Event Duality: Exploitation of Personal and Social Dimensions for Photo Indexing

IVAN TANKOYEU

KNOW DIVE

UNIVERSITÀ DEGLI STUDI DI TRENTO

October 21, 2013

Outline

- **Personal vs Social Events**
- **Event Mining Approaches: Overview**
- **Personal Event Detection** via spatio-temporal analysis
- **Social Event Detection** through the matching of personal events from different users
- Social **event co-participation analysis** to predict the existence of social ties (useful as a "friend recommendation" technique)

Personal vs Social Events

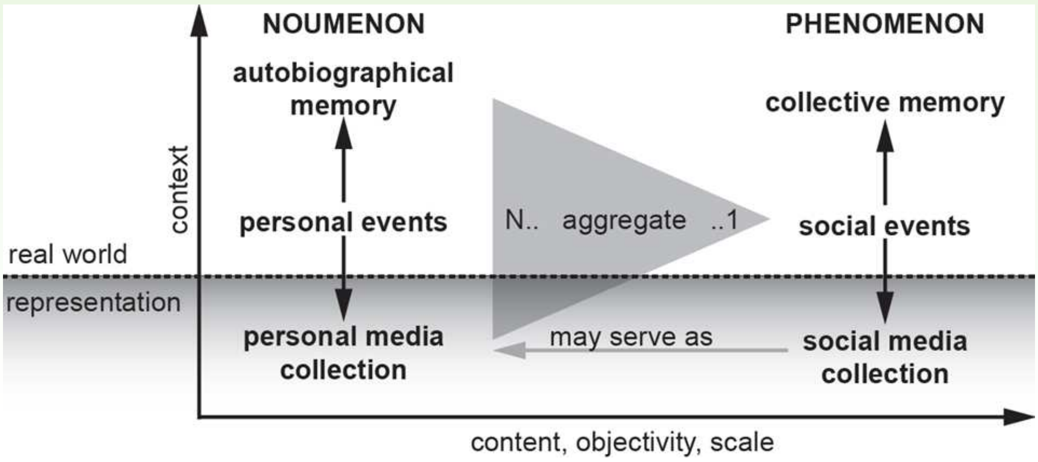

Personal and Social Event Detection: overview

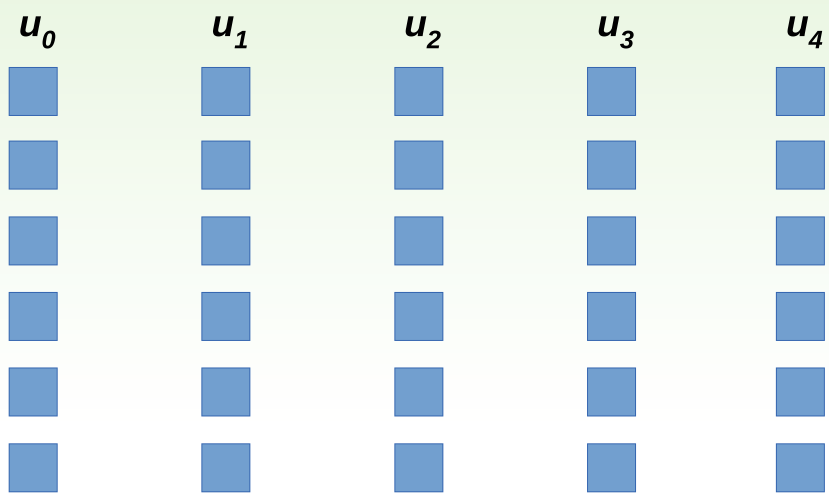

Personal and Social Event Detection: overview

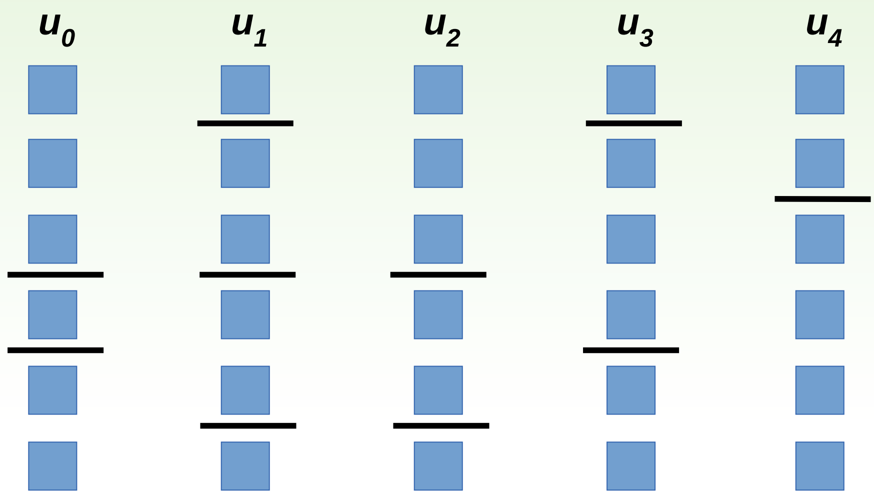

Personal and Social Event Detection: overview

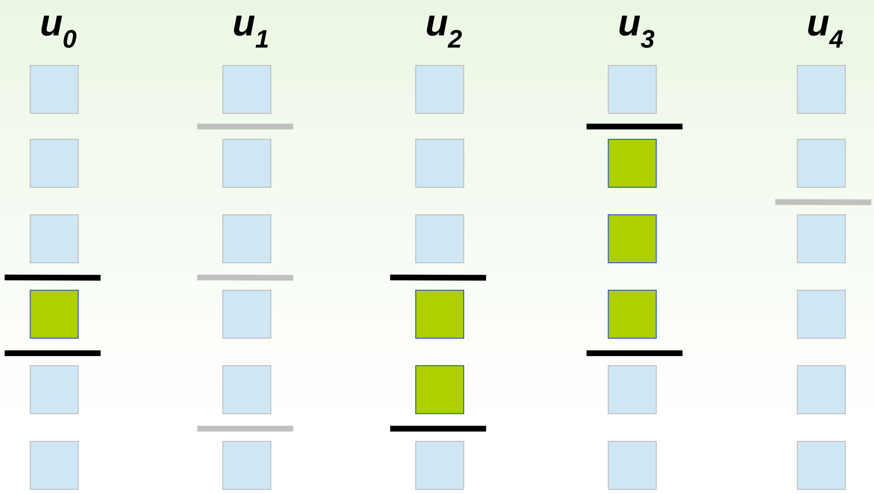

Social Events and Social Ties: overview

u2 and ***u3*** **co-participate in 1 event**.

We can find the probability of existence of a social tie between them**.**

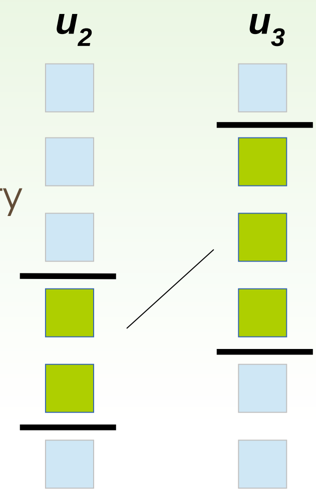

Social and Personal Event Detection: overview

u2 and ***u3*** **co-participate in 2 events**.

the probability of existence of a social tie is **higher** between them**.**

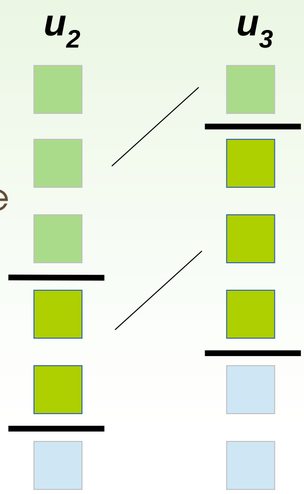

Personal Event Detection: Methodology

- **Input**: a personal photo collection where each photo has a timestamp and geo-coordinates
- **Output**: groups of photos corresponding to personal events

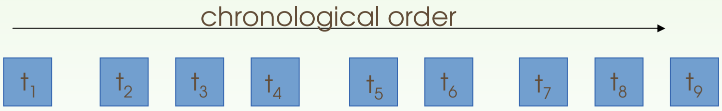

Personal Event Detection: Methodology

- **Input**: a personal photo collection where each photo has a timestamp and geo-coordinates
- **Output**: groups of photos corresponding to personal events

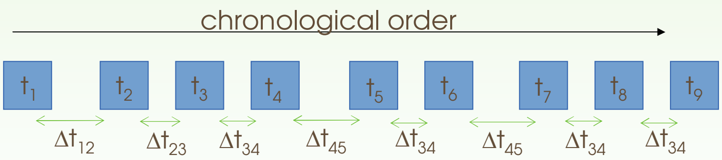

Find all the gaps:

- first in **space** (spatial distance between neighbours)
- then in **time** (elapsed time between neighbours)

For each step classify into small and big gaps using k-means **where k=2**

Personal Event Detection: Methodology

- **Input**: a personal photo collection where each photo has a timestamp and geo-coordinates
- **Output**: groups of photos corresponding to personal events

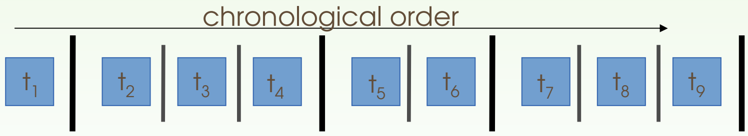

Find all the gaps:
- first in **space** (spatial distance between neighbours)
- then in **time** (elapsed time between neighbours)

For each step classify into small and big gaps using k-means **where k=2**

Personal Event Detection: Methodology

- **Routine location detection:** for each location compute the number of days in which photos was taken. The maximum number indicates the routine location for a period of time.

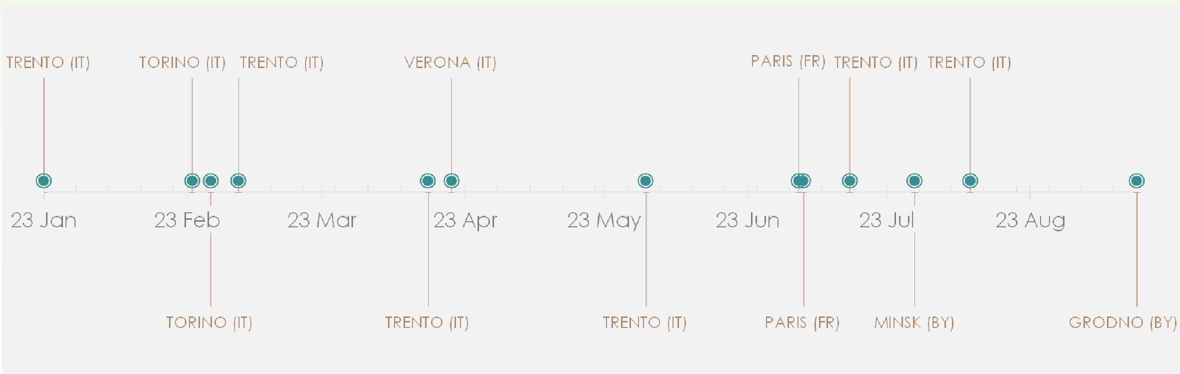

Personal Event Detection: Methodology

- **Routine location detection:** for each location compute the number of days in which photos was taken. The maximum number indicates the routine location for a period of time.

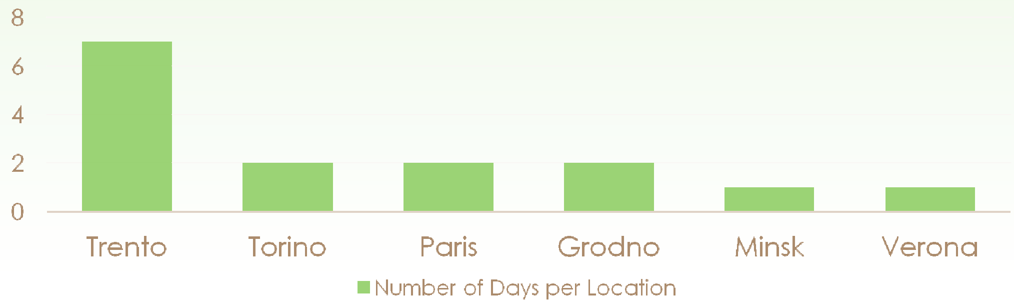

Personal Event Detection: Methodology

- **Routine location events:**

birthdays, graduation, Christmas, etc.

- **Non-Routine location events can be further separated to sub-events:**

event: ACM MM'13 Barcelona, trip to Europe

sub-event: EBMIP workshop, Workshop on Immersive Media Experience

Personal Event Detection: Dataset

- 6 users
- ~42 000 images
- Average duration per user 6.36 years
- All images are grouped in 726 folders (events) arranged by users. This folder structure is the **ground truth.**

Personal Event Detection: Results

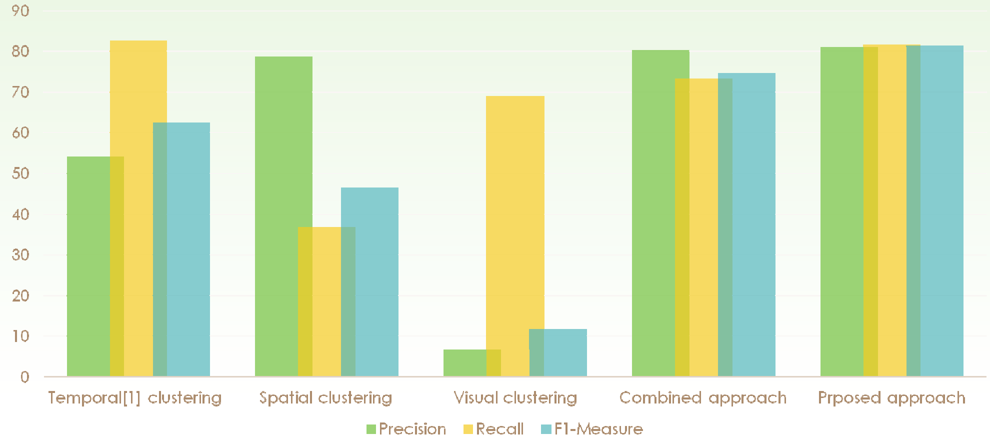

Social Event Detection: Methodology

- **Input**: personal events from different users
- **Output**: groups of personal events corresponding to social events

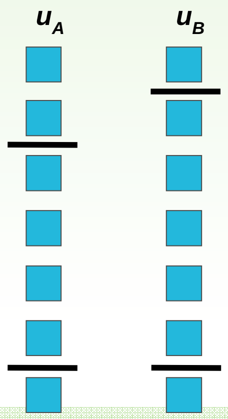

Social Event Detection: Methodology

- **Input**: personal events from different users
- **Output**: groups of personal events corresponding to social events

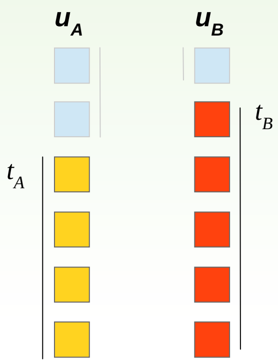

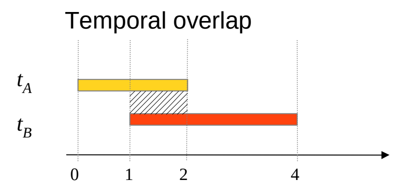

$$t_{sim}(t_A, t_B) = 1 - \begin{cases} \frac{t_A \cap t_B}{t_A \cup t_B}, & \text{if } t_A \cap t_B \neq \varnothing. \\ 0, & \text{otherwise.} \end{cases}$$

Social Event Detection: Methodology

- **Input**: personal events from different users
- **Output**: groups of personal events corresponding to social events

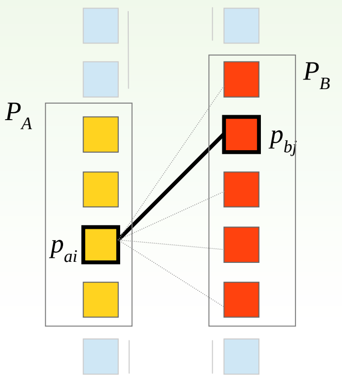

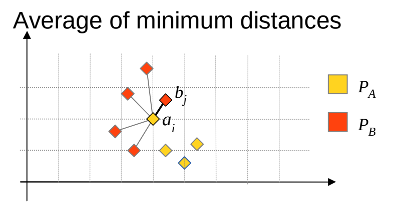

$$d_{min}(p, P) = \min_{i=1}^{n} d(p, p_i)$$

$$s_{sim}(P_A, P_B) = \frac{\sum_{i=1}^{n} d_{min}(p_{ai}, P_B)}{n}$$

Social Event Detection: Dataset

- **11 180** events from **Y! Upcoming**. Of these, **1291** events contain photos from more than 1 user
- For each Y! Upcoming event we have:
  - **Start-end dates**
  - **Location information** (Y! WOEIDs)
- Y!-Upcoming-tagged Photos in Flickr uploaded from 2007 to 2012
  - ID of owner
  - Timestamps
  - GPS coordinates
- For each user within the set of photos:
  - **Contact list**

Social Event Detection: Dataset

How do we select **c_T** and **c_S** values?

- We use **50%** of the Y! Upcoming events to find these thresholds and the remaining **50%** to test.

When we compare against Ground Truth it may happen that:

- There is a 1:1 relation between a Detected Social Event (DSE) and a Y! Upcoming Event (YUE). This is regarded as a correct detection.
- 1 **YUE** corresponds to more than 1 DSE: **under-joining**
- 1 **DSE** contains more than 1 YUE tag: **over-joining**

Social Event Detection: Results

c_T	c_S	U-joint	Correct	O-joint	Correct%
0.75	0.50	28	265	32	81.54
0.75	1.00	26	265	32	82.05
0.75	5.00	23	264	34	82.24
0.75	10.00	23	263	34	82.19
0.50	0.50	28	267	31	81.90
0.50	**1.00**	**26**	**267**	**31**	**82.40**
0.50	5.00	23	264	34	82.24
0.50	10.00	23	263	34	82.19
0.25	0.50	45	254	14	81.15
0.25	1.00	45	254	14	81.15
0.25	5.00	43	254	15	81.41
0.25	10.00	43	253	15	81.35

According to the parameter learning phase **c_T = 0.50** and **c_S = 1.00**
Using these parameters for the test set:
- Correct detection: **76.78%**
- Under-joining**: 9.56%**
- Over-joining: **11.68%**

Social Event Detection: Dataset

We analysed our dataset to see if there is a correlation between:

- **event co-participation** (two Flickr users have photos with the same Y! Upcoming tag)
- and the existence of **a social tie** (two Flickr users have each other in their contact lists)

A rapid analysis of our 1291 events with more than one participant shows that 1039 **(80.40%)** have at least 1 pair of participants that "know each other".

Social Event Detection: Results

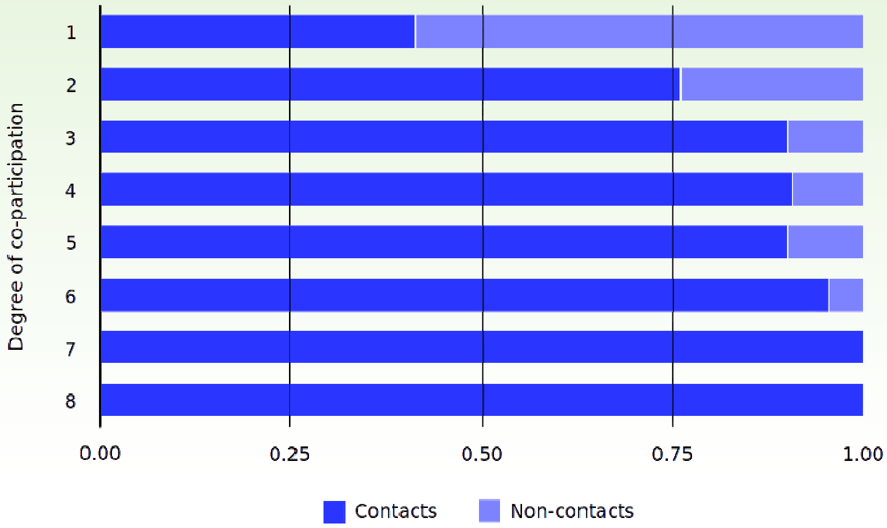

Social Event Detection: Results

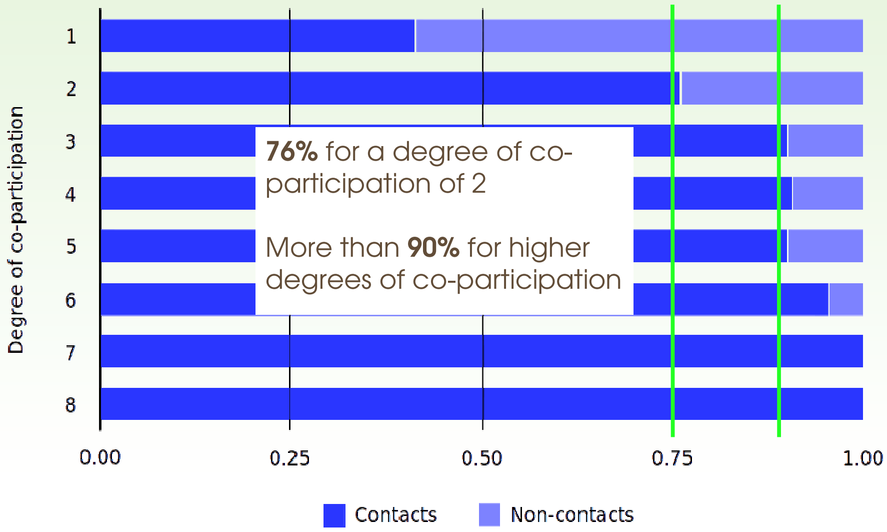

Conclusion

For the Personal Event Detection Task:

- Our approach detects events with F1-measure equal to **81.35%**

For the Social Event Detection Task:

- We are able to produce correct detections of social events in **78.76%** of the cases.
- This number rises to **88.32%** if we include cases in which we only achieve partial social event detection (under-joining).

For the Co-participation Analysis:

- Two users "know each other" in **76%** of the cases if they co-participate in at least **2 events**. This number rises to almost **90%** for higher degrees of co-participation, meaning it can be used for "friend recommendation".

Thank you!

Questions?

Used Material

- "Event detection and scene attraction by very simple contextual cues" I.Tankoyeu, J.Paniagua, , J.Stöttinger, F.Giunchiglia, *Proceedings of the 2011 joint ACM workshop on Modeling and Representing Events (JMRE'2011)*
- "Indexing media by personal events" J.Paniagua, I.Tankoyeu, J.Stöttinger, F.Giunchiglia, *Proceedings of the ACM International Conference on Multimedia Retrieval (ACM ICMR'2012).*
- "Personal photo indexing" I.Tankoyeu, J.Stöttinger, J.Paniagua, F.Giunchiglia
  Proceedings of the 20th ACM international conference on Multimedia, (ACM MM'2012).
- *"Event-Based Media Indexing"* I.Tankoyeu, *PhD Thesis.*
- "Social Events and Ties" talk by J. Paniagua at the *3rd ACM International Conference on Multimedia Retrieval (ACM ICMR'2013).*
- "Social Events and Social Ties" J.Paniagua, I.Tankoyeu, J.Stöttinger, F.Giunchiglia, *Proceedings of the 3rd ACM International Conference on Multimedia Retrieval (ACM ICMR'2013).*