

DISI - Via Sommarive 14 - 38123 Povo - Trento (Italy)
<http://www.disi.unitn.it>

CLUES OF PERSONAL EVENTS IN ONLINE PHOTO SHARING

Pierre Andrews, Javier Paniagua, Fausto
Giunchiglia

August 2011

Technical Report # DISI-11-474

Clues of Personal Events in Online Photo Sharing

Pierre Andrews, Javier Paniagua, and Fausto Giunchiglia

Dipartimento di Ingegneria e Scienza dell'Informazione
University of Trento, Italy
{andrews,paniagua,fausto}@disi.unitn.it

Abstract. There is currently a trend in media management and the semantic web to develop new media processing methods and knowledge representation techniques to organise and structure media around events. While this increased interest for events as the central aggregator when organising media is supported by strong research in the fields of knowledge representation and computer vision; it is not yet clear how the digital era users use events when sharing their personal media collection. In this paper, we explore how users share photos online and discuss the results of a preliminary automatic processing of the data collected. We show that while media sharing services do not support events as yet, users still share their media around personal events, either by providing explicit spatio-temporal metadata, or by using an event-centric vocabulary.

1 Introduction

With the increased availability of digital capturing devices, people now build large personal media collections; what is then done with these media has drastically changed in the last years with the emergence of popular photo sharing services like Flickr¹ or Picasa² and social networks sites such as Facebook³. Understanding how people organise and share their digital collections is key to building better tools that accommodate for the users' needs instead of forcing them to change their mental model to fit a fixed software workflow. This workflow has changed with the introduction of new technologies and ways to interact and share media online; the user's goal is now, not only to archive media for a personal use but also to share them with relevant contacts online. Therefore the issue is not only of organising the user personal media collection for better future search and retrieval, but also the one of organising shared media for visibility to the relevant people and future search and retrieval of these media not just for the author that built the collection but also by these relevant people.

Until recently, the “album” concept was one of the main metaphor for helping users to organise their personal collection, thus staying close to how photo prints

¹ <http://www.flickr.com>

² <http://picasaweb.google.com>

³ <http://www.facebook.com>

where organised previously. However, new metaphors of organisation are now emerging to leverage more complex indexing and search. Flickr for instance has introduced a very loose organisation system, focusing on tags to group photos, and with the availability of GPS technology, media management services have introduced the possibility to “geotag” media and to browse and search them with location based services. Some have also introduced search and navigation services based on who is in the photo and when it was taken, using the metadata provided by the camera.

This use of media metadata is moving away from the physical photo album metaphor. However, there is still a semantic gap between the low level metadata, the high level information of who is in the photo or where it was taken and how people group their media for personal archiving and sharing. In fact, most popular media management services still provide the “album” metaphor⁴ as people still have a need to group media together in ways that are more meaningful to them than just a location or time grouping. Researchers are thus focusing around the *event* metaphor to combine metadata seems to represent part of the higher level intent of the users when they group their media.

While this metaphor is backed by some early user studies, these were led before the large adoption of social media sharing services and there has been little recent research on how users actually use events digitally to organise and share their media. Discovering if this is the case is not an easy task, and in this paper we discuss an initial study of the sharing behaviour of users on Flickr and Picasa to see if they are currently using events when sharing their photos online. We first introduce the current work on event representation for media management and the current model of events (Section 2). In the following sections we discuss how we have collected data on Flickr and Picasa (Section 3), how people use event metadata when organising in the *album* metaphor of these sites (Section 4) and how, even when they do not explicitly use such metadata, they often share media by using an event-centric vocabulary (Section 5).

2 Events for Media Organisation

It seems that people take photos to archive important events and share within their close community. [16] and [11], found clues that people intentionally classify photos according to events in their lives. This fact has been observed even before the advent of digital photography in [2]; however, Chalfen argues that people do not share pictures per-se but use them to tell a story. More recently, Miller et al’s user study [12] similarly concluded that users took photos primarily to archive important events and share within their community; however, at the time of their study, they found that layman users did not share photos actively online and preferred to use prints or email.

This organisation of photos “chronologically by event” eases the search and retrieval of specific photos in personal collections as it aligns with the way memory is structured. According to [21], humans identify activity boundaries at

⁴ Flickr calls it a “Set”.

points that correspond to a maxima in the number of changing physical features, thus aggregating memories around events. [10] states that the brain operates in this way to cope with the increased difficulty brought by indexing new information when it is dissimilar from the “current moment” beyond a certain threshold. Some are thus proposing an event-centric models to characterise media in terms of the events they are associated with [6,8,4].

Last.fm and Upcoming.org are services that already try to link media and event. They do so for public events such as concerts or conferences, but still do not allow users to share their personal events (e.g. weddings, birthdays). [3] presents a user study to elicit requirements for such services and interaction paradigms that help discover and enrich public events. While this approach is interesting for public events, it does not clarify if users naturally use personal events to organise and share their personal media collections.

[21] recognises subject, actors and causal properties as components of the human perception of an event, stressing the importance of the *temporal and spatial* aspects to build the event structure. [1] defines events as having a close link to their *spatio-temporal collocation* and to the things that constitute their subject (e.g., a sparrow in the event “a sparrow falls”). Inter-event relations are studied in [17] that states that events may be composed of sub-events that are temporally, spatially and causally connected. [7] explores use case scenarios to show possible ways in which untrained users may organise media in terms of events with complex *spatio-temporal structure*.

Practical models for events can be found in the IPTC G2 family of news exchange standards are provided. EventML⁵ is one of these standards oriented at describing public events in a journalistic fashion, although support for media is limited, and this model is close to Chalfen’s idea that media are only used to support a story. A set of requirements for a base model of events is presented in [20] that categorises all the properties and relations of an event into six aspects: *temporal, spatial*, informational, experiential, structural and causal. The F event model [18] specifically addresses most of these requirements, [19] also addresses the *temporal, spatial* and informational aspects by integrating different ontological models. The Simple Event Model is proposed in [5] to represent not only who did what, when and where, but also to model the roles of each actor involved, when and for how long this is valid and according to whom. MediAssist [14] organises digital photo collections using time and location information combining it with content-based analysis (face-detection and other feature detectors). The work in [15] uses time and latitude/longitude data to analyse tags and unstructured text from photos on Flickr to extract place and event semantics. VisR⁶ is a smartphone application that detects events from photos and metadata available on the device. All these studies have in common the predominance of the spatio-temporal aspect of events as it is the one that helps users determine inter-event boundaries, recollect their memories and find

⁵ http://www.iptc.org/site/News_Exchange_Formats/EventsML-G2/

⁶ <http://www.visrapp.com/>

their media. Thus, events refers to “something that occurs in a certain place during a particular interval of time” [6].

3 Photo Sharing Websites: Data Collection

Event-centric services such as Upcoming.org or Last.fm are focused on public events such as concerts or conferences. While datasets [3] based on these websites already provide samples of media organised around event metadata, they do not represent personal events. That is, media of more personal events, such as a birthday or a holiday, are not shared on these websites. However, this kind of media can be found on photo sharing websites such as Flickr and Picasa where users share photos of personal happenings with their family and friends.

These websites do not provide a way to organise photos around events but provide a way to group photos in *albums*. These albums can only have a very small amount of metadata and are not presented as events to the user. On Picasa, albums can have a *title* and a *description*, and optionally a *date* and a *location*; on Flickr, sets can only have a *title* and a *description*.

We are interested in seeing how the users describe albums they share on Picasa and Flickr by using the title and description fields. Our hypothesis is that if they share media related to events, they will provide the event metadata in the fields that are available to them and we will find event references in the titles and descriptions of the albums. We are focusing on these two social sharing sites as they are some of the more popular available at the time of writing; while Facebook is also very popular, it provides very similar features (album based organisation of photos) and does not allow data collection.

We have thus collected a dataset of digital albums shared on Flickr and Picasa. To select users, we use the “explore” pages of each website that feature randomly selected photos; from these photos, we find a set of random users and collect all public albums that are shared by these users. For each album shared on Flickr we retrieve: (a) the *title* of the set, (b) the *user identification*, (c) the *URL* of the set and (d) the *number of photos and videos* within the set. For each album from Picasa we collect: (a) its *URL*, (b) the *date* specified for this album, (c) the *number of photos*, (d) the *title*, (e) the *description* and (f) the *user identification*.

Because both websites are international, many entries are not written in English. In this paper, we are only able to process metadata provided in English and thus want to filter out the other languages. The perl `Lingua::Identify`⁷ module was used to identify the language of the title and description (when available) in each album entry. The algorithm provided by this module was trained on the EuroParl corpus [9]; we have performed a manual annotation of a subset of the automatically processed entries from the Picasa dataset and have found that the algorithm labels English albums with a 89% precision.

We are interested to see if users refer to locations when they describe albums and have thus automatically processed the dataset to find references to geo-

⁷ <http://search.cpan.org/~ambs/Lingua-Identify/>

graphic locations. The Yahoo! Placemaker⁸ service is used to perform this task. This freely available geoparsing service can identify place references in unstructured text. While Yahoo! does not provide information on the accuracy of their algorithm, from our manually annotated sample, we found that Placemaker is able to detect if there is at least one location reference in an English title with 81.2% accuracy.

References to dates are also of interest to us as time is a main attribute of an event. To detect such references, we analysed each title with a custom date parsing algorithm that detects full dates but also partial dates (e.g. “Paris’08”) and periods (e.g. “40.5 miler in Sespe Wilderness April 2nd - 5th 2010”). On our manually annotated sample, this algorithm performed with 88.1% accuracy.

We have collected 32 168 sets from Flickr and 88 593 albums from Picasa over the month of July 2011⁹. We have kept only English albums, resulting in 5 339 (16.6%) sets from Flickr, and 11 355 (12.8%) albums from Picasa.

4 A Given Place and Time

According to the definition that we introduced in Section 2, the two main attributes defining an event are its location and when it happened. Thus, if users are to describe events using albums when sharing their photo, they will probably specify some of these attributes within the available attributes. We found that in the Picasa dataset, only 31% of the albums have a description and thus, in this paper, we focus on the title attribute of the albums as we do not have enough data to draw conclusions from the descriptions.

Table 1. Proportion of Albums with Titles Referring to Dates or Location

	FLICKR	PICASA	FLICKR+PICASA
Dates	33.9%	44.6%	41.2%
Locations	22.4%	26.7%	25.3%
Both	8.7%	12.9%	11.6%

Table 1 shows the proportions of albums where date or location references can be found, a test of equal proportion shows that Picasa and Flickr are comparable ($p < 0.01$) and we thus consider that there is no difference in users’ behaviour between the two services in the factors we analyse.

The number of albums where an explicit date reference can be found in the title makes for more than a third of the dataset. We can thus see that people do like to share their albums with metadata about the date when the photos were taken. Note that while the title is set manually by the users, the *date* field

⁸ <http://developer.yahoo.com/geo/placemaker/>

⁹ the random crawling collected albums posted between 2006 and 2011.

on Picasa is filled automatically with the album creation date if the user does not specify any value explicitly. When there is a date in the title on Picasa, it is often not consistent with the album *date* field. It seems that while the users are ready and interested to share their photos around dates, they are not motivated to fill in an extra metadata field. The reason behind this might be a limitation in Picasa’s interface or it can simply be because the users do not see the gain in filling this extra field.

While the date is an important attribute of events, albums with only a date reference are not always events according to our previous definition. In fact, from a preliminary manual annotation of the Picasa dataset, we can see that 27.3% of the albums with a date reference but no location reference are not really events. This is because there are catch-all albums for entire years or months, where users put photos of many different events in the album (e.g. “Misc. Apr. 2009”). The album is thus only a way to aggregate photos in a time range and not used to represent a specific event. This happens also when people share photos of their newborn child for milestone periods (e.g. “Jake - 9 months: March”).

There are less albums with an explicit location reference, but it still makes for a fourth of the dataset. From the manual annotation, we can see that 78.7% of the albums with only a location reference are actually events. In the same way as with the dates, users use locations for catch-all albums where they put photos of a location they visited multiple times but not for any specific event (for instance for photos of their home-town).

In these two cases, we can see that the dates and locations are sometimes used only as aggregators for media that could be replaced by automatic metadata based services. However, it seems that the users are not aware of, or willing to use, these services on the studied websites.

96.8% of the albums with a date and a location together were annotated as being events by the manual validators. While these albums represent a small amount of the dataset, we can already see that when space and time are specified in the title, the users wanted to share an important event.

5 An Event Vocabulary

In the previous section, we have looked at how users might use album attributes to describe explicitly an event location or date. However, there are many events represented on Picasa and Flickr that do not include explicit dates or locations. For instance “Janet and Ian’s wedding”, “father’s day”, “Michelle’s shower” or “Christmas Eve” are all titles of albums from our dataset that do represent important personal events with no explicit dates or locations. Thus, there might be more albums in this dataset that represent events than the previous section’s analysis hinted.

In fact, if we look at the most popular words used in the titles (see Table 2), many of them are references to events (e.g. “party”, “wedding”, “trip”) or time periods, without having explicit dates. Note that, while not shown in Table 2, the most popular words in the vocabulary are years, in fact on Flickr, 11.0%

of the vocabulary are numerals while on Picasa 17.5% of the words used are numbers. Table 2 reports figures in per-thousand, while the distribution of the vocabulary follows a very steep long-tail curve, the most popular words still do not cover a large part of the album vocabulary.

Table 2. Most Popular Words in Titles (% of the whole vocabulary)

FLICKR	%	PICASA	%	FLICKR+PICASA	%
spring	5.88	new	4.72	spring	1.76
city	6.24	trip	5.82	city	1.87
day	7.80	wedding	8.30	day	2.34
wedding	12.78	day	11.73	wedding	3.83

While it is easy to see that in the most used words in the dataset there are concepts representing events, it is not an exhaustive view of the dataset and it would be interesting to see how many albums refer to events by using such vocabulary. However, it is not easy to exhaustively list manually the whole vocabulary that could be used to refer to events. We take a semi-automatic approach, using WordNet [13] as a thesaurus, to find all terms that might refer to a concept representing an event. To do so, we have listed all inherited hyponyms of the synset `event#n#1` – which include the words “wedding”, “birthday”, etc. – and of the synset `calendar.day#n#1` – which include the words “Christmas”, “Thanksgiving”, etc. This provides us with a list of 11 092 words and 14 304 concepts combined in 15 389 word-concept pairs¹⁰ that we then searched in the titles of the albums in the dataset.

Table 3. Top Leaf Concepts Related to Events

FLICKR			PICASA		
	Events %	Overall %		Events %	Overall %
Sunday#n#1	3.33	1.64	marriage#n#3	3.42	1.74
Easter#n#1	3.41	1.68	Easter#n#1	4.37	2.22
Michigan#n#3	3.41	1.68	Halloween#n#1	4.51	2.29
Halloween#n#1	4.28	2.11	Christmas#n#2	5.30	2.70

We found that around half of the albums (Flickr: 49.4%; Picasa: 50.9%) have a title with at least one word that represents an event according to WordNet. Of these albums, only 29.6% have a date or a location (or both) in the title. There are indeed many albums that describe events without providing either an explicit date or a location reference (e.g. “Katie’s Swiss trip”, “Field trip - Farm”,

¹⁰ Note that because of homography, the same word can appear under different concepts.

Table 4. Top Concepts Related to Events – cumulating the hyponyms occurrences

	FLICKR		PICASA	
	Events %	Overall %	Events %	Overall %
calendar_day#n#1	24.2	11.9	27.5	14.0
activity#n#1	39.3	19.4	32.4	16.5
act#n#2	66.3	32.8	64.3	32.8
event#n#1	75.8	37.5	72.5	36.9

“Lily fathers Day”). From a preliminary analysis at these albums, it seems that many of them either refer to the third important attribute of an event: the participants; or to relative dates (e.g. “Father’s Day”, “My Birthday”) or locations (e.g. “Trip Home”). In fact, we can see in Tables 3 and 4 that the `day#n#3` and `calendar_day#n#1` are among the most used concepts. This is in line with Jain’s [6] definition of an event: “a significant occurrence or happening, or a social gathering or activity”. However, relative location or participant references are hard to detect automatically and further work is required to check how these are used in the album vocabulary.

WordNet is a very detailed vocabulary and many terms that it declares as relating to the *event* concepts might not be used by the users to refer to events. Indeed, there is ambiguity in the vocabulary and we have taken a naive approach where we count the occurrence of all possible words without applying disambiguation. For instance, `Michigan#n#3` appears as one of the most popular leaf concepts for Flickr; however, this concept represents a card game called “Michigan” but might have been used by users in their album title as the location. The other top concepts however represent less ambiguously event references.

This confirms Chalfen’s conclusions that people like to take photos around personal events that they then share with a community made of close relations ([2]). However, as we have discussed earlier, these photos are usually shared without description, and thus Chalfen’s hypothesis that people use photo to tell a story might not be exact on photo sharing websites.

6 Discussion

The results we found, while preliminary, show that there is a tendency for users to share photos around places and location. While this is not a guarantee that they are sharing albums about specific personal events, it seems to align with the previous observations of Zacks et al and Kurby et al [10,21] who found that users like to segment their memories around time and space.

While most event models discussed in the state of the art (for instance [19]) represent events around dates and locations too, they do not seem to fit perfectly the behaviour of the users that we observed on the sharing sites. In particular, some users seem to aggregate media around date or location without describing events (e.g. the newborn album cases pointed out earlier). While this could be

done automatically from the metadata of the photos, there might be a higher semantic to this grouping when sharing. As Chalfen [2] discusses, even if it was for printed photo, people group photos together to support a story and not always just for the content of the photos per-se. That is, the grouping of photos of the “second month” of a baby is not a specific event according to most of the existing metadata models but is still an event of importance for the users that share them online.

In addition, in accordance to Chalfen’s [2] and Miller et al. [12], people share photos around important personal events. These events (e.g. Christmas, trips, visits) are not always global events and their scope is limited to the close circle of personal relationships. This kind of sharing has probably a different purpose from the one of exploring concert or conference photos (for instance) as is discussed in [3], or from the news outlet use-cases for which the IPTC standards have been developed¹¹. Therefore, we need custom model and services for layman users.

As was pointed out in [12], there is also a stronger issue of privacy and access control when dealing with the sharing of personal events. On Picasa and Flickr, we were able to crawl public albums – featured on the website main pages – that were of highly personal nature but are accessible to anyone online. While this is not the scope of this paper, we believe that there is a need for better privacy services directly integrated with the event models to deal with the personal media sharing use-cases.

7 Conclusion

In this paper we present a preliminary study of a dataset of albums shared on Flickr and Picasa. As we show, while these two services have different interfaces and features, users tend to have the same behavior on both sites and we believe that this demonstrates some general intent of the users more than site-specific behaviors. While this is a raw analysis of the data and a more extensive manual annotation is required, we have found that a significant amount of users share media online illustrating personal events, and use time-location metadata to describe them. In fact, we have found that more than a third of the albums shared reference a date in their title and more than a fourth refer explicitly to a location. Users also seem to group their photos around important personal events (e.g. birthdays, wedding, festivities) without always specifying explicitly a location or date.

We are planning future work, in particular in analysing the user needs and habits directly with the users, it shows that they already try to use events when sharing media, even when the applicative workflow does not allow it explicitly. We are also planning to extend this work to study the current use of geo-tagging when sharing media. Therefore providing users with new interfaces and services using the event metaphor should improve their experience and the searchability of the media they share online.

¹¹ <http://www.iptc.org>

Acknowledgment

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n°248984 GLOCAL and n°247758 EternalS.

References

1. Casati, R., Varzi, A.: Events. In: The Stanford Encyclopedia of Philosophy. The Metaphysics Research Lab; Stanford, CA (2010)
2. Chalfen, R.: Snapshot Versions of Life. Bowling Green State University Popular Press, Bowling Green, Ohio (1987)
3. Fialho, A., Troncy, R., Hardman, L., Saathoff, C., Scherp, A.: What's on this evening? In: EVENTS'10 (2010)
4. Giunchiglia, F., Andrews, P., Trecarichi, G., Chenu-abente, R.: Media Aggregation via Events. In: EVENTS'10 (2010)
5. van Hage, W.R., Malaisé, V., Segers, R., Hollink, L.: Design and use of the Simple Event Model (SEM). *Journal of Web Semantics* (Accepted (2011))
6. Jain, R.: EventWeb: Developing a Human-Centered Computing System. *Computer* 41(2), 42–50 (2008)
7. Jameson, A., Buschbeck, S.: Interaction design for the exchange of media organized in terms of complex events. In: EVENTS 2010 (2010)
8. Kim, P.: Event-based Multimedia Chronicling Systems. *Computer Engineering* pp. 1–12 (2005)
9. Koehn, P.: Europarl: A parallel corpus for statistical machine translation. In: MT summit. vol. 5 (2005)
10. Kurby, C.A., Zacks, J.M.: Segmentation in the perception and memory of events. *Trends in Cognitive Sciences* 12(2), 72 – 79 (2008)
11. Lansdale, M., Edmonds, E.: Using memory for events in the design of personal filing systems. *International Journal of Man-Machine Studies* 36(1), 97–126 (1992)
12. Miller, A.D., Edwards, W.K.: Give and Take : A Study of Consumer Photo-Sharing Culture and Practice. In: CHI'07. pp. 347–356 (2007)
13. Miller, G.A.: WordNet: a lexical database for english. *Communications of the ACM* 38(11), 39–41 (1995)
14. O'Hare, N., Lee, H., Cooray, S., Gurrin, C., Jones, G., Malobabic, J., O'Connor, N., Smeaton, A., Uscilowski, B.: MediAssist: Using Content-Based Analysis and Context to Manage Personal Photo Collections. In: CIVR'06. pp. 529–532 (2006)
15. Rattenbury, T., Good, N., Naaman, M.: Towards automatic extraction of event and place semantics from flickr tags. In: SIGIR '07. p. 103 (2007)
16. Rodden, K., Wood, K.R.: How do people manage their digital photographs? In: CHI'03. p. 409 (2003)
17. Scheffler, U.: Events as shadowy entities. *Philosophy* 2, 35–53 (1994)
18. Scherp, A., Franz, T., Saathoff, C., Staab, S.: F—a model of events based on the foundational ontology dolce+DnS ultralight. In: K-CAP'09 (2009)
19. Shaw, R., Troncy, R., Hardman, L.: Lode: Linking open descriptions of events. *The Semantic Web* pp. 153–167 (2009)
20. Westermann, U., Jain, R.: Toward a Common Event Model for Multimedia Applications. *IEEE Multimedia* 14(1), 19–29 (2007)
21. Zacks, J.M., Tversky, B.: Event structure in perception and conception. *Psychological Bulletin* 127(1), 3 – 21 (2001)