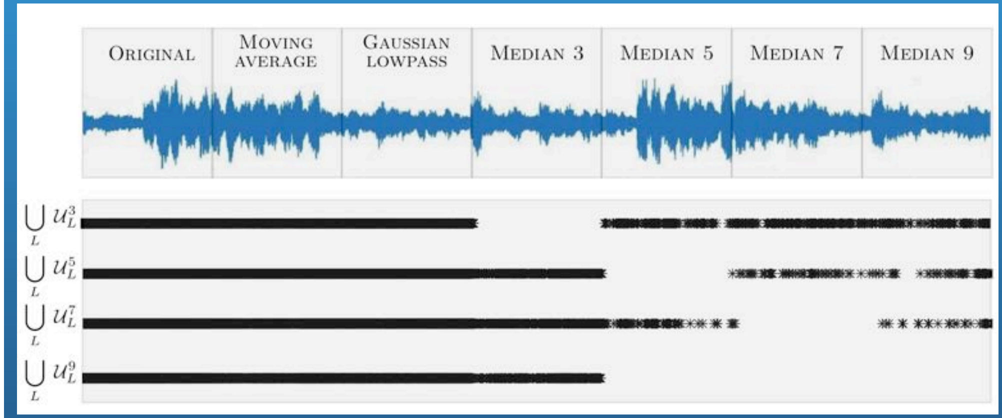

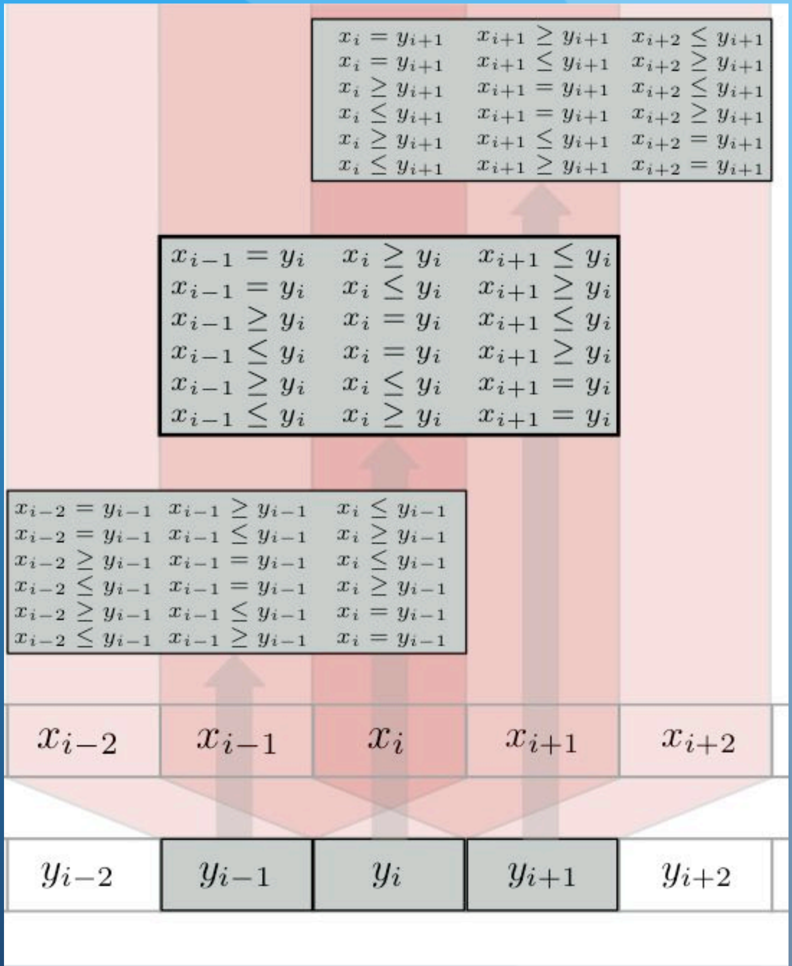

C.Pasquini, G.Boato, N.Alajlan, F.DeNatale, "A deterministic approach to detect median filtering in 1D data", IEEE Trans. On Information Forensics and Security, 2016

Filtering traces

F.DeNatale, G.Boato, "Detecting morphological filtering of binary images", IEEE Trans. On Information Forensics and Security, 2017

Light consistency traces

- When composing images, it is difficult to match lighting effects under which each image was taken
- Inconsistencies in lighting can be used to detect splicing

M.Johnson, H.Farid, *“Exposing Digital Forgeries in Complex Lighting Environments”*, IEEE Transactions on Information Forensics and Security, 2007

Perspective consistency traces

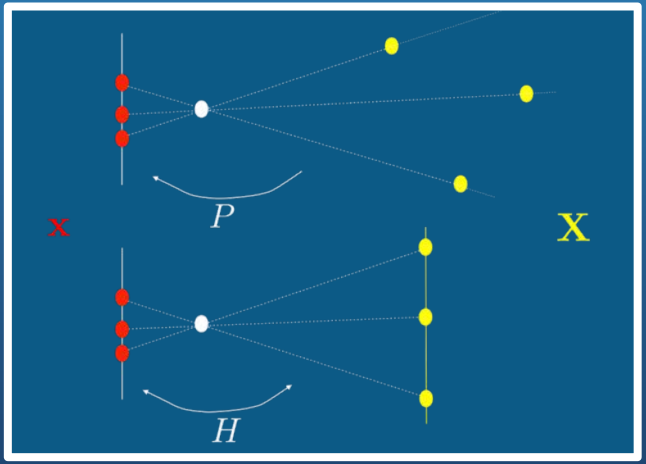

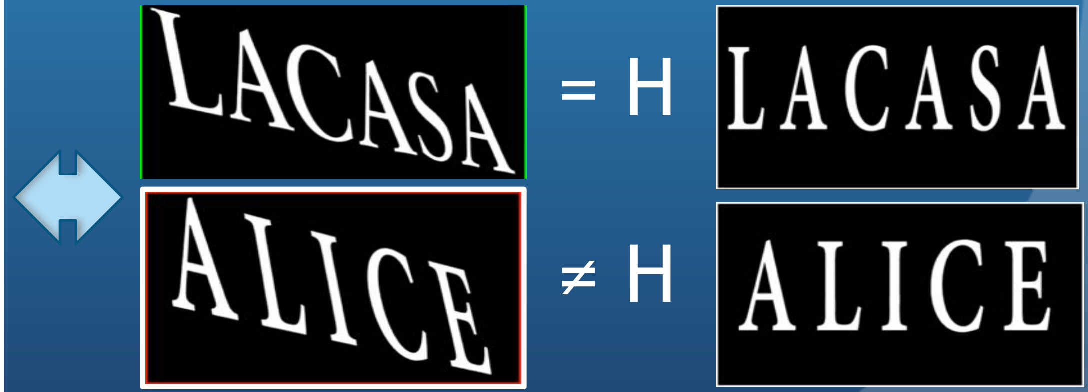

V.Conotter, G.Boato, H.Farid, *"Detecting photo manipulation on signs and billboards", IEEE ICIP, 2010*

Trajectory consistency traces

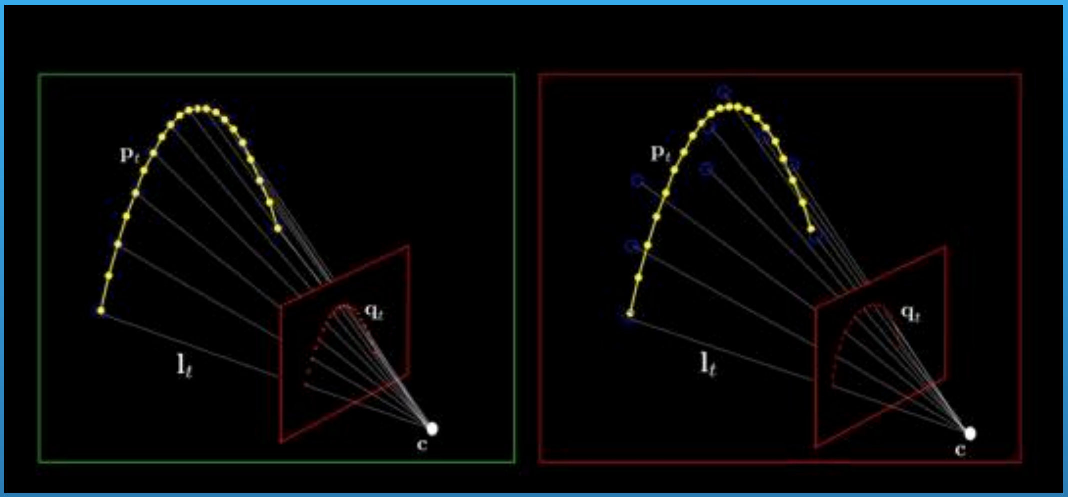

- Geometric techniques can detect physically **implausible trajectories** of objects in video sequences

V.Conotter, J.O'Brien, H.Farid, "Esposing digital forgeries in ballistic motion", IEEE Trans. on Information Forensics and Security, 2012

Compression traces

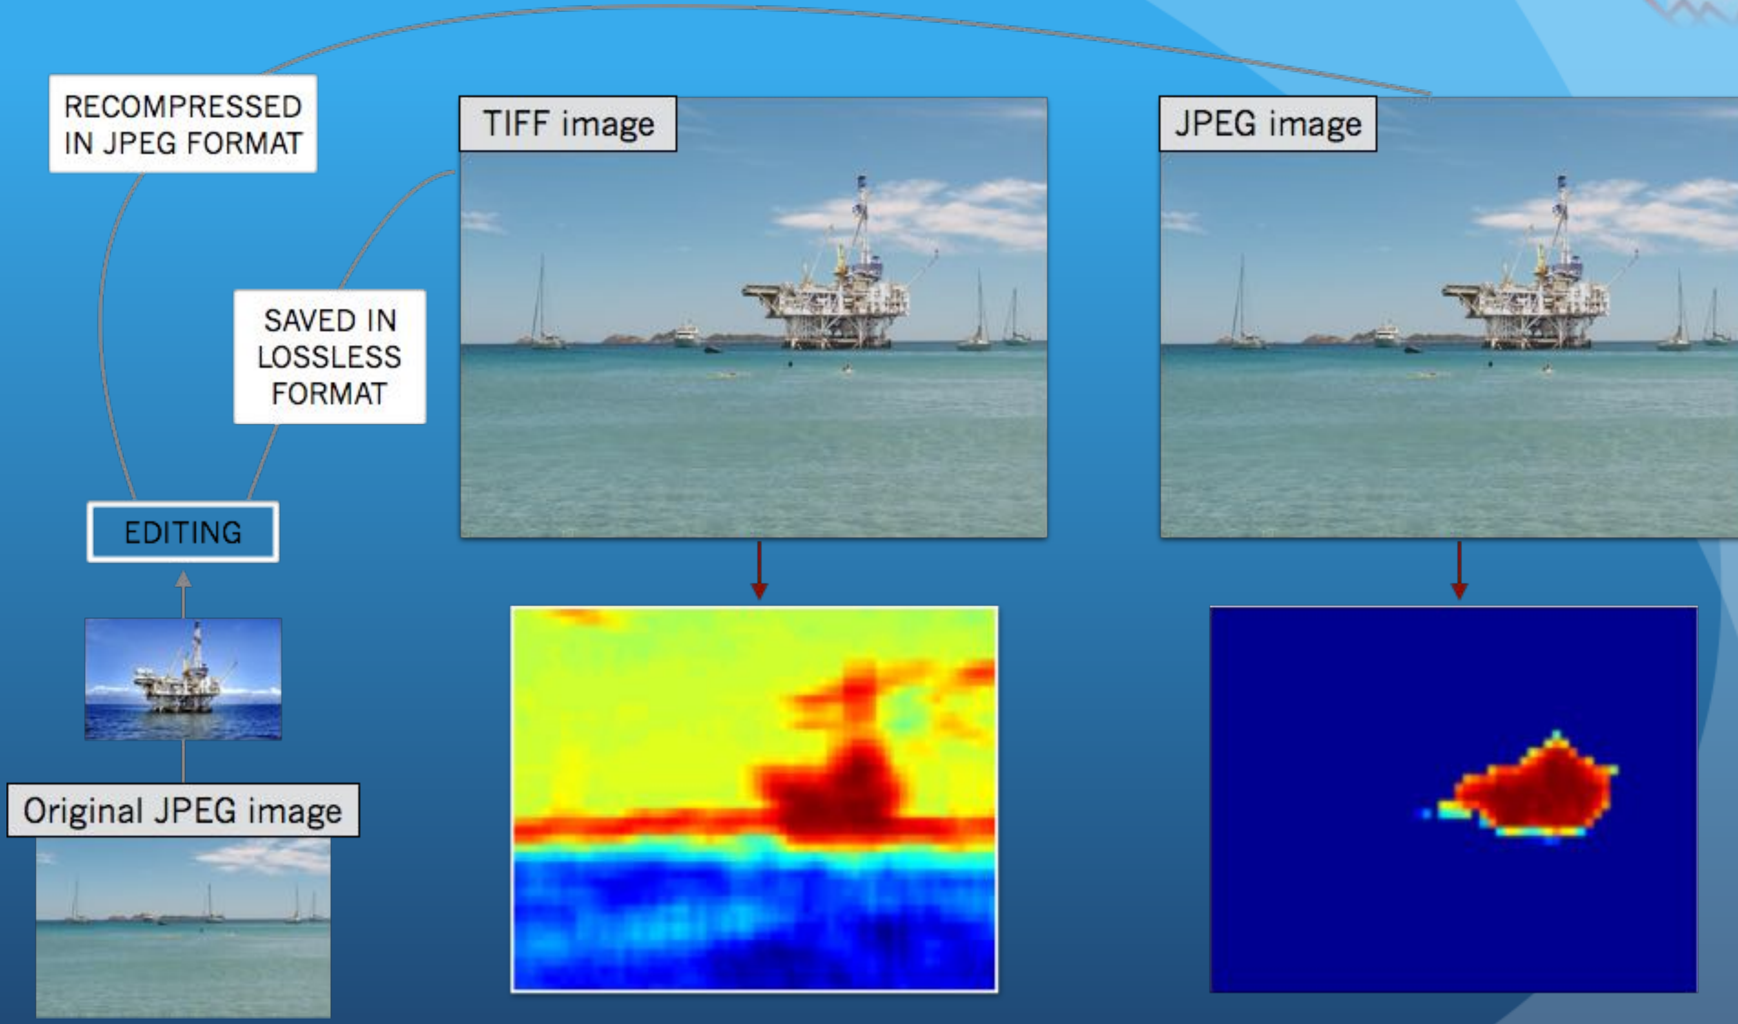

C.Pasquini, G.Boato, F.Pèrez-Gonzàlez, »Statistical Detection of JPEG Traces in Digital Images in Uncompressed Formats», IEEE Trans. on Information Forensics and Security, 2017

Outline

- Why social media are a good example of ‘big data’
- Challenges of social media:
  - Making structured the unstructured
    - Event-based media structuring
    - Media Synchronization
    - Media Diversity
  - Making trustworthy the untrustworthy
    - Source and provenance detection
    - Manipulation detection and biasing
- Open problems in big social media research

The Challenges:
Structuring multimedia information

- There is still a **gap in our capability to capture media semantics**
  - The new generation of neural networks are creating big expectations in research and industry, but…
  - …there is not yet a sufficient knowledge about how they work and how to manage the relevant models
- **Scalability and generalization** are still unsolved problems (managing the amount and variety of data)
- There are a number of opportunities for **applications**
  - Using big social media to create new user apps in different domains (health&well-being, mobility, turism, safety, etc.)

The Challenges:
Multimedia information trustworthiness

- **Robustness & Security**
  - Current techniques are mostly at research level, additional effort needed to make them robust/usable in the wild
- **Adversarial approaches**
  - It's a cat and mouse game, as far as forensic techniques improve, also manipulations become smarter.
  - Adversarial approaches mimic this process in advance
- **Generalization and Scalability**
  - Many specific techniques, no unique/standard framework
  - Need to check ever increasing amounts of information, this requires computationally efficient approaches

The challenges: Other (general) issues

- **Availability of data**
  - Only big players can fully access real data
  - The groundtruth is still a big problem
- **ANNs are defeating the competitors**
  - Current efforts are mostly focused on using neural approaches, is that the ultimate answer?
- **Is there an upper limit in “big data”?**
  - We are generating a lot of data, but what about information content?

How much "information" is really out there?

Extract from first page of Google image search for query "Coliseum"

S. Mc Cann, "3D Reconstruction from Multiple Images", 2014

How much “information” is really out there?

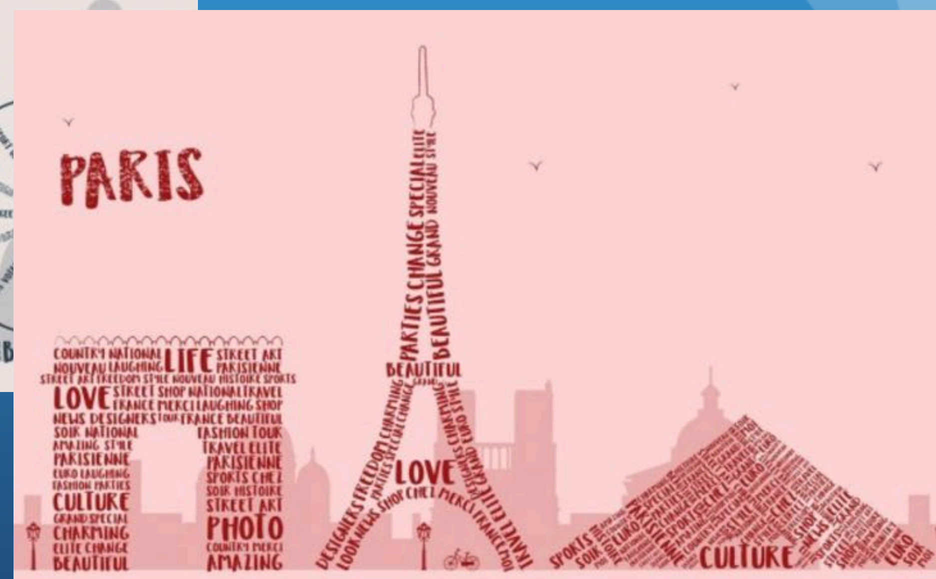

Most tweeted words in relation to famous cities

Source: hoppa.com

How much "information" is really out there?

- 4 Million "likes" every minute on Facebook
- Oxford Dictionary elected Word-of-the-Year 2015 an emoji

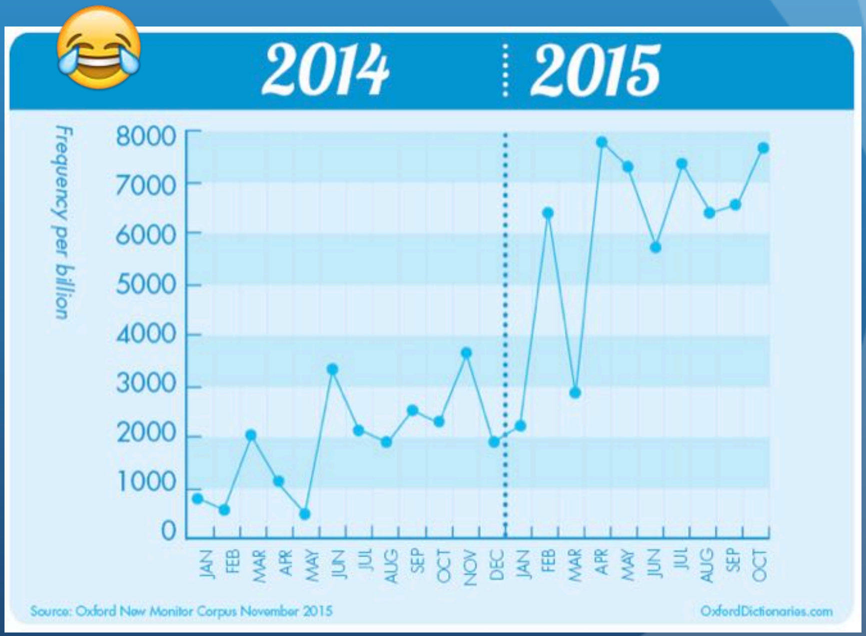

Frequency of use of "tears of joy" according to Oxford Dictionary

Does big data mean also big information?

- Time for a new chapter of Information Theory?

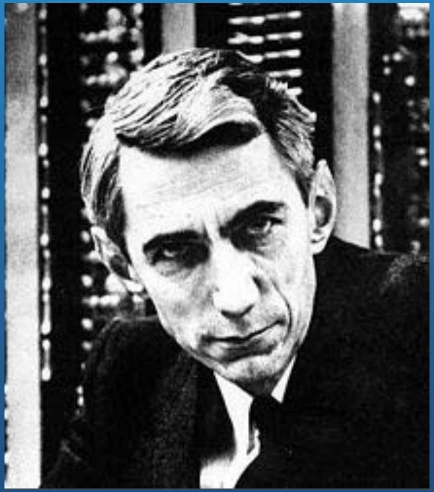

Claude Shannon (1916-2001)