The Big Picture: Lessons Learned from Collecting Shared Experiences through Lifelogging

Sarah Clinch[1], Nigel Davies[1], Mateusz Mikusz[1], Paul Metzger[1,3], Marc Langheinrich[2], Albrecht Schmidt[3] and Geoff Ward[4]

Lancaster University[1], University of Lugano (USI)[2], University of Stuttgart[3] and University of Essex[4]

[s.clinch | n.davies | m.mikusz]@lancaster.ac.uk, paul.felix.metzger@googlemail.com, marc.langheinrich@usi.ch, albrecht.schmidt@acm.org, gdward@essex.ac.uk

ABSTRACT

The emergence of widespread pervasive sensing, personal recording technologies and systems for quantified self are creating an environment in which one can capture fine-grained activity traces. Such traces have wide applicability in domains such as human memory augmentation, behavior change and healthcare. However, obtaining these traces for research is non-trivial, especially those containing photographs of everyday activities. To source data for our own work we created an experimental setup in which we collected detailed traces of a group of researchers over 2.75 days. We share our experiences of this process and present a series of lessons for use by other members of the research community conducting similar studies.

INTRODUCTION

Advances in pervasive sensing are enabling us to capture increasingly comprehensive records of users' activities. For example, photo capture technologies such as Microsoft's SenseCam [17] and more recently Narrative Clips [16] can be used to provide photographic records of activities, while GPS tracking and health appliances capture physical motion. The proliferation of such data enables researchers to explore new applications and interaction patterns with computers such as memory enhancement [5, 9] and health-behavior monitoring [13, 18]. However, obtaining pervasive sensing datasets, particularly photographic datasets, is extremely challenging.

For our own research on human memory augmentation we needed a comprehensive dataset for a group of individuals interacting in a shared physical space. In the absence of publicly available datasets we initiated an experiment to capture our own data by placing a group of researchers in a heavily instrumented house. Our experiment ran for 2.75 days and a range of worn and fixed cameras and sensors were used to provide group lifelogging data. To our surprise we found little research literature reporting experiences or methodologies for designing such a temporary experimental set up. Since we believe

that this type of experiment will be increasingly common we report on the data we captured and the lessons we learned from our experiments, and anticipate that these will be of value to those conducting future immersive experiments in pervasive sensing.

Note that we do not claim that our study results are particularly scientifically rigorous – participants were recruited from within the project and hence cannot be considered independent (though research in pervasive sensing and lifelogging does have a long tradition of self-experimentation). However, the lessons and guidelines we present in this paper are not reliant on this aspect of the study, rather they focus on the general challenges of creating immersive experimental environments and thus we believe will have widespread applicability.

RELATED WORK

Our study draws on two parallel avenues of data capture research: lifelogging and instrumented housing. During our study, a range of both worn and fixed cameras and sensors were used to provide group lifelogging data.

Typical lifelogging data originates from a range of sources including dedicated capture devices (e.g. wearable cameras) and digital histories gathered from traditional computing interactions (e.g. web browsing history). Our study is predominantly focussed around image capture—arguably a dominant feature of current approaches. Reviewing lifelogging data offers a chance to reflect on everyday experiences in a manner not easily afforded otherwise, providing the opportunity to draw new insights and learning, informing future behavior. Unsurprisingly, studies of lifelogging experiences are already an important focus for research [9, 10, 11, 19], particularly in the areas of human autobiographical memory [9, 15, 17, 19] and health-behavior monitoring [13, 18].

Lifelogging techniques currently focus on the individual, with photo capture taking place using wearables to capture a first-person view. As a complimentary approach, we have previously proposed the use of infrastructure cameras to capture a third-person view in lifelogging images [3]. Such an approach to lifelogging may draw on lessons learned in instrumented housing research in which the creation of sensor-rich environments allows the study of user activity [12, 14]. Study of these smart environments has identified lessons informing the design of these future spaces including awareness of both social (e.g. privacy, preventing disruption) and technical challenges (e.g. retrofitting of spaces, making inferences about behavior) [6]; our work complements these lessons by focusing on the design of successful trials and studies. The majority of instrumented housing research has been dependent on established pieces of fixed infrastructure (one notable exception was the portable investigation kit used by researchers at MIT in the early 2000s [1]). By contrast, our work focuses on the process of running temporary experiments in lifelogging. Creating temporary immersive environments for studying subjects in-situ has

also appeared in the popular media in the form of the "Big Brother" global franchise.

Our work combines both wearable and infrastructure-based approaches, and represents a novel attempt to document the process of conducting an experiment to collect contemporary lifelogging data from multiple co-located participants.

THE GRASMERE HOUSE EXPERIMENT

Methodology

Location Selection

An early challenge in our study design was the selection of an appropriate location to conduct the experiment. This proved to be a non-trivial task as we needed to reconcile multiple distinct requirements:

1. relatively accessible for international travellers,

2. capable of hosting ~20 participants,

3. remote enough to minimize capture of bystanders,

4. served by good communications infrastructure, and,

5. willingness by the owner to allow deployment of a large number of technologies during the experiment.

We had initially focused on self-catering cottages which provide remoteness but rarely cater for the required numbers and are typically served by poor communications infrastructure. In contrast, hotels offer scale and better infrastructure but are almost impossible to instrument without capturing (and requiring consent from) hotel staff.

Our compromise was a small hostel in the village of Grasmere, in the Lake District National Park in the UK. The hostel allowed housing of 24 people in relative comfort yet provided the possibility for researchers to instrument the environment and to engage in household activities such as cooking. A remote location and guarantee of exclusive use allowed deployment of cameras without inadvertently capturing bystanders.

Experimental Setup

The hostel featured five shared bedrooms and communal facilities including two common rooms, kitchen/dining areas, a large bathroom and shower area, and a conference room. Bedrooms, bathrooms and shower areas were designated (and clearly signed) private with no lifelogging permitted in these spaces [Figure 1].

We aimed to simulate a setting in which capture devices are ubiquitous within the house. All public spaces were video- and audio-taped using IPTV cameras, and instrumented with iBeacons for gathering proximity data from nearby mobile phones. Workstations were provided for central data collection and post-processing. In addition to these fixed devices, all participants were given a wearable lifelogging camera that automatically captured images and GPS data every 30 seconds; participants were requested to wear these throughout the study except in the designated private areas. Participants were also invited to use any additional lifelogging devices they wished and contribute the data to the experimental dataset (this included 2 additional wearable cameras, 4 GoPro video capture devices, and some smartphone activity trackers); installation of the iBeacon location tracking smartphone application was also optional.

During the 2.75 days of study, participants engaged in a variety of tasks following a prepared schedule; this allowed us to assemble a wide variety of lifelogging data in order to understand the challenges collecting and analyzing data from different everyday activities. Household tasks (e.g. cooking, eating, cleaning) were supplemented with additional activities such as hiking, jogging, meetings, and workshops.

Figure 1: Hostel floor plan. Bedrooms were located on the second (b) and third (c) floor, and shared social spaces on the first (a) and third (c) floor. Areas marked green were recorded throughout the study; red areas were deemed "private".

Subject Recruitment

We recruited 21 participants – each agreed to participate in some portion of the study, with the majority staying for the full duration; participants were not financially compensated for their participation. Our aim was to gain an understanding of how to conduct an experiment in this space, and to produce sample datasets rather than to obtain scientifically rigorous experimental data. As a result, we essentially "experimented on ourselves" – continuing a long science tradition – with all participants recruited from within the project. Participants were briefed on the aims and objectives of the study prior to participation. Reflecting their research background, participants were actively encouraged to contribute their own experiments to the schedule.

Results

During the experiment we collected 42,959 images (34.9 Gigabytes) and 248.15 hours of video (245.2 Gigabytes). We gathered indoor location data from four participants (using iBeacons and compatible smartphones that ran a small custom app) and GPS data from all participants (our worn lifelogging cameras all included a portable GPS tracker, and some participants also collected traces using smartphones). Location data was, depending on the type of device, stored as either plain metadata (indoor) or raw GPS signal strength data.
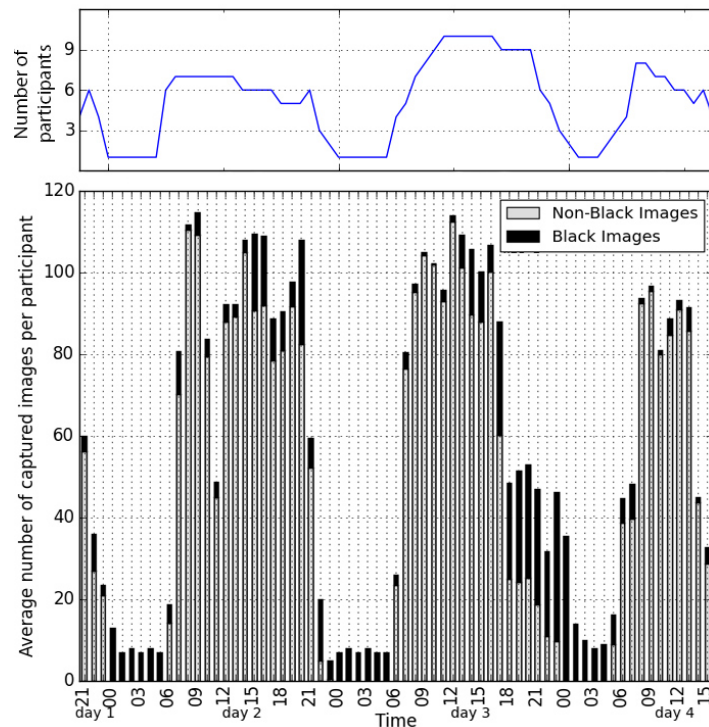

Figure 2: Data captured from wearable cameras per hour of study.
Top: number of participants capturing images per hour.
Bottom: average number of images per participant per hour.

KEY ISSUES

Based on our experiences, we now consider in detail three important questions relating to the value of the different types of data collected and the ethical implications of our work.

What is the value of wearable camera data ?

During the study 13 participants used a wearable camera for at least some portion of the time—however few supplied complete sets of data from their device. This was a result of a number of factors including: forgetting to wear or failing to charge the device, and forgetting to reattach the device when leaving a private space. Given our experimental duration of 42.5 waking hours (66.5 hours total), 13 cameras, and a capture interval of 30 seconds, the theoretical maximum number of worn camera images was ˜66,300. In practice we recorded just 26,218 readable images (16.21 GB), 40% of this maximum. The median number of images contributed by each participant was 1,762 (IQR 2,638-1,225) corresponding to 14.68 hours of capture [Figure 2]. 3,228 (12.31%) of the captured images were programmatically identified as "black" – corresponding to photographs taken when the device was covered or face down. The remaining images had varying levels of utility – some were totally or partially occluded (e.g. by the wearer's arm) while many were shots of irrelevant background objects (e.g. the ceiling or floor). In addition, the photos inevitably incorporated considerable duplication.

To gain an understanding of the overall utility of the photos captured by the wearable cameras we manually analyzed a sample image set selected by dividing the first full day of the study into 2 hour blocks (i.e. 00:00-01:59, 02:00-03:59 etc.) and taking up to 5 images, two minutes apart from between 30 and 39 minutes past the first hour. This generated a sample of 193 images from 9 participants whose devices captured at least some of the sampled time. Unsurprisingly, this sample was dominated by images captured during the day and early evening: 4 (2.07%) of the sampled images were captured between 10pm and 5:59am; 90 (46.63%) between 6am and 1:59pm; and 99 (51.30%) between 2pm and 9:59pm). We programmatically identified 22 black images– these images were distributed throughout the day but included all 4 of the images between 10pm and 5:59am (18.18% of the total set of 'black' images)–this left a final sample of 171 non-black images.

We began by judging the extent to which each image was occluded (measured in 25% increments), whether the image was blurred and if the image duplicated an immediately preceding image. We found 4 duplicates (2.34% of sample) and 11 blurred images (6.43%). In 41 images (23.98%) at least 25% of the image was considered to be occluded, 10 images (5.85%) were at least 50% occluded, 6 (3.51%) at least 75% occluded, and 3 completely occluded (1.75%).

To understand the content we considered, for each image, the people, places, objects and actions that the image contained. This distinction was motivated by the classification

scheme used by Lee and Dey (2007) [15]. Reflecting the scheduled activities (meetings, hiking, general living) we recorded a large number of personal devices, household objects and items of food and drink. Self-occlusion and problems with angling the wearable cameras were reflected in 95 instances of body parts in the images. A total of seven actions were identified ('cooking', 'driving', 'eating', 'sitting', 'walking', 'working', 'writing') accounting for 136 images (79.53%); for the remaining 35 photographs, no action could be identified. Reflecting the scheduled programme, the most common activities identified included walking (47 images), eating (38 images) and working (30 images). The majority of blurred images were associated with 'walking' (7 of 11 blurred images). A total of 116 images were tagged as being indoor; 52 outdoors; 3 unknown.

The Grasmere House experiment was distinctive in that it included multiple co-located users of wearable cameras. Such data enables new forms of experiments and analysis. For example, to explore relationships between data captured by multiple users, we considered the specific case of reciprocity—in which two participants take photographs of each other, typically while engaged in a conversation. 142 of the sampled images contained some identifiable portion of another participant. For each individual in these photographs we extracted a larger sample from the featured individual's wearable camera data for ten minutes either side of the timestamp of the original photograph (1 image per minute). We examined this new sample to identify photographs containing the original photographer [Table 1]. In approximately 10% of cases, individuals photographed by a worn camera recorded for themselves an image of the photographer within 10 minutes, but over a shorter period few reciprocal relationships are seen (<1% within a period of 2 minutes) [4]. Examination of the sample images demonstrated that group interactions were more complex than the reciprocal relationships represented in Table 1; Participant A's camera might capture Participant B, but Participant B would capture Participant C, who in turn may then capture Participant A.

Time Interval	Number of images showing reciprocity	
+-1 minute	1	(0.70%)
+-2 minute	1	(0.70%)
+-5 minute	7	(4.93%)
+-10 minute	14	(9.86%)

Table 1: Personal image reciprocity (based on an image sample of 142 images containing some identifiable portion of another participant). Over 10 minutes, approximately 10% of these lifelogging images featuring another participant are matched by a corresponding image taken by the lifelogging device of the person captured containing the original photographer.

In determining the overlap between images produced by wearable cameras there are a number of factors to consider. What we ideally wished to understand was the extra value that each additional wearable camera brought to a scene in terms of its ability to record new information to aid in subsequent interpretation of the images [Figure 3]. In practice this proved extremely hard to quantify as the figures were heavily influenced by factors such as the size of the space, the duration of the experiment and the relative mobility of the participants.

What is the value of infrastructure cameras ?

In contrast to traditional lifelogging research, we wanted to understand how data from multiple viewpoints – including fixed infrastructure – can be used to improve an individual's recollection of events. Our experimental setup therefore included five fixed video cameras and two time-lapse cameras. Given an experimental duration of 2.75 days the theoretical amount of data we would expect to record is 67,050 minutes. In practice we recorded 14,227 minutes of video from the five fixed cameras, and 180 minutes of video from the time-lapse camera (total 159.208 GB). (An additional 480 minutes of video captured was from worn Go-Pro devices). The shortfall was due to occasional failures of the fixed cameras (requiring reboots) and network issues that meant that the camera streaming rate was much lower than expected (camera specification was 15fps vs. an average of 4.04fps in the field). As shown in Figure 4 the number of frames per second correlates to the file size due to compression used by the cameras (many movements in the video results in a large size which might have a negative impact on the frames per second if the network is being overloaded).

These cameras provided a radically different set of images to those produced by wearable cameras [Figure 5]—notably, images from the infrastructure cameras were all usable (i.e. not blurred or occluded). However, initial setup was needed to frame the image captured by infrastructure cameras (e.g. to choose appropriate viewing angles and ensure good coverage of the hostel).

Figure 3: Examples of ten different perspectives on an event as captured by the participants' wearable cameras. All participants have given consent for publication of these photographs.
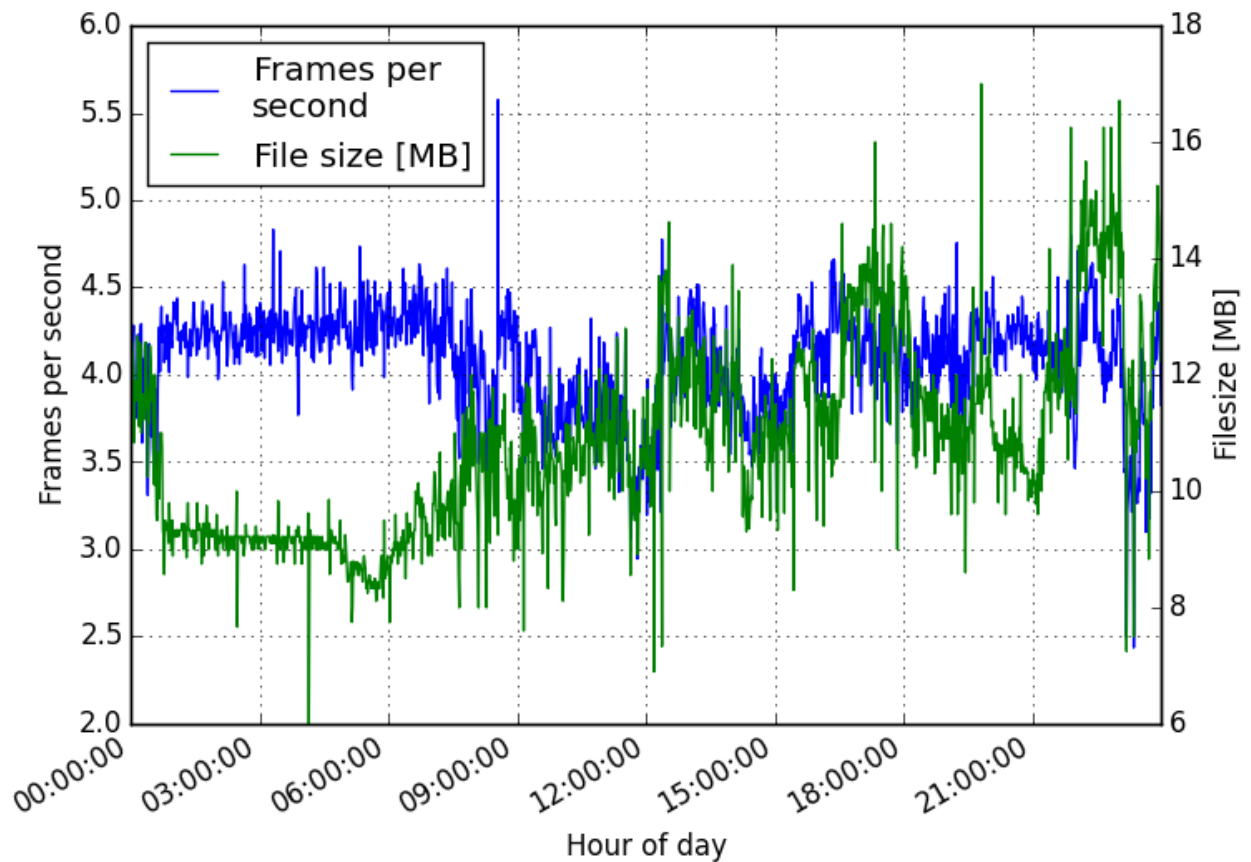

Figure 4: Frame rate and file size for fixed infrastructure cameras (aggregated values) grouped by time of day.

Figure 5: Infrastructure camera photos (top) vs. wearable camera photos. All participants have given consent for publication of these photographs.

What Privacy Issues are Likely to Arise ?

We were interested in understanding the social and ethical issues of conducting a study such as the Grasmere House, and to learn first-hand how it felt to participate in this type of experiment. At the outset, participants were provided with information sheets and consent forms indicating that data would be shared within the project but that consent would be sought before publication of photographs featuring identifiable individuals. Participants could withdraw at any time and could withdraw any data they had contributed but not images of themselves captured by others or by the fixed infrastructure. This caveat is illustrative of a unique challenge of this kind of study, in which data can be considered to have multiple owners and stakeholders. Since participants were recruited within the project we were conscious of the possibility of a sense of obligation to participate; we took care to stress that participation was entirely optional and our participants had no problems with the acceptability of our ethics process.

The hostel was divided into public and private spaces [Figure 1]. Private spaces (bedrooms, bathrooms) were not instrumented and were clearly signed to indicate that the use of recording devices was forbidden. Participants were frequently reminded of the need to respect the different types of space and we had expected that this signage would be sufficient to remind people. Nevertheless many participants reported failures to remove their personal recording device when entering these areas. The shared nature of these spaces made this particularly problematic. Discussions with long-term lifeloggers [2] indicate that the likelihood of violating the rules of private space does not reduce with experience and hence technical solutions (e.g. [7]) may be required to provide bystanders with the levels of privacy that they might reasonably expect. Reflecting on our previously described wearable camera image sample however, we note that only 0.60% of our

captured images could actually be considered to violate privacy by depicting private areas with an additional 0.30% featuring bystanders/non-participants.

A second problem that occurred repeatedly during our study was that lifelogging devices captured the screens of participants' computers and phones (seen in 12.54% of our sampled worn camera images). For most participants this was both unexpected and unacceptable – particularly given that the resolution of the images was sufficient that individual emails or messaging sessions could be read. Crucially, both the capture of private spaces and personal screens made participants more concerned about sharing their photo streams (concerns about sharing images with personal computer screens have also recently been observed by Hoyle et al. [10]). Technological solutions such as automated blurring of laptop and computer screens could in principle address this latter issue, yet in many cases screens may be precisely the type of information that a user may wish to capture.

A final concern related to the capture of audio. As a group we initially believed that the fixed cameras captured only images. During the study it became clear that they also captured audio; this discovery gave rise to significant concerns amongst participants. Similar privacy concerns for audio recordings have been seen in prior work (e.g. [8, 12]); during trials of Georgia Tech's Personal Audio Loop [8] this concern was seen to diminish over time and with further consideration. During the Grasmere House participants' concerns about audio were not sufficient to impact their participation nor for anyone to propose restricting this audio capture.

Our participants could review their lifelogger images prior to submitting them to the shared repository. However, providing subjects with realistic opportunities to filter their images posed considerable challenges. We had explicitly set aside time for participants to review images and remove those that they did not wish to share. Yet, despite having over two hours in which to review a single day's images, most participants made little headway; many abandoned the task within an hour. This highlights the difficulty of relying on users filtering captured data and suggests that solutions that *(i)* automatically filter data, or *(ii)* those that restrict capture, may be more appropriate. Prior work indicates that automatic filtering of lifelogging photos based the presence of faces can be effective at reducing bystander privacy risks, as can cropping techniques that remove the bystanders but keep the remaining image content [18]. However, privacy infringements vary with person and context, and in many cases those very items that pose a privacy risk are themselves most valuable as a memory trigger (e.g. faces, computer screen content). For situations where privacy risks can be clearly defined, automatic discontinuation of lifelogging capture seems like a promising solution (e.g. [7]), and for less clearly defined privacy criteria manual (user) in-situ control over capture may offer greater flexibility. Unsurprisingly Hoyle et al. [11]'s recent study suggested that users preferred in-situ marking of photos for deletion rather than post-hoc processing and removal. Unlike Hoyle we did not explicitly ask our participants whether they would be prepared to share

the images captured, but Hoyle's findings that 92.4% of lifelogging images were deemed shareable with diverse groups of people [11] is comparable with our finding that ~13% of images contained some form of objective privacy issue.

LESSONS LEARNED

So what lessons can be learned from our experiences of running the Grasmere House? One of our most striking observations is that the quality of data captured by mobile devices is, in general, poorer than expected. This can be attributed to a number of factors – some of which can be addressed, and some that are inherent in the current generation of devices. In the first category, our participants were very lackadaisical in dealing with wearable devices. As a result, our data contained significant holes because participants simply forgot to wear devices or had let them discharge. This issue of charging occurred with all of the mobile devices and appeared to be independent of the actual battery life of the device. Of course this may be an issue of perceived value by the participants – after all most were able to keep their mobile phones charged. However, there is an important difference – a mobile phone can still be used while it is charging whereas this is generally not possible with a lifelogging device and hence discharged lifelogging devices lead to significant loss of data capture.

More problematic is the fact that for many worn camera images there was typically little correlation between what participants *thought* they were capturing and what was in practice captured. This was the result of a range of factors including occlusion, blurring, poor camera angle, and a limited field of view. Although devices with a larger field of view exist, such images feel less natural due to image distortion. Problems of occlusion, blurring, and poor camera angle are more intrinsic to the nature of these devices and are difficult to overcome. From these experiences with mobile lifelogging devices we generate our first two lessons: *(L1) researchers' expectations for mobile data sources must be carefully managed, as data is likely to be of a poorer quality than expected* and *(L2) charging solutions should be provided that lower effort for overnight charging (e.g. wireless charging) or that allow subjects to continue logging during charging.*

A further insight comes from having multiple lifeloggers sharing a single space. Lifelogging studies to date have focused on isolated individuals, but as such devices become more common we expect to see instances in which wearers are co-located. Understanding the benefits of combining individuals' datasets to improve capture of shared experiences are not yet well understood but we have used reciprocity as case study of how such combined datasets may be used to gain new insights. We therefore suggest that *(L3) designing studies with co-located lifeloggers offers opportunities for richer data capture*.

In contrast to mobile devices, our fixed infrastructure provided much better data than

expected, and images were of a consistently high quality. However, the infrastructure did require time to install and calibrate (e.g. to ensure the good coverage of the space) and made considerable demands on the communications network. We saw problems with the network throughout the study, despite investing time in setting up what we considered to be an appropriate solution. Thus, we note that *(L4) fixed infrastructure can be harnessed to provide high quality ground truth* and *(L5) the infrastructure requirements for lifelogging experiments increase rapidly*.

The demands on our infrastructure were further highlighted when we endeavoured to start manipulating the captured data. In order to create a comprehensive lifelogging dataset, we copied data from devices to a single shared storage system. Due to the heterogeneity of devices used, some with proprietary communication protocols, participants were required to first use a variety of tools to export data from their devices to local storage (e.g. laptop), from which they would then copy it to the shared storage. This process was time consuming, particularly during periods of congestion on the wireless network (e.g. when multiple participants transferred their data simultaneously); these periods also had implications for capture since during upload no meaningful data could be gathered as the devices were removed from their positions and were busy with the task of data transfer. Real-time syncing of this data could have reduced the likelihood of these congestion periods (and, would also provide a means for introducing centralised timestamps [1]). Finally, for devices that allowed direct access to their data, participants typically used different ways of storing and structuring the data, despite being encouraged to follow a pre-defined structure. Our sixth lesson is thus that *(L6) automation and standardization are key for centralized data storage – agreement on metadata, data exchange formats, and toolchains needs to take place well ahead of the experiment to provide an opportunity for development and testing of appropriate tools*.

Most of the lifelogging devices used provided metadata for location and time of capture. However, we found limitations in usage of the metadata. For example, in an attempt to conserve battery the worn cameras enriched each image with raw sensor readings from a GPS chip—this data could only be resolved to actual coordinates by uploading images and metadata to a 3rd party cloud service. Another issue occurred with incorrectly configured devices, (e.g. DSLRs and GoPros with incorrect date settings); this significantly complicated data analysis, as timestamps had to be manually synchronized across devices. Similarly, merging multiple data streams would had been greatly simplified if we had been able to clearly label people, places and things at the time of capture. For example, even something as simple as having all participants wear distinctively colored t-shirts would have improved the manual tagging process. Labeling wearable cameras so that individual devices (and thus the person wearing them) could be identified in photographs would have been a big help in testing for reciprocity. Our seventh lesson is *(L7) care should be taken to ease the task of post-experiment data processing – particularly through ensuring good quality meta-data and through annotation of the physical space*.

Our work on privacy appears to suggest conflicting courses of action. On the one hand it is clear that participants are poor at adhering to instructions on where lifelogging equipment can be used; most had anecdotal examples of having worn their devices in private areas. However, data analysis revealed few incidents of devices actually capturing such private settings. Much more common was the inadvertent capture of personal screens (laptops, phones). While our sample was potentially pre-filtered by participants, the quantity of privacy violations suggest that this filtering was, in practice, quite minimal. In light of our experiences with privacy issues we suggest the following: *(L8) researchers should recognize that participants are unlikely to comply with "take off your logging tool" signage and must be prepared for privacy incidents* and *(L9) researchers should advise participants to avoid capturing screens with sensitive data.*

Our final lesson relates to the logistics of running data capture experiments. We found that having an appropriate location was critical for our experiments, yet obtaining this location was extremely challenging. In particular, the need to balance the desire for remoteness with quality infrastructure greatly limits one's options. Equally, selection of an appropriate venue impacts the nature of activities than can be studied. Our tenth lesson is thus *(L10) researchers should budget significant time and money to source an appropriate location and ensure that infrastructure is over-specified for the planned experiments.*

CONCLUSIONS

Accurate lifelogging datasets are increasingly important for researchers. Stringent ethics constraints and specific requirements mean that traces produced as part of a research project can rarely be shared beyond that project. As a result, many researchers are collecting their own traces. However, there is little in the way of experience reports from researchers involved in this collection process.

In this paper we report on our experiences of running an immersive data capture experiment and provide ten key lessons that we hope will prove useful to future researchers in the area.

ACKNOWLEDGMENTS

REFERENCES

1 Beaudin, J., Intille, S., and Tapia, E. M. Lessons learned using ubiquitous sensors for data collection in real homes. In *CHI '04 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '04, ACM (New York, NY, USA, 2004), 1359–1362.

2 Billinghurst, M., Davies, N., Langheinrich, M., and Schmidt, A. Augmenting human memory - capture and recall in the era of lifelogging. **http://www.dagstuhl. de/en/program/calendar/semhp/?semnr=14362**.

3 Clinch, S., Metzger, P., and Davies, N. Lifelogging for 'observer' view memories: An infrastructure approach. In *Proceedings of 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing Adjunct Publication*, Ubicomp '14 Adjunct (2014), 1397–1404.

4 Clinch, S., Mikusz, M., and Davies, N. Understanding interactions between multiple wearable cameras for personal memory capture. In *Proceedings of the 16th International Workshop on Mobile Computing Systems and Applications*, HotMobile '15, ACM (New York, NY, USA, 2015).

5 Davies, N., Friday, A., Clinch, S., Sas, C., Langheinrich, M., Ward, G., and Schmidt, A. Security and privacy implications of pervasive memory augmentation. *Pervasive Computing, IEEE 14*, 1 (2015).

6 Edwards, W. K., and Grinter, R. E. At home with ubiquitous computing: Seven challenges. In *Proceedings of the 3rd International Conference on Ubiquitous Computing*, UbiComp '01, Springer-Verlag (London, UK, UK, 2001), 256–272.

7 Fan, M., Adams, A. T., and Truong, K. N. Public restroom detection on mobile phone via active probing. In *Proceedings of the 2014 ACM International Symposium on Wearable Computers*, ISWC '14, ACM (New York, NY, USA, 2014), 27–34.

8 Hayes, G. R., Patel, S. N., Truong, K. N., Iachello, G., Kientz, J. A., Farmer, R., and Abowd, G. D. The personal audio loop: Designing a ubiquitous audio-based memory aid. In *Proceedings of the 6th International Conference on Human Computer Interaction with Mobile Devices and Services*, Mobile HCI 2004. Springer, 2004, 168–179.

9 Hodges, S., Williams, L., Berry, E., Izadi, S., Srinivasan, J., Butler, A., Smyth, G., Kapur, N., and Wood, K. Sensecam: A retrospective memory aid. In *Proceedings of the 8th International Conference of Ubiquitous Computing (UbiComp 2006)*, Springer Verlag (September 2006), 177–193.

10 Hoyle, R., Templeman, R., Anthony, D., Crandall, D., and Kapadia, A. Sensitive lifelogs: A privacy analysis of photos from wearable cameras. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, ACM (New York, NY, USA, 2015), 1645–1648.

11 Hoyle, R., Templeman, R., Armes, S., Anthony, D., Crandall, D., and Kapadia, A. Privacy behaviors of lifeloggers using wearable cameras. In *Proceedings of 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*,

Ubicomp '14, ACM (New York, NY, USA, September 2014), 571–582.

12 Intille, S. S., Larson, K., Tapia, E. M., Beaudin, J. S., Kaushik, P., Nawyn, J., and Rockinson, R. Using a live-in laboratory for ubiquitous computing research. In *Proceedings of the 4th International Conference on Pervasive Computing*, PERVASIVE '06, Springer-Verlag (Berlin, Heidelberg, 2006), 349–365.

13 Kerr, J., Marshall, S. J., Godbole, S., Chen, J., Legge, A., Doherty, A. R., Kelly, P., Oliver, M., Badland, H. M., and Foster, C. Using the sensecam to improve classifications of sedentary behavior in free-living settings. *American Journal of Preventive Medicine 44*, 3 (2013), 290–296.

14 Kidd, C. D., Orr, R., Abowd, G. D., Atkeson, C. G., Essa, I. A., MacIntyre, B., Mynatt, E., Starner, T. E., and Newstetter, W. The aware home: A living laboratory for ubiquitous computing research. In *Cooperative Buildings. Integrating Information, Organizations, and Architecture*, N. A. Streitz, J. Siegel, V. Hartkopf, and S. Konomi, Eds., vol. 1670 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 1999, 191–198.

15 Lee, M. L., and Dey, A. K. Providing good memory cues for people with episodic memory impairment. In *Proceedings of ASSETS '07*, ACM (2007), 131–138.

16 **http://getnarrative.com/**.

17 Sellen, A. J., Fogg, A., Aitken, M., Hodges, S., Rother, C., and Wood, K. Do life-logging technologies support memory for the past?: An experimental study using sensecam. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '07, ACM (New York, NY, USA, 2007), 81–90.

18 Thomaz, E., Parnami, A., Bidwell, J., Essa, I., and Abowd, G. D. Technological approaches for addressing privacy concerns when recognizing eating behaviors with wearable cameras. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '13, ACM (New York, NY, USA, 2013), 739–748.

19 Woodberry, E., Browne, G., Hodges, S., Watson, P., Kapur, N., and Woodberry, K. The use of a wearable camera improves autobiographical memory in patients with alzheimer's disease. *Memory* (2014).

AUTHORS

Sarah Clinch is a post-doctoral researcher at Lancaster University, UK, where she also completed her PhD. Sarah's research interests include pervasive displays and personalisation in ubiquitous computing systems.

Nigel Davies is a Professor in the School of Computing and Communications at Lancaster University. His research focuses on experimental mobile and ubiquitous systems. Nigel is active in the research community and has co-chaired both Ubicomp and MobiSys conferences and is a former editor-in-chief of IEEE Pervasive Magazine. He is currently the coordinator of Recall -- a large multinational project that is exploring new technologies for augmenting human memory.

Mateusz Mikusz is a PhD student in the School of Computing & Communications at Lancaster University, UK. His research focusses on pervasive analytics, mainly in the context of digital signage.

Paul Metzger was a visiting researcher at Lancaster University and student at the University of Stuttgart. He studied Software Engineering and has a strong interest in Human Computer Interaction.

Marc Langheinrich is an Associate Professor at the Università della Svizzera italiana (USI) in Lugano, Switzerland, where he works on privacy and usability in pervasive computing systems. Marc has served as General Co-Chair for PerCom 2012 and Program Co-Chair for UbiComp 2013. He is on the editorial board of IEEE Pervasive, Elsevier's Personal and Mobile Communications, and Dagstuhl's Open Access Series in Informatics. Marc holds a PhD from ETH Zurich. He is a member of IEEE and ACM and can be reached at langheinrich@ieee.org.

Albrecht Schmidt is a professor of human-computer interaction and cognitive systems at the University of Stuttgart. His research interests are at the intersection of ubiquitous computing and human-computer interaction, including large-display systems, mobile and embedded interaction, and tools to augment the human mind.

Geoff Ward read Natural Sciences at the University of Cambridge, and then completed his doctorate in Experimental Psychology at the University of Oxford. Geoff was appointed Lecturer in Psychology at the University of Essex in 1993, where he is currently a Professor. He is a member of the British Psychological Society and the Experimental Psychology Society and an Associate Member of the Psychonomics Society. His current research includes encoding and retrieval processes in human memory and augmenting human memory through technology.