



HAL
open science

Who's Time Is It Anyway?

Bart Thomee, José G. Moreno, David A. Shamma

► **To cite this version:**

Bart Thomee, José G. Moreno, David A. Shamma. Who's Time Is It Anyway?. The 22nd ACM International Conference of Multimedia, Nov 2014, Orlando, United States. pp.4 Pages, 10.1145/2647868.2655066 . hal-01070351

HAL Id: hal-01070351

<https://hal.science/hal-01070351>

Submitted on 1 Oct 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Who's Time Is It Anyway?

Investigating the Accuracy of Camera Timestamps

Bart Thomee
Yahoo Labs
San Francisco, CA, USA
bthomee@yahoo-inc.com

Jose G. Moreno
Normandie University
Caen, France
jose.moreno@unicaen.fr

David A. Shamma
Yahoo Labs
San Francisco, CA, USA
aymans@acm.org

ABSTRACT

People take photos all over the world at all times of day; each photo depicting a place and a moment worth capturing. In the context of multimedia analysis and social computing, accurate location and time information about where and when these photos were taken is of importance for understanding event semantics, image content and many other purposes. While location information associated with photos is known to be relatively accurate, time is not. From a sample of 10 million public Flickr photos, we observe that 37% of the photos differ more than an hour between their camera timestamps and GPS timestamps with respect to local time at the locations where the photos were taken. Erroneous time information may adversely influence the correctness of any kind of temporal analysis that relies on camera timestamps, as well as research and real-world applications that require accurate knowledge of when and where photos were captured. In light of our observations we propose a simple yet effective metadata-only technique for improving the accuracy of camera timestamps.

Categories and Subject Descriptors

H.3.1 [Information Storage and Retrieval]: Content Analysis and Indexing

Keywords

Georeferenced photos; GPS information; time accuracy

1. INTRODUCTION

Spatial and temporal metadata of photos is frequently used in the research community to analyze the behavior of how people act and move in the world [3, 7, 10]. In the absence of GPS information, such techniques rely on the location that is manually supplied by the user and on the timestamp of the internal clock of the camera, which at the moment of capture gets written as EXIF metadata alongside the photo. The implicit assumption in these scenarios is that people are generally able to retrospectively place their photos at the

location where they were taken and that they are able to keep the clock of their camera up to date. This assumption is particularly important when location and time form the vital ingredients of a proposed technique, for example when recommending time-dependent travel itineraries for tourists visiting a city for the first time [2]. Some online photo sharing services also include or rely on a non-representative photo uploaded time instead.

The alternative to camera time is to focus on GPS data, which is designed to provide highly accurate location and time information. Many portable devices have integrated GPS receivers, such as modern photo cameras and smartphones which are the most popular cameras on many websites like Flickr [4]. These devices have a limited battery life and need to trade off spending energy against refreshing the data [5]. At the moment a photo is taken the GPS measurements may thus be several minutes or even hours old. The GPS receiver may additionally continue to broadcast its last known location and time, even when it is not able to obtain updated information at all, for instance when the camera-phone is indoors. While GPS is claimed to currently have a root mean square error of just 4 meters [9], the actual accuracy obtained can vary depending on manufacturing quality of the receiver, structural interference, atmospheric conditions and signal reception. For example, GPS devices have been shown to provide inaccurate positioning with maximum errors exceeding 300m across a wide range of urban areas, where even with good visibility the maximum error could be greater than 100m [6], while a study on mobile Foursquare check-ins revealed a median error of 70m and a mean error of 551m [8]. A recent ethnographic study has also shown reported cameraphone GPS positioning to deviate off target [1] during navigation tasks.

To the best of our knowledge, no one has yet investigated whether camera timestamps and GPS timestamps correspond with local time at the locations where the photos were taken. Intuitively, it is particularly easy to forget to change the clock of the camera when time shifts between summer and winter time, or when one travels between different time zones. The internal clock may further automatically reset itself when the batteries of the camera run out. Even though smartphones, which are increasingly used as cameras, are capable of obtaining local time information directly from the cell towers they are connected with, this functionality is overridden when the clock is manually set. In the case of the Apple iPhone, if it is set in 'Airplane mode'

<https://www.flickr.com/cameras/> Accessed 4/2014.

the last cached GPS information is stamped on every photo taken. Erroneous time information may adversely influence the correctness of any kind of temporal analysis based on photo timestamps. Gaining an understanding of the error range is therefore important, and in this paper we therefore set out to assess the extent of which camera timestamps, GPS timestamps, and local time differ from each other. To reduce the timestamp error we further propose a simple yet effective metadata-only technique that can correct the camera timestamps of all photos in a session when at least one of the photos is also associated with a GPS timestamp.

2. TIMESTAMP ANALYSIS

The correctness of the camera timestamp depends entirely on the user keeping the internal clock of the camera up to date. While we assume the camera timestamps refer to local time at the location of capture, the GPS timestamps, according to the EXIF specifications, refer to Coordinated Universal Time (UTC) instead and as such need to be converted to local time first. We used the `tz_world`² collection of shape files that capture the boundaries of the IANA time zones³ across the world to assign the GPS location associated with a photo to its designated time zone. Each time zone is associated with a time offset with respect to the UTC time zone, where the offset may vary between summer time or winter time. The GPS location and timestamp together thus allow us to determine the correct offset in seconds to apply to the GPS timestamp to convert it to local time. From this point forward we will refer to the GPS timestamp converted to local time as the `GPSL` timestamp.

If one forgets to adjust the camera clock when traveling to a different time zone then the camera timestamps will be ahead or behind local time by at most 24 hours (e.g. flying from Kiritimati to Hawaii). Furthermore, the transitions from summer time to winter time and vice versa also change the time of day, so that if the camera clock is not adjusted accordingly then its timestamps will differ from local time by half an hour (Lord Howe Island, Australia) or one hour (rest of the world, if daylight savings is observed). In addition, if the GPS device was not able to refresh its data at the moment of capture (e.g. the user is indoors), the `GPSL` timestamps will be behind local time with an amount that depends on the previous instant it managed to perform a refresh, which may be minutes, hours or even months.

In the next section we first analyze the discrepancies between the camera and `GPSL` timestamps, which we follow up with an investigation into the extent of which both timestamps differ from local time given a set of reference photos.

2.1 Differences

We queried the Flickr API for photos that contained a camera timestamp, a GPS timestamp and a GPS location in their EXIF metadata, yielding a total of 10 million photos. For each photo we computed the differences between its camera and `GPSL` timestamps, as shown in Figure 1. Our results indicate that 60% of the camera and `GPSL` timestamps differ less than 10 minutes from each other, 63% less than 1 hour, 90% less than 24 hours and 93% less than 48 hours. In 19% of the cases the camera and `GPSL` timestamps are exactly the same. We further observed an interesting hourly

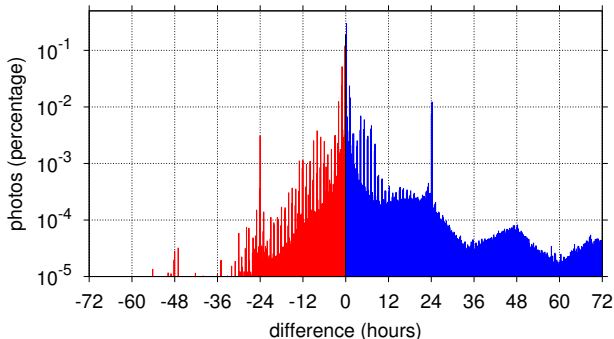


Figure 1: Histogram showing the differences between camera timestamps and `GPSL` timestamps based on a sample of 10 million public Flickr photos. The negative (red) values refer to camera timestamps that are behind the `GPSL` timestamps in time, whereas the reverse is the case for the positive (blue) values. The plot is shown on a logarithmic scale to emphasize the hourly peaks. Each bar represents a 10 minute window.

pattern, suggesting that changes in time zone and daylight savings do affect whether the camera and `GPSL` timestamps correspond with each other, while two other high peaks occur at a difference of 24 hours, suggesting a mistake made by the user in setting the date of the camera clock one day forward or backward compared to the actual date. Overall, we see that in 28% of the cases the camera timestamps are behind the `GPSL` timestamps in time, clearly indicating that these camera timestamps are incorrect, since `GPSL` timestamps either exactly reflect local time or are already a certain amount of time behind local time due to staleness. The other 53% of the time that the camera timestamps are ahead of the `GPSL` timestamps does not allow us to draw any conclusions yet, since this hinges on the correctness of both timestamps with respect to local time.

2.2 Accuracy

We now set out to investigate whether the camera and `GPSL` timestamps actually reflect the local time at the locations where the photos were taken. To do so, we require a ground truth dataset of photos captured at known moments in time. To this end, we filtered our earlier dataset for photos that were associated with the tag `clock` or with the name of a famous landmark that contains a clock, such as the Big Ben in London and the Steam Clock in Vancouver. To avoid a bias towards particular users, cameras and time zones, we retained at most one photo per user+camera+time zone combination. We then visually inspected the remaining photos to see whether they showed a clock and, if so, from the clock face manually coded the actual time of day when each photo was taken. We discarded photos of clocks that could be ambiguously interpreted or when they appeared to be broken, see Figure 2, unless the photo’s title, description, tags or comments suggested otherwise. After performing the three aforementioned filtering steps we ultimately ended up with 750 photos of clocks in our ground truth collection.

We show the results of the comparison between camera timestamps, `GPSL` timestamps and local time in Table 1. As can be seen, while both timestamps can be very accurate, they can also produce errors in the order of hours. Even though the `GPSL` timestamps are slightly more ac-

²<http://efele.net/maps/tz/world/> Accessed 4/2014.

³<http://www.iana.org/time-zones> Accessed 4/2014.



Taken by J



Taken by North Charleston



Taken by cobra



Taken by Dauvit Alexander

Figure 2: Examples of photos inspected for inclusion in the ground truth collection. The time of day of the clock in the first photo can be accurately estimated at 17:00, whereas in the second photo it is not clear whether the time expressed is 7:47 or 19:47. The third clock is shown at an angle that makes it difficult to accurately estimate the time of day, while the last clock seems to be broken. Of these four photos we would only have retained the first one.

Table 1: Difference in minutes between the perceived local time of the clocks in our ground truth collection and the camera and GPS_L timestamps.

	Camera	GPS_L
Minimum	0	0
1 st Quartile	1	1
2 nd Quartile	3	2
3 rd Quartile	61	59
Maximum	668	672
Average	76	65

curate than the camera timestamps, the difference between them is small. Note that due to clocks typically only showing the time but not the date, our current evaluation is limited to assessing the correctness of the timestamps with respect to time of day only, and therefore the differences between the actual and the reported times may in reality amount to several days, weeks, or even months.

3. TIMESTAMP CORRECTION

Our results indicate that the camera and GPS_L timestamps can yield good approximations of the actual local time, but the errors can also be large. Thus, for reliable application in real-world scenarios the accuracy of the timestamps needs to be further improved. Based on a manual inspection of sessions produced by different users and cameras, we often observed cases in which the GPS_L timestamp correctly reflected local time when the camera timestamp did not, and vice versa. We will exploit this observation to automatically generate a corrected timestamp that is able to reduce the error with respect to local time even further. However, we remain with the issue that only a relatively small percentage of photos has GPS metadata. To address this, we exploit the notion that users will typically not just take a single photo, but are likely to take many more within a short timespan of each other—we call such a sequence of photos a *session*. We hypothesize, when considering all sessions that include at least one photo with a GPS_L timestamp, that we can reduce the error of all camera timestamps in a session with respect to local time by propagating the most accurate GPS_L timestamp to all photos; this way we mitigate the effects of any stale GPS_L timestamps by focusing on only the photo that has the most recently updated GPS information.

3.1 Sessions

To form sessions, we first queried the Flickr API for all available photos taken by the same users and their cameras that

provided the photos of the clocks in our ground truth dataset. Per user+camera combination we then aggregated all photos that were captured less than τ seconds apart according to their camera timestamp. While the camera timestamps may or may not accurately represent local time, we assume they do correspond with the order in which the photos were taken and correctly reflect the period of time that passed between two consecutively taken photos. We only focus on sessions containing at least one photo with GPS information, since otherwise we have no point of reference we can exploit for correction. When a session spans multiple time zones—the odds of which increase as τ increases—we split the session into as many sessions as it has time zones and place photos without GPS information in the same session to which its time-wise closest photo with GPS information was assigned.

3.2 Correction offset

When a session contains more than one photo with GPS metadata, we address the potential staleness of this metadata by first determining which photo has the most recently refreshed GPS information. We do so by picking the photo where the difference obtained by subtracting the camera timestamp from the GPS_L timestamp is maximal (i.e. least negative or most positive). Namely, the more recently the GPS information was refreshed the smaller its difference with the camera timestamp, and vice versa; this holds irrespective of whether the camera timestamp correctly reflects the actual time or not. We then use this difference as the correction offset that we will apply to the camera timestamps of all photos in a session. While the most recently updated GPS_L timestamp may still not correctly reflect the time at which the photo was actually captured, it nonetheless provides the optimal time estimate given the circumstances. Our earlier manual inspection of sessions produced by different users and cameras yielded the following features that we will also exploit for producing improved timestamps.

Staleness. When the most recently updated GPS_L timestamp refers to a moment in time that is more than 24 hours behind the camera timestamp, we found that it was more likely for the GPS_L timestamp to be stale than for the user to have accidentally set the camera clock more than one day ahead of the actual time. In such situations we do not apply a correction offset to the photos in the session and simply use the original camera timestamps instead.

Reorder. As part of pre-share photo organization [4], we noticed that users occasionally reordered their photos by manually changing the camera timestamps before uploading

Table 2: Difference in minutes between the perceived local time of the clocks in our ground truth collection and the camera and GPS_L timestamps, as well as our technique using different values of τ to create the sessions.

	Camera	GPS _L	$\tau = -$	$\tau = 1h$	$\tau = 2h$	$\tau = 3h$	$\tau = 4h$	$\tau = 6h$	$\tau = 12h$	$\tau = 24h$	$\tau = 48h$
Minimum	0	0	0	0	0	0	0	0	0	0	0
1 st Quartile	1	1	1	1	1	1	1	1	1	1	1
2 nd Quartile	3	2	2	2	2	2	2	2	2	2	2
3 rd Quartile	61	59	55	50	51	51	51	52	52	58	59
Maximum	668	672	671	665	665	665	665	665	665	665	665
Average	76	65	55	54	55	55	55	56	56	57	57

them to Flickr. Since a manual reordering of photos can lead to an incorrect correction offset, we detect such sessions by ordering all photos in a session by increasing camera timestamps and verifying that the GPS_L timestamps also monotonically increase. When we identify a reordered session, we do not apply a correction offset to the photos.

Shift. We observed that occasionally the GPS timestamps seem to already refer to local time rather than to UTC time. To address this, when the difference between the GPS_L timestamp of a photo and its corresponding camera timestamp is approximately identical to the time zone offset, we subtract the time zone offset from the correction offset.

3.3 Evaluation

We created sessions using different values of τ between 1 and 48 hours to assess the effect of session length on the accuracy of the timestamps that were corrected by our technique. The advantage of sessions created with a large τ is that these will generally contain more photos, such that the timestamps of a larger number of photos can be corrected than is the case for sessions created with a smaller τ . Furthermore, with more photos available, it is more likely that at least one of them will have a very accurate GPS_L timestamp that can be propagated to all photos. Yet, if an inaccurate timestamp is propagated to all photos in a long session, then all of them will end up incorrectly reflecting local time. We also evaluate the accuracy of the corrected timestamps when each photo is assigned to its own session ($\tau = -$), which effectively only applies the **staleness** and **shift** features of our technique.

For each of the sessions containing one of the ground truth clock photos, we compare our manually annotated local time with the camera timestamp, the GPS_L timestamp and the corrected timestamp produced by our technique. We show the results in Table 2. We can see that our technique is able to reduce the errors of the corrected timestamps with respect to local time, on average by approximately 20 minutes compared to the camera timestamps and 10 minutes to the GPS_L timestamps. When each photo forms its own session ($\tau = -$) it already reduces the error observed at the 3rd quartile and the average error, but not the maximum error. A τ of 1 hour yields the lowest errors, whereas an increase of τ yields an increase in error, albeit still lower than those obtained when using the camera or GPS_L timestamps. Yet, the average session length with $\tau = 1h$ was 15 photos, whereas with $\tau = 48h$ the average session length increased to 66 photos, which means that for a given application scenario the optimal τ should be selected based on the trade-off between the desired accuracy of the timestamps and the quantity of timestamps of which the error can be reduced.

4. CONCLUSIONS

We presented an analysis into the accuracy of the camera and GPS timestamps with respect to local time. We observed that 37% of the photos differ more than an hour between their camera and GPS timestamps with respect to local time at the locations where the photos were taken. We further showed that, while the majority of the timestamps closely approximated local time, they on average exhibiting errors of more than an hour. In light of our observations we proposed a simple technique that reduced the errors by on average 20 minutes compared to camera timestamps and 10 minutes compared to GPS timestamps. While we acknowledge our improvements are modest, we achieved the error reductions solely by applying a fully automatic metadata-only technique to a session of photos. In future work we aim to involve textual analysis, visual analysis, and daily sunrise/sunset time tables to further reduce the errors. We additionally intend to obtain photos that have been taken at a known date, time, and location, in order to fully assess the accuracy of camera and GPS timestamps, and how GPS staleness affects capture location accuracy.

5. REFERENCES

- [1] B. Brown, M. McGregor, and E. Laurier. iPhone in vivo: video analysis of mobile device use. In *CHI*, pages 1031–1040, 2013.
- [2] M. D. Choudhury, M. Feldman, S. Amer-Yahia, N. Golbandi, R. Lempel, and C. Yu. Automatic Construction of Travel Itineraries using Social Breadcrumbs. In *HT*, pages 35–44, 2010.
- [3] L. S. Kennedy, M. Naaman, S. Ahern, R. Nair, and T. Rattenbury. How Flickr helps us make sense of the world: context and content in community-contributed media collections. In *MM*, pages 631–640, 2007.
- [4] D. Kirk, A. Sellen, C. Rother, and K. Wood. Understanding photowork. In *CHI*, pages 761–770, 2006.
- [5] M. B. Kjærgaard, J. Langdal, G. Torben, and T. Toftkjær. EnTracked: energy-efficient robust position tracking for mobile devices. In *MobiSys*, pages 221–234, 2009.
- [6] J. Paek, J. Kim, and R. Govindan. Energy-efficient rate-adaptive GPS-based positioning for smartphones. In *MobiSys*, pages 299–314, 2010.
- [7] T. Rattenbury, N. Good, and M. Naaman. Towards automatic extraction of event and place semantics from Flickr tags. In *SIGIR*, pages 103–110, 2007.
- [8] B. Shaw, J. Shea, S. Sinha, and A. Hogue. Learning to rank for spatiotemporal search. In *WSDM*, pages 717–726, 2013.
- [9] United States Department of Defense. *Global positioning system standard positioning service performance standard*. 4th edition, 2008.
- [10] Y. Zheng, L. Zhang, X. Xie, and W.-Y. Ma. Mining interesting locations and travel sequences from GPS trajectories. In *WWW*, pages 1–10, 2009.