IAPMA 2011

2nd Workshop on Information Access for Personal Media Archives

Organised by:
Aiden Doherty
Kieron O'Hara
Kiyoharu Aizawa
Niamh Caprani
Alan Smeaton

Workshop of the 33rd Annual BCS-IRSG European Conference on Information Retrieval

ECIR 2011
18-21 APRIL 2011
DUBLIN, IRELAND
THE 33RD EUROPEAN CONFERENCE ON INFORMATION RETRIEVAL

Information Access for Personal Media Archives

At ECIR 2011, 18th April 2011 CLARITY

Towards e-Memories: challenges of capturing, summarising, presenting, understanding, using, and retrieving relevant information from heterogeneous data contained in personal media archives.

Welcome to IAPMA 2011, the second international workshop on “Information Access for Personal Media Archives”. It is now possible to archive much of our life experiences in digital form using a variety of sources, e.g. blogs written, tweets made, social network status updates, photographs taken, videos seen, music heard, physiological monitoring, locations visited and environmentally sensed data of those places, details of people met, etc. Information can be captured from a myriad of personal information devices including desktop computers, PDAs, digital cameras, video and audio recorders, and various sensors, including GPS, Bluetooth, and biometric devices.

This year’s workshop highlights the latest research on advancements towards the goal of effective capture, retrieval and exploration of e-memories. Dr. Chaminda De Silva and colleagues from The University of Tokyo and Microsoft Research Asia provide an overview on a visual lifelog dataset of 4,179 time and location-stamped images, which will prove a valuable resource for the personal media archives community (page 4). Yi Chen and colleagues from Dublin City University consider the challenges in constructing a new event segmentation algorithm to identify high-level abstracted events across longer time periods in lifelog datasets spanning multiple years (page 8). Saskia Koldijk and colleagues from Radboud University Nijmegan and TNO in Delft consider the various forms of personal media that knowledge workers are confronted with and discuss novel activity logging techniques to provide feedback to these knowledge workers with the goal of aiding self-coaching. (page 12). Karen Lee and colleagues in the University of Ulster consider ambient assisted living scenarios and present the development of an ambient framework to augment and extend the capabilities of the current OpenAAL platform in recognising people activities (page 16). Dr. Jochen Mayer and Benjamin Poppinga from Oldenburg University recognise the significance of healthcare applications with respect to personal media archives, and in particular automatically deriving health related features from web scale distributed personal archives (page 20).

Organising Committee

Aiden R. Doherty (Oxford University)

Kieron O’Hara (Southampton University)

Kiyoharu Aizawa (University of Tokyo)

Alan F. Smeaton (Dublin City University)

Niamh Caprani (Dublin City University)

www.iampa2011.dcu.ie **Published at http://doras.dcu.ie/16291/** **ISBN: 1872327974**

Workshop Program

Location: Guinness Storehouse, Dublin, Ireland

Opening Session	**2.00 – 3.30 pm**
	Opening Remarks
Chaminda De Silva, Kiyoharu Aizawa, Yuki Arase and Xing Xie	*Capturing, Using and Sharing a Liflog Dataset (Page 4)*
Yi Chen, Gareth Jones and Debasis Ganguly	*Segmenting and Summarizing General Events in Long-term Lifelogs (Page 8)*
Saskia Koldijk, Mark Van Staalduinen, Stephan Raaijmakers, Wessel Kraaij and Iris Van Rooij	*Activity-logging for self-coaching of knowledge workers (Page 12)*
Afternoon Coffee/Tea Break	**3.30 – 4.00 pm**
Session 2	**4.00 pm – 5.00 pm**
Karen Lee, Tom Lunney, Kevin Curran and Jose Santos	*Context-Aware Support for Assistive Systems and Services (Page 16)*
Jochen Meyer and Benjamin Poppinga	*Distributed Personal Archives: Deriving Health Profiles from the Web (Page 20)*
	Closing Remarks

Program Committee

Yuki Arase (Microsoft Research Asia, China)

Jenny Benois-Pineau (LABRI, Bordeau, France)

Kevin Curran (University of Ulster, N. Ireland)

Chaminda De Silva (University of Tokyo, Japan)

Lothar Fritsch (Norwegian Computing Center)

Jim Gemmell (Microsoft Research, USA)

Fausto Giunchiglia (Uni. Trento, Italy)

Cathal Gurrin (Dublin City University, Ireland)

Josef Hallberg (Luleå University of Technology)

Alexander Hauptmann (CMU, USA)

Wilko Heuten (OFFIS, Oldenburg, Germany)

Jeremy Leighton John (British Library, U.K.)

Vaiva Kalnikaité (GE Digital Energy, U.K.)

Birger Larsen (RSLIS, Copenhagen, Denmark)

Siân Lindley (Microsoft Research Cambridge)

Claire Lobet (University of Namur, Belgium)

Tom Lunney (University of Ulster, N. Ireland)

Daniela Petrelli (Sheffield Hallam Uni, U.K.)

Mike Sinclair (Microsoft Research, USA)

Kristin Tolle (Microsoft Research, USA)

Capturing, Using and Sharing a Lifelog Dataset

Chaminda de Silva and Kiyoharu Aizawa
Interfaculty Initiative in Information Studies
The University of Tokyo
(chamds, aizawa)@hal.t.u-tokyo.ac.jp

Yuki Arase and Xing Xie
Microsoft Research Asia
Beijing, China
(yukiar, xing.xie)@microsoft.com

ABSTRACT

We present a large collection of location based lifelog data that is publicly available for research. This dataset contains location data and digital photos captured by one of the authors over a period of more than a year. The location data were captured by continuously carrying a GPS receiver and recording the location coordinates at all times when they can be estimated. The photos, 4179 in total, were taken during sightseeing and other daily-life events.

We first describe our effort and experience in capturing the data. We also outline how the data have been used in both personal and research capacity, with a brief description of the systems used. After describing how the data were prepared to be shared as an open dataset, we reflect on the overall experience and conclude the paper with possible future directions.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous; H.3.3 [**Information Storage and Retrieval**]: Information Search and Retrieval

Keywords

Personal Multimedia Archives, Lifelog, Multimedia Retrieval

1. INTRODUCTION

With recent advances in sensing and storage technologies, it is now possible to continuously or regularly capture and store data regarding a person's life. Such data collections are often called *Lifelog* data. Researches related to lifelog data focus on diverse objectives such as more effective acquisition [8], fast and accurate retrieval of data [3], and easy-to-interpret visualization of experiences and events [3].

The most important step in a lifelog related research project is data acquisition. Capturing data for a long time is essential. This requires a large amount of effort since it is necessary to continuously carry and power the sensing devices, while dealing with privacy and security issues.

There are several research projects that work on large, multimodal datasets [5, 6, 3]. However, due to the highly personal nature of the collected data, researchers are not so willing to share their data with others. This makes it hard to compare the performance of algorithms developed by different researchers. A few moderately sized and non-continous datasets, such as GeoLife [9], have been shared by some researchers. However, a large, continuously archived lifelog dataset is not publicly available at the moment.

In this paper we introduce an open multimedia dataset for lifelog research. This dataset spans a year of a person's life. It contains continuously archived location data (approximately one million GPS data points), and around 4200 digital photos taken during travel and events. The dataset is available for download at http://www.hal.t.u-tokyo.ac.jp/~chamds/pics_n_trails/. It is open for use in research projects, under a license agreement. To the best of our knowledge, this is the largest continuously-archived, publicly-available lifelog data collection. The rest of this paper describes how these data were captured, how they have been used, and the effort required for sharing them.

2. CAPTURING

The first author (hereinafter referred to as "the lifelogger") decided to create a lifelog data archive using location data and images. The location data were continuously captured using a GPS receiver. While several types of wearable cameras are available for continuously or regularly capturing image data, they are seen as an intrusion of other people's privacy. Therefore, we decided to use ordinary digital cameras for image acquisition.

2.1 Location Data

The main criteria for selecting a GPS receiver were high sensitivity, high accuracy, and the ability to store a large amount of data. After testing a few models of GPS receivers, we selected a Garmin® *GPSmap60CSx* GPS receiver. Its drawbacks were the relatively large size (6.1×15.5×3.3cm) and average battery life (approximately 20 hours with two AA-sized batteries). The lifelogger continuously carried this receiver and kept it powered, from November 2007 to January 2009. Data were collected at all times when signals were received. No special travel was arranged for data acquisition. However, the ligelogger is a frequent traveler and this helped creating a spatially distributed data set.

The sampling interval of the GPS receiver is one second. Samples corresponding to the same GPS coordinates are

Table 1: format of the GPS data captured by the receiver.

Date and Time	Altitude	Distance	Duration	Speed	Direction	Latitude	Longitude
2007/11/22 9:22:02	29 m	24 m	0:00:13	7 km/h	199°	N35° 44.713'	E139° 44.755'
2007/11/22 9:22:15	29 m	39 m	0:00:21	7 km/h	200°	N35° 44.700'	E139° 44.749'
2007/11/22 9:22:36	29 m	31 m	0:00:17	7 km/h	198°	N35° 44.681'	E139° 44.741'

combined to form a location record in the format shown in Table 1. The speed is calculated assuming constant velocity between the current and previous samples. The direction recorded is the angle of the vector from the previous location to the current, measured clockwise from North.

2.2 Photos

No special effort was made for capturing the photos. Being an amateur photographer, the lifelogger takes about 8000 digital photos per year. Photos were taken during sightseeing, business travel and other daily-file events. The photos were taken using three digital cameras: a Canon® EOS Kiss 400X (digital SLR camera), an Olympus® C700 Ultra-Zoom (mid-range digital camera), and a Sony® CyberShot DSC-G1 (slim digital camera). Therefore the photos cover a wide range of photographic quality and capture modes. The clocks of the cameras were kept set to the home country's time zone, even when traveling abroad. A simple back-calculation can synchronize GPS data with the photos when they are taken outside this time zone.

3. USING

3.1 Retrieval Systems

The data collection captured as above was quite large. Therefore, It was necessary to design interaction techniques and retrieval systems to facilitate fast and accurate retrieval of GPS data and photos. We summarize two such systems that were built over this data set. These two applications made it possible to search the multimedia collection much faster, according to user studies.

We developed a system for retrieving locomotion patterns from the collection of location data [4]. A clustering algorithm segmented the GPS data according to the nature of the lifelogger's movement. Freehand sketches made on a map and a calendar displayed on a computer screen were used for specifying spatial and temporal queries. We implemented algorithms to analyze a sketch made by a user, identify the query, and retrieve results from the collection of data. Figures 1a and 1b show temporal queries for finding movements recorded during a particular date. The results are shown in a directed graph-like format as seen on the map in Figure 1b. Figure 1c shows a spatial query for finding instances where the person left the city of Tokyo, and the summary of results.

We used both the location data and photos to design a system that retrieves multimedia travel stories from the data collection [2]. The interaction techniques proposed for locomotion pattern retrieval were adopted for image and video retrieval. Figure 2a shows the results of a query for "travel around Tokyo." The presence of the camera-like icon on a circle or an arrow indicates that photos have been taken during the corresponding segment. The user can select each segment to view more detailed results. Figure 2b shows the detailed visualization of location data and photos. The location data plotted on the map change from blue to red with time, indicating the direction of movement. If street view panoramas are available, the system also creates a virtual drive video for the trip (Figure 2c). Figure 2d shows the detailed visualization for a collection of photos taken at a single location. The user clicks on a thumbnail to get an enlarged view of the corresponding image.

Table 2: a summary of the GPS data collection.

Attribute	Values (approximate)
Locations	most parts of the Honshu region (main island) of Japan, and a few locations in South Korea and Germany
Area	a land area of 200,000 km^2
Altitudes	0 to 10,000 m above mean sea level
Speeds	up to 1219 km/h
Total distance	30,000 km
Data points	900,000
Signal reception	8-24 hours per day

3.2 Non-research Use

While the location data were collected primarily for research, the lifelogger found the data very useful in daily life. Combined with the retrieval systems, the data make a "complete and searchable travel log." The most common use has been finding photos taken while traveling. The lifelogger, like many other digital camera users, stores his photos in a collection of folders that are named by date. Not remembering the exact travel dates make a search tedious, especially for somebody who takes photos regularly. While the lifelogger has stopped continuous capture of location data, he still continues to carry the GPS receiver during travel due to this advantage. Another use of the data collection is to recall places and find directions to places that the lifelogger has been to, while collecting the data.

4. SHARING

We decided to share a selected set of the data captured as above, for research and non-commercial purposes. The following subsections decribe how we prepared the data and the issues we had to deal with.

4.1 Preparing Location Data

We selected the GPS data captured during the calendar year 2008, for sharing. Table 2 summarizes this data collection. The location records were grouped in to time slots of two-three days, and then saved as ASCII text files in tab separated value format. This format was selected to allow easy importing to spreadsheets or SQL databases.

4.2 Preparing Photos

The lifelogger captured approximately 8500 photos, amounting to nearly 24 GB of disk space, in 2008. This photo collection could not be shared "as is" for two reasons. The first is the problem of privacy. We should be able to share all the photos if all the people who appear in them consent to sharing their photos with others, and publishing the pho-

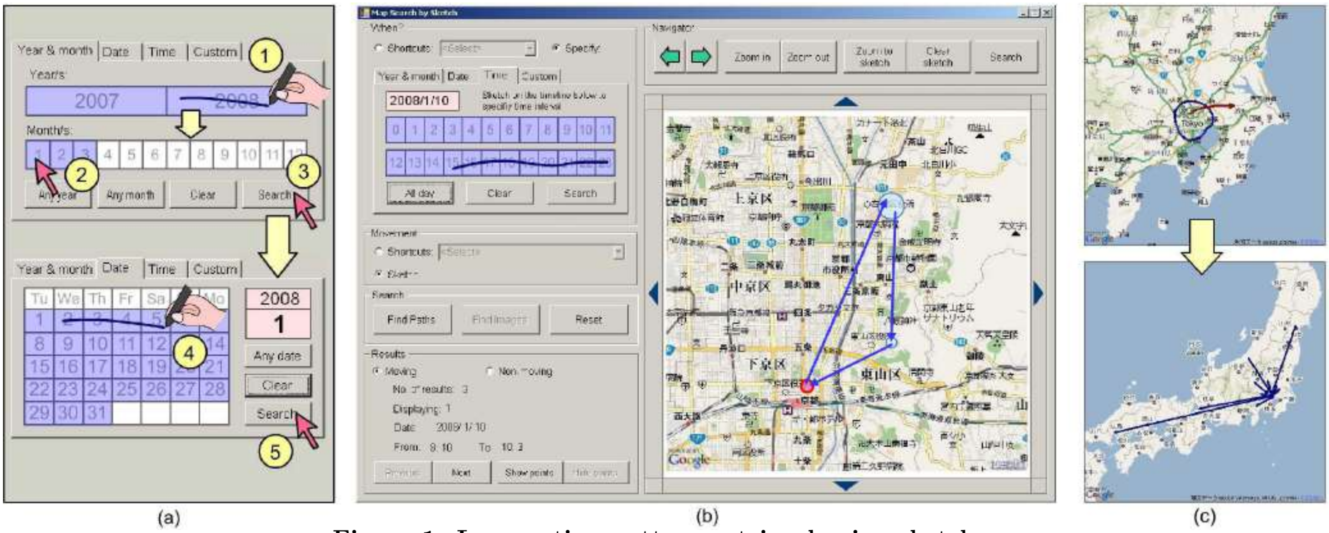

Figure 1: Locomotion pattern retrieval using sketches.

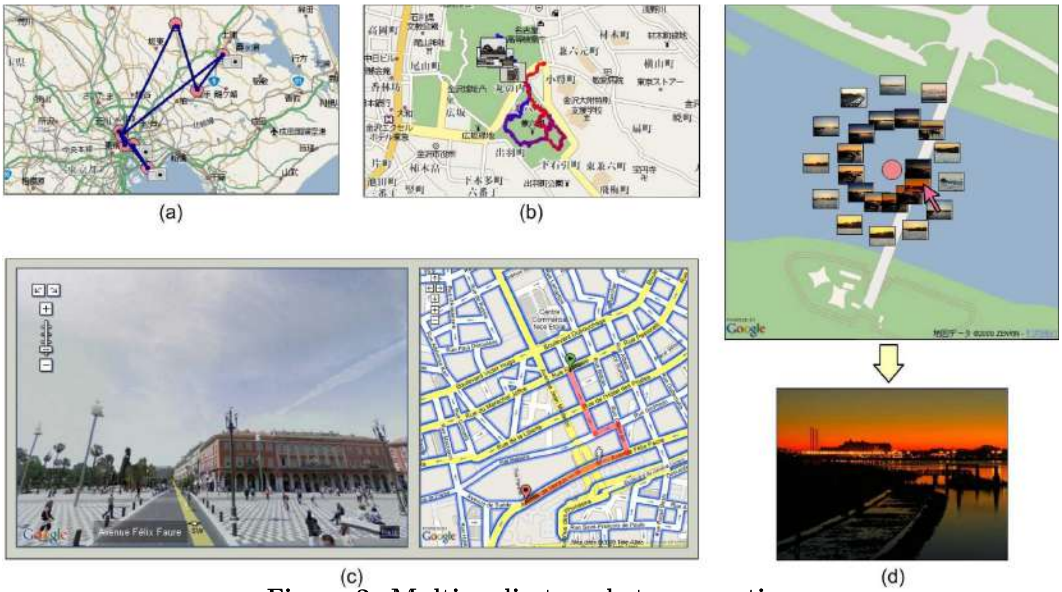

Figure 2: Multimedia travel story creation.

tos. It is impossible to get such a large number of people's consent for sharing photos. It might not be effective to ask not to publish certain photos, after making a dataset open. Therefore, we had to remove most "people-photos" from the collection. We left some photos such as those of public performances and those with small face sizes, with a guideline on the maximum resolution of a photo when used in a publication. These measures ensured that the publishing the photo does not require a model release in most countries [7].

The second reason is the large size of the dataset. Even after removing some photos as described above, there were 4179 photos amounting to approximately 12 GB. The original photos had resolutions ranging from 2 to 10 Megapixels. The quality settings varied from RAW (uncompressed) to JPEG with Q=70. We first converted the RAW images to JPEG with Q=100. Thereafter, we re-sampled all images so that the shorter dimension of each photo is 1200 pixels. Lanczos Filter was used to achieve good quality in re-sampling images. This reduced the size of the photo collection to approximately 5.8GB. All photos contained complete EXIF Metadata identical to those from the original images. These metadata include the date and time of capture, model of the camera, ISO speed, lens aperture, shutter speed, focal length and and other parameters that may be useful in research. Figure 3 shows some examples from this photo collection. The photos were stored in folders named by the date of capture.

4.3 Preparing the Event Index

We created an index to selected types of activities and events that we think are interesting and less common. The index consisted a set of entries such as "*From 2008/02/25 09:53 to 2008/12/25 13:33 - Skiing.*" This was provided mostly as a guide to explain the nature of some of the recorded data. We did not prepare a complete index, as it would be subjective and time consuming.

5. REFLECTIONS

Continuously acquiring sensor data and compiling a lifelog dataset to be shared by everyone is a unique experience. This section covers our impressions and observations that did not fit the previous sections.

5.1 Data Acquisition

The main problem in collecting data was carrying the GPS receiver all the time. The receiver was a bit large, therefore cumbersome to carry. There were several occasions when airport security was not happy with the receiver's presence in the hand luggage. However, we do have data from a few flights in which hand carrying of the receiver was permitted. There was a two week period where the lifelogger was on vacation in a country that had a civil war at the time. Possession of high accuracy GPS receivers was banned, therefore data could not be captured. However, it should also be noted that carrying a GPS receiver was not seen as an "odd" by the friends and family of the lifelogger.

Another problem was to keep the GPS receiver continuously powered. It was necessary to have batteries in hand and change them once or twice a day. Using rechargeable batteries and connecting the receiver to USB ports helped reduce the cost and the burden. However, it still needed an additional effort to remember to replace batteries.

The GPS data were generally accurate, but had location dependent noise. The error was less then 5m outdoors, in areas with fewer buildings. For indoor locations and areas with a large number of tall buildings, the accuracy deteriorated. The barometric altimeter of the GPS receiver provided altitudes with an accuracy of 1-2 meters. However, it gave wrong elevation data during flights, due to the artificially controlled air pressure inside aircrafts.

5.2 Other Issues

Privacy is one of the most important issues in lifelog research. Sharing a year from one's life, even in terms of merely location data and images, releases a large amount of personal information. While the location data are continuous and generally accurate, the lifelogger felt alright to share the data without any filtering. It was apparent that location data were not sufficient to interpret most experiences. Images, on the other hand, had to be filtered carefully. The main problem there was that the lifelogger's location data combined with the photos could expose the privacy of other people in the photos. Most of the group photos in the original photo collection had to be removed due to this reason. However, removal of photos taken with other people has had a bad effect on the dataset. After removing such photos, the remaining photos from a given event do not convey to the viewer the same experience as before. For our own research, and that of who directly collaborate with us, the original image collection is available.

The applications we developed for searching the data made the data very useful. However, the effort in carrying and maintaining the GPS and transferring the data to a computer has been a discouraging factor. Acquiring location data using a mobile phone can solve these problems. Further, most smart phones can accurately estimate indoor locations using Wi-Fi signal strength, and create a more continuous log. In 2008, NTT Communications of Japan started "Kiseki", a service for logging a mobile phone user's movement [1]. However, this was discontinued as continuous location acquisition was too demanding on the batteries of most mobile phones. We hope this problem will go away with technological advances.

6. CONCLUSION AND FUTURE WORK

We presented a multimedia dataset that spans a year of a person's life. The dataset consists of continuously archived location data and 4179 digital photos with EXIF context data. We described our effort and experience in capturing and using the data, and sharing the dataset for public use. The applications we developed showed that location data can be very useful for an ordinary user, if there is a way to capture them without an additional effort.

Figure 3: Sample photos from the dataset.

The dataset has so far been downloaded by more than ten research institutions for research on different areas such as multimedia management, community photo collection analysis, and creation of guided virtual tours. Work is in progress to expand the dataset by adding location data and images for different types of activities and events.

7. ACKNOWLEDGMENTS

This research is funded by a faculty-specific research grant from Microsoft Research Asia.

8. REFERENCES

[1] N. T. T. Communications. Service uses gps handsets to record movments for blogs. http://www.ntt.com/release_e/news08/0005/0529.html, 2008.

[2] G. C. de Silva and K. Aizawa. Retrieving multimedia travel stories using location data and spatial queries. In *Proceedings of the 17th ACM international conference on Multimedia*, MM '09, pages 785–788, New York, NY, USA, 2009. ACM.

[3] G. C. de Silva, T. Yamasaki, and K. Aizawa. An interactive multimedia diary for the home. *Computer*, 40:52–59, May 2007.

[4] G. C. De Silva, T. Yamasaki, and K. Aizawa. Sketch-based spatial queries for retrieving human locomotion patterns from continuously archived gps data. *Trans. Multi.*, 11:1240–1253, November 2009.

[5] J. Gemmell, G. Bell, and R. Lueder. Mylifebits: A personal database for everything, 2006.

[6] C. Gurrin, A. F. Smeaton, D. Byrne, and G. J. F. Jones. An examination of a large visual lifelog. In *In Proc. AIRS 2008*, 2008.

[7] D. Heller. *A Digital Photographer's Guide to Model Releases: Making the Best Business Decisions with Your Photos of People, Places and Things*. Wiley Publishing, 2008.

[8] S. Hodges, L. Williams, E. Berry, S. Izadi, J. Srinivasan, A. Butler, G. Smyth, N. Kapur, and K. Wood. Sensecam: A retrospective memory aid. In *Proc. 8th International Conference on Ubicomp*, pages 177–193, 2006.

[9] Y. Zheng, Q. Li, Y. Chen, X. Xie, and W.-Y. Ma. Understanding mobility based on gps data. In *Proceedings of the 10th international conference on Ubiquitous computing*, UbiComp '08, 2008.

Segmenting and Summarizing General Events in a Long-Term Lifelog

Yi Chen
CDVP, School of Computing
Dublin City University
Dublin, Ireland
ychen@computing.dcu.ie

Gareth J. F. Jones
CDVP & CNGL, School of Computing
Dublin City University
Dublin, Ireland
gjones@computing.dcu.ie

Debasis Ganguly
CNGL, School of Computing
Dublin City University
Dublin, Ireland
dganguly@computing.dcu.ie

ABSTRACT

Lifelogging aims to capture a person's life experiences using digital devices. When captured over an extended period of time a lifelog can potentially contain millions of files from various sources in a range of formats. For lifelogs containing such massive numbers of items, we believe it is important to group them into meaningful sets and summarize them, so that users can search and browse their lifelog data efficiently. Existing studies have explored the segmentation of continuously captured images over short periods of at most a few days into small groups of "events" (episodes). Yet, for long-term lifelogs, higher levels of abstraction are desirable due to the very large number of "events" which will occur over an extended period. We aim to segment a long-term lifelog at the level of general events which typically extend beyond a daily boundary, and to select summary information to represent these events. We describe our current work on higher level segmentation and summary information extraction for long term life logs and report a preliminary pilot study on a real long-term lifelog collection.

1. INTRODUCTION

Lifelogging uses digital devices to capture a person's life experiences. Current digital technologies are making it possible to record things one has seen or heard and to even detect what one was doing by analysis of one's digital activity records and sensor data. Examples of lifelog include [1,2]. Capture of an individual's lifelog can potentially last for many years. Over this time the lifelog might contain several years worth of video material, millions of personal images, many thousands of other files including emails, text messages, and various context data. It is unrealistic to expect people to easily browse such a vast collection of items. For this reason, applications such as that described in [3], are being developed to group certain types of data into small meaningful units, which are generally referred to as "events".

This type of segmentation can enable people to quickly scan for relevant sections of a lifelog covering a day, several days, a week or a longer period. However, according to [3], there are about 20 event per day, that is about 140 events a week, and more than 7000 events a year. Thus there would still be a very large amount of information to browse in a lifelog lasting several years. Thus, higher levels of segmentation and abstraction are desirable. Although grouping of events by dates, months and years can reduce the amount of items (events) that need to be displayed in a time period, it is unlikely that people always know the exact date associated with their required information. Further people do not necessarily want to browse their data using boxes defined by days, for example they may prefer to browse for a higher level "event", e.g. a holiday. We suggest that higher level segmentation should follow the way people remember their past experiences, so as to help them recognize which group (directory) they need to browse.

ECIR'10, Workshop on IAMPA, April 18, 2011, Dublin, Ireland.

In order to browse a lifelog collection, once items have been grouped into events, some form of surrogate summary is needed to represent the event. For example, a keyframe image may be selected from the event to help people recognize it. To enable people to recognize content associated with a certain activity based the information presented, it is important that the selected information is remindful enough to the user. Since the likelihood of recognizing the features of segments depends on how much they resemble the structures of the information in one's autobiographical memory, it is desirable that the segmentation algorithm can follow some general mechanisms of human autobiographical memory.

Autobiographical memory is the memory system responsible for the memory about one's individual's life. According to autobiographical memory theories [4], there are generally three levels of autobiographical memory: *lifetime periods*, *general events*, and *event-specific knowledge*. A lifetime event describes an extended period such as "when I was working at M company" or "when I was living in Y". A general event refers to a more specific period, which is usually in the form of a summary of repeated events of the same theme, such as "working on a small project", or an extended event like "a holiday in Italy". Event-specific knowledge usually contains vivid sensory-perceptual information of a specific event which happened in a consecutive time period (usually less than a day). Most of current event segmentation research has focused on this final level.

In the remainder of this paper, we report on a preliminary study to investigate the segmentation of lifelogs on the second level: general events, and an algorithm for the extraction of summary information for each general event.

2. PROTOTYPE DATA COLLECTION

It is essential to have a long-term lifelog data set to explore the effectiveness of general event level segmentation. As part of our ongoing lifelogging research, several participants have collected 20 months lifelog data collections containing items recording what they have seen and their activities at the computers, together with context information such as their location. The details of the data captured are suuarmarised as follows:

SenseCam photos: A wearable camera called Microsoft SenseCam was used to continuously capture what the lifelogger saw from the first personal perspective. The SenseCam proactively captures up to 6 images per minute.

Computer activities: Each window which comes to the foreground on a computer desktop was recorded by a software application called S'life. Information captured includes the title (name) of the window, the application which the window is opened by, the time when the window came to foreground and the time when it was closed. It also records the textual content and path of files in the window whenever applicable.

Context: The following context information was recorded:

Location: Location information was captured based on GPS and WiFi using Nokia 95 mobile phone. The location information is processed and stored in the form of five separated fields: country, country code, region (province, states, county), city, and street.

People: Names of Bluetooth devices near the lifelogger were captured by the Nokia N95 phone. It is expected that people name their Bluetooth device (e.g. mobile phones with Bluetooth) with their own names, so that the names of people near the lifelogger can be captured. In practice, not all people give their device a personal name or have their Bluetooth enabled at all times.

3. SEGMENTATION ALGORITHMS

Since general events usually take place over a period which contains repeated activities or share the same themes, we believe that it is possible to segment lifelog data to meaningful general events by detecting activity or theme distributions. We assume that in a general event, the information, attributes or features for the theme of the event occur densely, while such features may occur much less frequently in other parts of the lifelog (at least in its adjacent periods). Therefore, the segmentation of a lifelog into general events is similar to segmenting a textual document into several parts with differing main topics.

One method for segmenting textual documents into topical regions is the *TextTiling* algorithm originally developed by Hearst [5]. designed to segment expository texts, which are viewed as being composed of a sequence of main topics and a series of short, into subtopics regiosn. The TextTiling algorithm decomposes the text into blocks of a predefined size w; a pairwise comparison is made between the adjacent blocks using the standard vector space similarity measure, which greater similarity indicating a stronger match. Points with minimal similarity (i.e. valleys in the plotted graph of the similarities between the adjacent blocks) are chosen as the most likely candidates for segmentation boundaries. The valleys in the plotted graph are smoothed with a low pass filter (typically an averaging filter) before applying a threshold. In this preliminary long-term lifelog segmentation study, we explore segmentation based on computer activity information, location data and SenseCam images.

3.1 Segmentation of Computer Activities

We assume that people who spend a considerable time on their computers each day tend to have a periodic focus on computer activity themes. For example, in a certain 3 day period one may be interested in one topic, and read many items about it. Or one may spend some hours every day for several days working on a specific report. Of course, there are “gap” periods when the person is not focused on this topic or activity. We propose an approach of segmenting general events in (at least one aspect of) a person’s lifelog by looking into the distribution of computer activities. The detailed procedure is as follows:

1. **Representing computer activities:** We use titles of active windows to represent computer activities. We assume that the window title is most repetitive for that activity.

2. **Creating a document:** A straightforward approach to creating a composite document is to merge all the window titles into a single string (document) in the order in which they occurred. However, such a document cannot fully represent the distribution of activities since it ignores the duration of activities. For example, one may spend 5 hours each day working on a document, but there may be only one record of this activity. At the same time, some other activities (e.g. loading their homepage on a web browser), which the person only spends 2 seconds on may be repeated many times each day. For this reason, we need to give long duration activities higher weight than short duration ones to reflect their importance. To this we repeat an item N times, where N=normalized duration of activity/normalized mean (if the duration of an activity is longer than the average).

3. **Segmenting documents:** We use Hearst's TextTiling algorithm [5] with a window size = 25 and smoothing parameter s = 1000.

3.2 Location based Segmenting algorithm

We assume that people may consider travelling or a holiday to be a general event in their lives. If they are the type of person who tends to have holidays in a place different from their regular location, these types of general events can be distinguished from others based on the location information.

3.2.1 Simplified approach

A straightforward approach to identifying such events is to read the location information one by one until it changes. As we described in section 2, location information in our database includes names of country, region, city and street names. Since street level location can change very frequently in short time periods (a few minutes or seconds), we excluded street level locations from this algorithm. City or region level change patterns can be variable. For example, some people may travel to an adjacent city or region to work every day. In this case, we consider that frequently occurring location names as “regular location”, and only start segmenting when an unusual location name at the city level or above occurs. Details were as follows:

1) The most frequent city, region and country (which took more than 30% of the time) are extracted.

2) For each occurrence of a region or country which is different from these names, we start a new temporary segment. This segments ends when the region or country changes again.

3) Since when we are travelling it is possible that we pass by part of a region, city or country, if the duration of the segmentation is less than 2 hours, we did not consider that as a general event. Such segments are either joined with the previous segment or the latter. If the previous and latter segments share the same location information, the temporary segment is ignored.

3.2.2 TextTiling segmentation approach

Another approach to segmenting the lifelog based on location is to use the TextTilling method as we described above. We merge the location texts in the order of their timestamp with a format of: [country code][region][city][street]. Since the sample rate of capturing the location data is fixed, no additional repetition is needed. With this format of document, we anticipated that the segmentation could automatically detect stable locations for a given period. For example, a general event of having a holiday in France may involve frequent changes of street name, or even the city names, but the country level information would be stable (and is the most frequently occurring feature in this period), but different from the country names for the rest of time (the surrounding periods). Thus these stable location features could be used to distinguish this event from others.

3.3 SenseCam Concepts based Segmentation

We hypothesize that visual features may change in different general events. So we applied a similar document segmentation approach to segmenting a lifelog with the content of images, or more precisely, concepts identified content in the images, The application we used to detect the concepts in images is described in [6]. It returns a list of confidence scores for the presence of each of a set of 27 concepts typically found in SenseCam images in each image. We adopt the following procedure:

1) Sum the confidence scores for each concept for images in an event segmented by [3].
2) Merge the concepts (repeat N times) from all events in the lifelog to form a document, N = integer (the total confidence score of a concept in an event).
3) Segment with above TextTiling algorithm [5].

4. EVALUATION OF SEGMENTATIONS

In this pilot study, one of the three lifeloggers participated together with her lifelog collection. This includes about 450,000 images, 80,000 Slife records of about 2,000 hours of computer activities, and 18 months of context data (350,000 records). The latest data of the collection is about 14 months prior to the date of experiment. The data was segmented into four parallel sets:

S0) Segmented on weekly basis (baseline)

S1) Computer activity based segments

S2) Location based segments

S3) Simplified location based segments

S4) SenseCam concepts based segments (27 concepts)

The week based segmentation was used as a baseline partially because this time is the most straightforward segmentation base, but also because week, unlike month, is a perceptible temporal circle if a person distinguishes weekday and weekend. Yet, for some people, month may be a better segmentation base if they have more monthly events,

4.1 Method

A five point rating scale was used to evaluate each set of segmentation results (1=definitely a wrong segment point, to 5=definitely a correct segment point). Different materials were provided to assist the participant in rating the segmentations generated by the above algorithms. The average scores were compared with that of weekly baseline segmentation.

Judging Location-Based Segments: To rate the segmentation made based on location information, a list of merged location information was displayed with timestamps of the first and last captured records at that location before the location changes.

Judging Segmentation of Computer Activity Records: To assist the participant in judging the segmentation, we developed an experimental platform which shows:

1) daily computer activity records with: title and time of the top 5 activities in that day with the longest duration.
2) a list of segmentation points with timestamps.

Judging of Segmentation based on SenseCam Image Concepts: The DCU SenseCam browser[1] was used to assist the participant in recalling what was happening around the time of each segment point, so as to make judgments and ratings.

Judging Baseline Segments (Weekly): An experimental platform was developed based on the one used in judging computer activity based segments. Names of locations were added to the daily activity box.

4.2 Results

Five segmentation sets were evaluated using the data owner's manual judgment. Segments were considered to be bad if they were made between:

1) two identical locations: at city level or above

2) two identical computer activities: since we repeated the titles of long duration computer activities when creating the document for segmentation, it is possible that bad segmentation points can be placed between two of the same items which we repeated at document creation stage.

Table 1. Comparing segmentations results

	S0	S1	S2	S3	S4
Total segments	78	81	17	21	46
Bad segments	N/A	7	14	N/A	N/A
Average rating score	2.24	2.22	1.00	4.78	2.71

The simplified location-based approach (S3) was the most precise, while the TextTilling (S2) location-based segmentation did not get a single correct segment. In the latter case most of the segmentation points appeared between two identical records (with the same location at city level). The computer activity based and SenseCam concepts based segmentation approaches tend to have similar satisfactory scores as the weekly based segmentation, which is generally not very satisfactory. However, the score is entirely based on the participants subjective rating, which may not be very reliable, since the user may not always be able to tell what was happening based on the information provided when making the judgments.

5. SUMMARY GENERATION

After the events are segmented a remindful surrogate for the event needs to be generated or extracted to represent the segment, so as to help users decide which event to explore.

It has been found that salient items/events or things a person spent more effort on or repeated many times are better remembered. We hypothesize that activities or features, which occur frequently in one period, but less frequently in the remainder of the lifelog, are a good representative for that period. We developed summarizing algorithms based this hypothesis, and tested them through a self-rating scale regarding: how easily the participant could recognize the periods with the summary information.

5.1 Algorithms

The summary of each general event included the title of the main computer activity, a SenseCam image, and location information.

1. Summary Information from Computer activities

[1]http://sensecambrowser.codeplex.com/.

The representativeness score of a computer activity for a given period is calculated by: The total time of that activity during that period / total time of that activity during rest of the time in the lifelog collection. The top five highest score computer activities' titles of each general event were selected in the evaluation.

2. Key SenseCam image:

1) The key concept is calculated by the: (sum of the likelihood of the concept for images in the given period)/(sum of likelihood of the concept for images in the rest of the lifelog collection);

2) Key images were selected from the images in the period with highest likelihood for this concept.

3. Summary location information

The top two regions, countries, cities (if there were more than one of the above), and street names during the given period were selected to represent the location in that period.

5.2 Evaluation

5.2.1 Method

The summarizing algorithms were evaluated using the lifelog segmented into regions of general events. Since we did not get satisfactory results with the approaches in section 4, we decided to have the data manually segmented by the participant based on computer activities. This segmentation was combined with segmentations from location-based approach (S2). Summary information including key computer activity title, location and key SenseCam image was selected for each segmented general event.

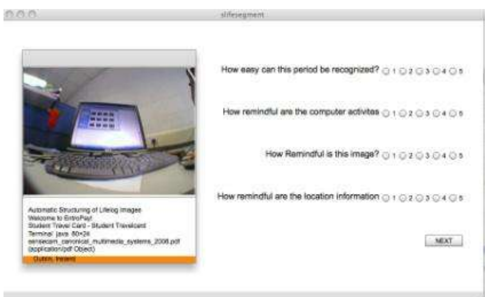

Figure1. Summary evaluation interface

An experimental platform was developed to present the summary information for the user to evaluate (Figure 1). A five point rating scale was used to investigate how easily a general event could be recognized based on the summary information, ranging from 1=unable to recognize the given period at all, to 5=it was extremely easy to recognize the general event with the summary information. Three five point rating scales were used for evaluating the effectiveness of each of the three types of summary information as memory cues for recognizing the periods, with 1=extremely ineffective to 5=extremely strong cue.

5.2.2 Results

The general easiness of recall effectiveness as memory cues for computer activities, location, and key SenseCam images achieved average rating scores of 3.6, 3.5, 1.9, and 2.8 respectively. The location information was particularly representative for the events when the person is away from routine location. Most of the selected computer activities were remindful for the activities that the user was doing, but they were not good cues for the exact time period, e.g. the year, month, etc. While some SenseCam images could be considered as representative for certain periods (mostly routine events), they are not good cues for recalling context information such as what time period an event was in. For general events which took place in a distinctive environment (e.g. a holiday abroad), although some key images were good cues, they did not usually concern the most representative scenes.

6. CONCLUSION

In this study, we examined automatic segmentation of general events in a prototype lifelog collection from an individual life logger. We segmented the lifelog using three types of data: records of computer activities, location, and visual concepts detected through content analysis of SenseCam images, and compared these with week based segmentation against a manual segmentation provided by the life logger. The main approach used for segmentation was to merge the text of records to form a single "document" representative, and then to segment it into sub-documents using the TextTilling algorithm. The results of the pilot study do not provide an immediate solution to the challenges of segmentation of long-term lifelogs. This may due to improper parameters of the TextTilling algorithm, the choice of the algorithm itself or the way in which documents were generated.

We also proposed an approach to selecting summary information to represent the events, and evaluated these using rating scales. We found that computer activity summaries were generally remindful regarding activities during the represented period, but they were not good cues for time attributes. Location provided good cues when it was distinctive from routine locations.

Further work will explore varying the segmentation parameters and alternative ways of creating the "document" to segment. Since how people segment temporal general events may be influenced by the context such as the current task and initial cues, future work may also explore dynamic segmentation which could cater for different tasks.

7. ACKNOWLEDGMENTS

This work is funded by a grant under the Science Foundation Ireland Research Frontiers Programme 2006, and supported by Microsoft Research under grant 2007-056.

8. REFERENCES

[1] Gemmell, J., Bell, G. and Lueder, R, 2006. MyLifeBits: A Personal Database for Everything. *Commun. ACM*.

[2] Lee, H., Smeaton, A., O'Connor, N., Jones, G., Blighe, M., Byrne, D., Doherty, A. and Gurrin, C. 2008. Constructing a SenseCam visual diary as a media process. *Multimedia Systems*, 14, 341-349.

[3] Doherty, A. R. and Smeaton, A. F. , 2008. Automatically Segmenting LifeLog Data into Events. In *Proceedings of 9th International Workshop on Image Analysis for Multimedia Interactive Services*

[4] Conway, M. A. and Pleydell-Pearce, C. W.,2000. The Construction of Autobiographical Memories in the Self-Memory System. *Psychol Rev*, 107, 2

[5] Hearst, M. A. ,2009. TextTiling: Segmenting Text into Multi-Paragraph Subtopic Passages. *Comput. Linguist.*, 23, 33-64.

[6] Byrne, D., Doherty, A. R., Snoek, C. G., Jones, G. J. F. and Smeaton, A. F. ,2009. Everyday Concept Detection in Visual Lifelogs: Validation, Relationships and Trends. *Multimedia Tools Appl.*, 49, 1, 119-144

Activity-logging for self-coaching of knowledge workers

Saskia Koldijk
Radboud University Nijmegen
The Netherlands
saskiaKoldijk@student.ru.nl

Mark van Staalduinen
TNO
Delft, The Netherlands
mark.vanstaalduinen@tno.nl

Stephan Raaijmakers
TNO
Delft, The Netherlands
stephan.raaijmakers@tno.nl

Iris van Rooij
Donders Institute for Brain, Cognition, and Behaviour
Radboud University Nijmegen
i.vanRooij@donders.ru.nl

Wessel Kraaij
TNO & Institute for Computing and Information Sciences
Radboud University Nijmegen
w.kraaij@cs.ru.nl

ABSTRACT

With an increased societal focus on a sustainable economy and a healthy population, well-being of knowledge workers has become an important topic. This paper investigates techniques to support a knowledge worker to manage his well-being. A possible solution is to monitor the workers' behaviour and use this information for giving feedback as soon as his well-being is threatened. Knowledge workers use a broad range of communication means to achieve their goals, like a computer and mobile phone. Our research aims at using features like mouse clicks, active applications or key presses, because these are rather simple features to obtain instead of more invasive tools like a heart-rate monitor. This paper presents the first results of our research. First, logging of low-level features is developed. Based on these features the behaviour of different users is investigated. At first sight, this behaviour seems to be rather chaotic, but by taking into account different tasks, more structure is observed within the data. This paper shows that different behaviour is observed for different users and different tasks, while the same characteristics are observed when a user is performing the same task. This suggests that also anomalous behaviour might be recognized, which is an important result for developing self-coaching tools.

1. INTRODUCTION

In the modern knowledge economy, the demands for productivity of knowledge workers are steadily increasing. At the same time, information sources and communication means are more fragmented than ever. Real-time communication means, such as e-mail, (micro)blogging and other social media have generated an overflow of information, lacking a structure that is adapted to the user's tasks. Networked information systems and portable devices make it possible to work anywhere, posing challenges to context aware net centric organisation of documents, task lists etc. Finally, since the work force in Western countries is ageing it is increasingly important to develop supportive techniques that help people having a reduced work capacity due to a medical condition to maintain a healthy work-style. The project User Centric Reasoning for Well-working (UCR4W)[1] is investigating the key determinants for well-being at work. One of the guiding hypotheses of the project is that logging the activities of knowledge workers can be the basis for an effective computer based coach. The objective of the project is to develop user-centric sensing and reasoning techniques that help to improve well-being at home and at work. Technology should help people improve their sense of being and feeling in control, with a positive impact on work efficiency and effectiveness, work pleasure, mental and physical health status. An example of empowerment is to have relevant information from personal data collections available 'just-in-time'. We think that understanding the activities and tasks of individuals is a key condition to achieve this.

In this paper we describe a study that is carried out as a precursor to the UCR4W project and present initial results of an experiment. The underlying idea of the study is that knowledge workers could possibly be helped to adapt their work style by providing them neutral feedback about their work style and activities. Section 2 discusses the background and assumptions of the study. Experimental results are presented in Section 3. We conclude with the implications of our results and future research in Section 4.

2. FEEDBACK FOR SELF-COACHING

The information overload and context switches of knowledge workers can be a threat for productivity and well-being at work. Let us consider the following scenario:

Scenario

Imagine a typical working day as a knowledge worker. You have different projects running and today your plan is to work on project A and B. For project A you do some internet search and start typing a document. While you are busy

[1]UCR4W is a 8 MEuro 4 year project co-funded through the Dutch national research programme COMMIT. Partners are: TNO, Novay, University of Twente, Radboud University Nijmegen, Philips research, Ericsson, Roessing R&D, Noldus Information Technology, Netherlands Centre for Social Innovation, Dutch Ministry of the Interior.

a colleague asks your help for project C. You interrupt your work to help her search for some information. When a mail about project B arrives you decide to switch to this project as it is quite urgent to finish the required document. Suddenly you notice it is 5 o' clock and you did not finish your work on project A as planned. You might start wondering: How did I spend my time today?

The relation between work-style and well-being at work

Knowledge workers typically have various tasks and deadlines and they have to produce results. The possibility to easily switch tasks makes working very fragmented. The course of action is not always self planned but also determined by external causes, like phone calls, mails, information requests, other persons or appointments [3]. Knowledge workers typically have to self-manage their work and make a good planning in order to be able to accomplish all their tasks. All in all this way of working easily causes a feeling of stress and it is quite difficult to keep a good overview what it is one has done over the course of a day, weeks or even months. A study has shown that knowledge workers often spend effort in tracking their own tasks (13%). Automating this process would be of great benefit for the working process. A system that could monitor and provide overviews of performed activities could support the worker with her self-management, adapt her work-style [1] and in this way diminish cognitive load and stress. More awareness of ones own working process might also have beneficial effects on the on-task behavior and adherence to scheduled activities [9].

Feedback based on action recognition

As a first step to test the hypothesis that tracking activity and work-style can improve well-being at work a simple feedback tool is under development which can automatically infer and log the tasks a user is performing. The log information could be presented in the form of a daily or weekly overview, showing the amount of time spent on tasks and the number of interruptions or task switches. Such a tool requires the following steps: i) design of an activity ontology, ii) automatic logging of low level computer interaction data iii) developing an inference module that maps low level activity to the activity ontology level iv) developing an effective presentation mode for feedback purposes. In this paper we report work on the first three steps, the main contributions of our study are related to i) and iii). The first step necessary in this research is the creation of a taxonomy of tasks people could be performing. Several taxonomies of tasks have already been proposed in the literature. The taxonomies about internet use by Morrison, Pirolli and Card [7] indicate that a distinction between actions on three different levels might be appropriate: The method the user adopts, the purpose of his actions and the specific content. The next step of this project is task recognition. A model will be made for the inference process from simple logging data to higher level tasks. We intend to compare several types of features, both static and temporal in combination with various classifier schemes. In section 3 we will report initial work, since the experiments are still ongoing. A final step is connecting the tasks recognition module to a graphical user interface. It is important to make the interaction with the system as pleasant as possible. Myers et al. [8] state that the system should be directable, personalizable, teachable and transparent. So in order to work well the system should optimally cooperate with – and adapt to – to the user. The tool should provide a means for the user to give feedback without irritating him and it should keep learning.

Related work

There has already been done much theoretical and applied research in the field of action understanding (e.g. [2, 4, 11]). Research on pattern recognition in sensor data, multimodal fusion and models for human goal directed behaviour are relevant for our work. Some research has specifically focused on recognizing patterns of user activities on PCs [6, 10]. These studies all focus on the detection of a specific kind of information to trigger certain actions. Our research differs as we want to log all kinds of activity in order to make a human understandable overview and categorization of tasks. Our intent is to give the user a better overview and more awareness about his working process and in this way help him improving his performance.

3. PILOT STUDY: WORK-LOGGING

3.1 Task analysis

Taking a user centred approach, a questionnaire was used to investigate the typical way of working of knowledge workers and their demands on software supporting them in their daily practices. From the responses by 47 knowledge workers at TNO we concluded that for a tool to be usable for support, the captured activity data should be aggregated to a higher level in order to provide the user with valuable information. The recognition of the task a user is performing is a useful first step towards providing the user understandable feedback and insights about his working process. On the basis of the questionnaire a set of tasks that knowledge workers perform was identified. The answers to the questions 'What tasks do you perform and how do you use your computer for this?' and 'Describe a typical working day' were manually grouped into sets of similar answers to derive a set of typical task types. The appropriateness of the identified set of task types was confirmed by several knowledge workers. From all task types, the tasks performed at the computer were finally selected for automatic task recognition (see Table 1 for the task labels used).

3.2 Data collection

After identification of the software demands by the users, our next step consisted of investigating whether the computer could possibly fulfill these demands. In an experimental phase the computer activities of three knowledge workers were logged using uLog.[2] An additional tool was created that reminded the user every 10 minutes to annotate his or her activity by selecting one of the labels from the task list (and indicating his level of wellworking). About two weeks of data collection resulted in a labelled raw data set. The labelled raw data set was processed to extract several features, for example how often the user clicked or which application was mainly in focus within a five minute time frame (cf. Table 3 for a full list of extracted features; cf. Table 1 for the amount of data points per label). In total 20, 180, 66 labelled segments were recorded for the three users respectively.

[2]http://www.noldus.com

Table 1: Dataset - amount of data per label

Task label	# data	as percentage	F-value
read mail	11	4%	0.583
write mail	12	5%	0.348
organize/archive data	5	2%	0
plan	14	5%	0
make presentation	3	1%	1
create visualisation	4	2%	0.857
program	63	24%	0.977
write report/paper	82	31%	0.8
search information	17	6%	0.654
read article/text	17	6%	0.746
make overview	31	12%	0.621
analyse data	7	3%	0
TOTAL	266	100%	0.656

3.3 Analysis of the labelled data

First analysis showed that distinguishable patterns of computer activity arose per assigned task label. The most indicative feature seems to be the application that was mainly in focus, which is logical as specific tasks require specific applications, as for example ‘programming’ is done in a programming application. But there is not always a simple one to one mapping between application and task. For both the tasks ‘write report’ and ‘search information’ Word has main focus, but someone ‘searching for information’ additionally uses an Internet browser and AcrobatReader (see Figure 1). Therefore the distribution of all applications in the time frame should be considered.

Besides the used applications the keyboard and mouse activity can be used to further distinguish tasks. Figure 2 shows the distribution of clicks and typed characters for the different task labels. Some features alone already have discriminative power (see Figure 3 for an indication of information gain ratio per feature), for example the amount of typed characters is about 0 for searching information, about 50 for mail writing and about 200 for report writing. Combining more features increases the discriminative power, for example tasks not discriminable by number of typed characters (for example writing mail and making an overview, both about 50 typed characters) could be recognized on basis of the number of clicks (about 40 vs. about 80).

A final useful feature that could indicate the task a user is performing is the amount of switching between different applications. Figure 3 plots the typical distribution for various users to show that there are clear individual differences.

3.4 Experiment: Automatic activity labelling

Some initial results about automatic activity labelling are available (see Table 2). We used Weka (see Hall et al. [5]) to train some classifiers and tested their performance by means of 10 fold cross validation. Labelling each activity simply as the majority class ’write report/ paper’ with Weka’s ZeroR classifier yielded a baseline accuracy of 30.83% (F=0.145). Using Weka’s Naive Bayes classifier with just the feature mainApp to classify tasks resulted in an accuracy of 59.77% (F=0.468), so we can conclude that the application that was mainly in focus is a very strong feature. Adding the other features with mouse and keyboard information and indications about active applications and application switches

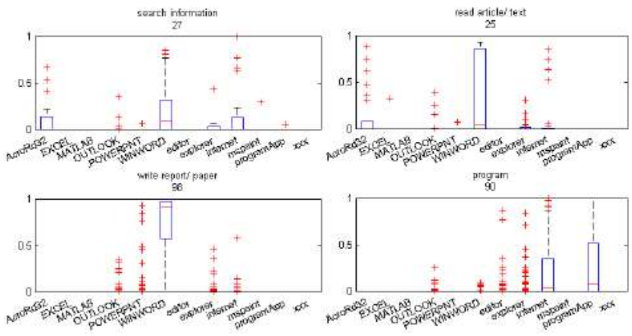

Figure 1: Distribution of application usage (as percentage of time that the application was in focus) per task

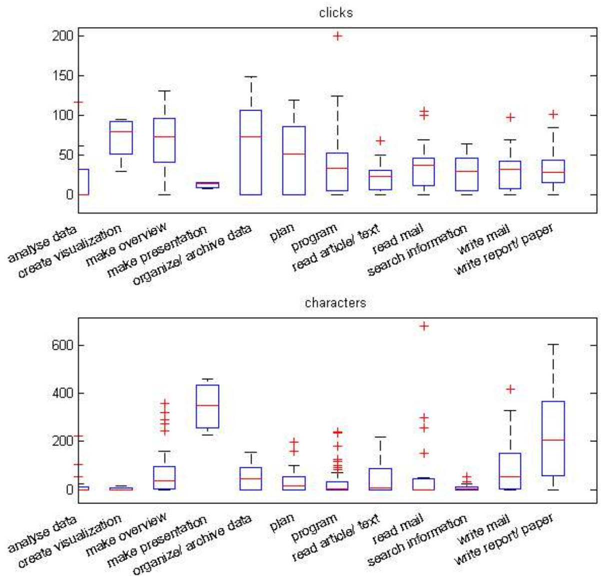

Figure 2: Distribution of #clicks and #characters per task

(discretized with Weka’s preprocessing option) improved the classification accuracy to 70.30% (macro-averaged F=0.656; F values per task can be found in Table 1). Leaving out all features that address the use of specific applications, classification accuracy drops to 52%, with an average F=0.45, which stresses that application-dependent information is important as well for task identification.

4. CONCLUSION AND FUTURE WORK

We have found promising first results showing that it is feasible to log the activities of knowledge workers and use this information to classify the tasks that they are performing. Future research at TNO and Radboud University will focus on extending this research to larger data sets and more systematic comparisons of different task classifiers. The present results suggest individual differences between users, indicating that personalization may be an essential feature

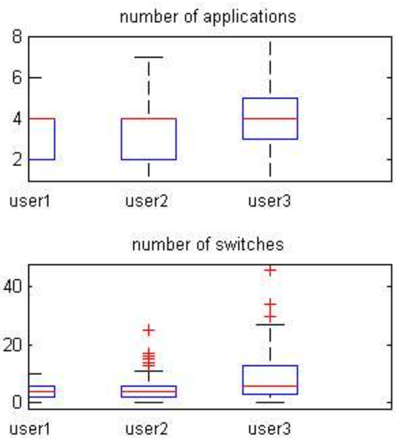

Figure 3: Application usage and switching behavior per user

Table 2: Classification - initial results

Classifier	Accuracy	Averaged F
Baseline (classify as main class)	30.83%	0.145
Naive Bayes (use only mainApp)	59.77%	0.468
Naive Bayes (use all features)	70.30%	0.656

of a task classification tool.

The next steps within the UCR4W project aim at the recognition of anomalous behaviour for each task that might signal a decreasing well-being of a worker. The data collection will probably be extended with a component that captures some semantic content that helps to model the interaction of well-being with an activity related to a particular project. Furthermore, the project will evaluate the self-coaching tools together with end-users in order to improve its acceptance. Finally, proper privacy protection mechanisms and procedures will be an integral part of the project, because these tools are based on personal data.

Table 3: Features extracted within 5 minute timeframes, sorted by information gain ratio (GR)

Feature name	Description	GR
mspaint	% time that mspaint had focus	0.847
user	the user who logged and labelled the data	0.765
programApp	% time that a programming application had focus (eclipse, cmd...)	0.677
OUTLOOK	% time that OUTLOOK had focus	0.626
WINWORD	% time that WINWORD had focus within the timeframe	0.597
mainApp	application that was most of the time in focus	0.522
AcroRd32	% time that AcroRd32 had focus	0.472
spaces	# spaces typed	0.39
characters	# characters typed	0.364
backspaces	# backspaces (inc. 'delete'-key)	0.361
specialKeys	# special keys typed	0.344
clicks	# clicks within the timeframe	0.29
switches	# switches between applications	0.26
internet	% time that an internet application had focus (iexplorer, firefox...) within the timeframe	0.251
daytime	time of the day (as hour, i.e. 9-18)	0
scrolls	# scrolls	0
nrApps	# different applications used within the timeframe	0
time	% time this mainApp was in focus	0
editor	% time that an editing application had focus (notepad++, wordpad...)	0
POWERPNT	% time that POWERPNT had focus	0
explorer	% time that the explorer had focus	0
EXCEL	% time that EXCEL had focus	0
MATLAB	% time that MATLAB had focus	0
label	task label for the activity given by the user	

5. REFERENCES

[1] C. Argyris and D. Schön. *Organizational learning: A theory of action perspective.* Addison Wesley, 1978.

[2] C. L. Baker, R. Saxe, and J. B. Tenenbaum. Action understanding as inverse planning. *Cognition*, 113(3):329 – 349, 2009. Reinforcement learning and higher cognition.

[3] M. Czerwinski, E. Horvitz, and S. Wilhite. A diary study of task switching and interruptions. In *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 175–182, New York, NY, USA, 2004. ACM.

[4] T. Duong, H. Bui, D. Phung, and S. Venkatesh. Activity recognition and abnormality detection with the switching hidden semi-markov model. volume 1, pages 838 – 845 vol. 1, jun. 2005.

[5] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I. H. Witten. The weka data mining software: an update. *SIGKDD Explor. Newsl.*, 11:10–18, November 2009.

[6] E. Horvitz, J. Breese, D. Heckerman, D. Hovel, and D. Rommelse. The lumiere project: Bayesian user modeling for inferring the goals and needs of software users. In *Fourteenth Conference on Uncertainty in Artificial Intelligence*, pages 256–265. Morgan Kaufmann Publishers, July 1998.

[7] J. B. Morrison, P. Pirolli, and S. K. Card. A taxonomic analysis of what world wide web activities significantly impact people's decisions and actions. In *CHI '01: CHI '01 extended abstracts on Human factors in computing systems*, pages 163–164, New York, NY, USA, 2001. ACM.

[8] K. Myers, P. Berry, J. Blythe, K. Conleyn, M. Gervasio, D. McGuinness, D. Morley, A. Pfeffer, M. Pollack, and M. Tambe. An intelligent personal assistant for task and time management. *AI Magazine*, 28(2):47–61, 2007.

[9] G. S. Richman, M. R. Riordan, M. L. Reiss, D. A. Pyles, and J. S. Bailey. The effects of self-monitoring and supervisor feedback on staff performance in a residential setting. *J Appl Behav Anal.*, 21(4):401Ű409, 1988.

[10] J. Shen, L. Li, T. G. Dietterich, and J. L. Herlocker. A hybrid learning system for recognizing user tasks from desktop activities and email messages. In *IUI '06: Proceedings of the 11th international conference on Intelligent user interfaces*, pages 86–92, New York, NY, USA, 2006. ACM.

[11] K. Tahboub. Intelligent human-machine interaction based on dynamic bayesian networks probabilistic intention recognition. *Journal of Intelligent & Robotic Systems*, 45:31–52, 2006. 10.1007/s10846-005-9018-0.

CONTEXT-AWARE SUPPORT FOR ASSISTIVE SYSTEMS AND SERVICES

K.M. Lee, T.F. Lunney, K.J. Curran & J.A. Santos

ISRC, University of Ulster, N.Ireland, email {lee-k1@email.ulster.ac.uk}

Keywords: Ambient Intelligence, Context-aware framework, assistive technology, elderly, sensor technology.

Abstract

An assistive living system that gives our elderly population the comfort and confidence necessary to remain independently in their own homes is essential for enhanced longevity. Ambient Assistive Living technology that provides intuitive and context-sensitive support presents researchers with additional challenges. This paper describes the development of an ambient framework which augments and extends the open source framework OpenAAL with enhanced reasoning, intelligent monitoring of the person and decision-making capabilities. It overviews an "Ambient Assistant" application that showcases the capabilities of the framework. The aim of this research is a generic framework for assistive living with fully interoperable components such as multi-parameter sensors and decision making.

1 Introduction

There is a very strong need to support ambient assistive living technologies in the home, Turner *et al.* [15]. Age UK [1], in a recent study for the Department of Work and Pensions reports that nearly a fifth of people living in the UK today are expected to celebrate their 100th birthday and more than 10 million of the UK's current residents, the equivalent of 17% of the population, are expected to live until they are at least 100. As the population ages, elderly people are more likely to suffer from reduced mobility, disability and mental health problems such as memory impairments. According to statistics published by the Alzheimer's Society UK [3], there are currently over 750,000 in the UK with an age-related cognitive illness and this figure is estimated to reach one million by 2025. The financial cost to the government of illnesses such as Alzheimer's disease and Dementia is 20 billion pounds each year whilst family carers save the economy an additional 6 billion pounds plus per year looking after their loved ones at home. With figures and costs predicted to rise in the future it is important to develop successful Ambient Assistive Living (AAL) solutions which can help people stay independent in their own homes for as long as possible thus reducing avoidable entry into hospitals or care homes. The definition of Assistive Technology (AT) taken from the Alzheimer's Society and The Foundation for Assistive Technology (FAST), [5, 3] refers to 'any device or system that allows an individual to perform a task that they would otherwise be unable to do, or increases the ease and safety with which the task can be performed'. It ranges from the simplest calendar, clock and pill-boxes to high tech solutions such as satellite navigation tracking systems which can locate someone who has wandered. Developing and applying AT within the home that can adapt and intelligently react to the users contextual needs is an extensively researched area. In spite of the significant research most AT on the market today requires the end user to adapt their behaviour to suit the limited intelligence or capabilities of the assisted living device or system, Wichert, [16]. This is inadequate if we are to successfully support the social care of today's elderly population. This paper highlights the main challenges associated with assisted living technologies and describes the approach taken in this research to overcome these challenges. This research aims to address the limitations of existing frameworks by augmenting and extending the open source framework OpenAAL which was developed by Wolf *et al*, [10] with enhanced reasoning, intelligent monitoring of the person and advanced decision-making capabilities. The extended framework namely AMiCA (Ambient Middleware for Context-Awareness) will interface with sensor technology such as the Vicon Revue [15] to model the user's behaviour and environment to develop advanced reasoning and intelligence for an AAL application. It also considers the design of an AAL "Ambient Assistant" application for assisting older adults with activities of daily living. This approach has the attractive characteristic that it offers a supportive environment for both inside and outside the home. It does not rely on a large supporting sensor infrastructure, which means it can be easily deployed in a range of contexts and locations without any changes. This facilitates new classes of ubiquitous AAL applications and devices can be developed with the framework thus enhancing the range of useful assistive technologies for older adults.

2 Assistive living technology

The Royal Commission on Long Term Care, [11] report that enhancing the social care of our elderly population can be achieved through the preferred strategy of aging in place. To date fulfilling the requirements of this approach has been difficult due to the diverse needs of this social group. Enabling an elderly person to remain in their own home has been shown to enhance their quality of life, as reported in the findings of a systematic assessment of the social, ethical and privacy issues involved in ICT and Ageing by the EU Senior Project [12]. Designing dynamic AAL systems for an actively aging population requires highly adaptable and intelligent context-awareness of users and their behaviours. A greater number of challenges appear due to the complexity involved in designing dynamic AAL context-aware systems and services. Henricksen *et al*, [7], describe the challenges as including the integration of a wide range of software and hardware components whilst at the same time maintaining the need for user privacy and personalised interaction with the right levels of context-sensitivity. The complexity involved in designing and developing context-aware systems and services makes context-aware middleware an essential requirement. Context-aware middleware is recognised as the enabler and simplifier of technology and services for dynamic systems such as those necessary in AAL. The need for an efficient middleware framework to seamlessly join the required sensor, network protocols, hardware, and software components together is well accepted [7]. To date various context-aware middleware platforms and frameworks have been designed however none have become an accepted standard in this area [2,10,13]. The intention of this framework is to enable the seamless provision of context-aware services to highly dynamic environments such as AAL environments. AMiCA aims to simplify the design process and address the challenges by integrating sensor technology, intelligent algorithms and ontologies in its multi-layered architectural design approach.

2.1 Challenges

The following computer science challenges arise in the field of assisted living:

Sensing the users' environment: involves gathering data about the users' daily living activities in the home and outdoors. Current sensors often gather imprecise data.

Data: acquisition, integration, modelling, presentation and distribution of data from multiple sources can often be difficult and lead to uncertainty.

Interaction: supporting effective communication between the user, family/carer and social care provider.

Personalised interaction: recognising the users social care needs and how they change over time. Interaction which is unique to the user needs to be context sensitive, adaptable and flexible in order to accommodate their changing needs.

Analysis: involves context data, current environment conditions and surrounding objects which are used to exploit and derive the users' activity. This data is often 'noisy' and difficult to apply intelligence and inference mechanisms.

Confidentiality: appropriate measures to support individual privacy and dignity, confidentiality, information and network security, and appropriate use of sensors and data collection.

Integration: interoperability issues between the heterogeneous hardware and software components in AAL systems and devices. No accepted standard to date.

In view of the diverse challenges in assisted living, this research aims to build on existing work and provide a framework which will support dynamic context-aware services optimised for use in an assistive living environment.

2.2 Addressing the challenges

Developing an efficient framework from scratch is a significant task, mostly due to the length of time involved and level of expertise required to design and develop an efficiently generic system. A practical alternative is the adoption and improvement of an existing middleware platform. Based on the timeframe and scope of this research, the OpenAAL [9] framework with additions such as an enhanced intelligent reasoning module, a dynamic inference module and an activity prediction module will add the necessary functionality required to build context sensitive applications for AAL domains. The AMiCA framework is a multi-layered design with intelligent reasoning and decision-making support. This approach aims to seamlessly and opportunistically provide the connectivity required by services of highly dynamic environments such as AAL clients. By delivering appropriate timely context-aware services new classes of ubiquitous AAL applications and devices can be developed thus increasing the quality of life of an increasing elder population. The key components within the AMiCA framework (see Figure 1) are the environmental sensing layer, the sensor integration layer or sensor fusion layer, the context management layer, the application logic layer and the service layer. The lowest architectural level of the architecture consists of location and context sensors in a person's home environment. To enable the dynamic integration of context sources, it is essential to keep the sensing as a separate layer. Details related to data acquisition from the various sensors can be made available on demand.

At the lowest level of the system direct and indirect context is sensed using the Vicon Revue and if necessary other suitable sensors such as Java Sun SPOT's and Parallax ultrasonic sensors may be employed. The users existing assistive technology systems such as telecare systems etc can also be incorporated into the framework if appropriate. This will address integration issues found in existing frameworks. Systems which are interoperable and can co-exist with other plug and play devices offers another unique aspect to this work as yet there are no accepted standard in this area. The sensor integration layer combines the data that has been acquired by the sensors and other relevant context sources and intelligently groups the data together. This approach simplifies application development by promoting reuse of software. Additionally higher-level context data can be more easily inferred as a result of intelligently managing the data at lower levels. This multi-layered approach to the handling of context data facilitates decision making at the highest levels. This intelligence is further enriched with context related facts about the user or service and passed to the context management layer for intelligent reasoning, inference and activity prediction. A combination of ontological reasoning and rule-based reasoning is used to efficiently group and reason over the context data in the reasoning module. Domain modelling and a combination of Bayesian networking and a Hidden Markov Model (HMM) algorithm approach is used to optimise inference over uncertain data as part of the analysis of users' behaviour. This will enable prediction of profiles to assist with intelligent monitoring. The contextual manager stores details on all the context sources that are available or become available during discovery. The contextual data store records the context data received from the various context sources and also records Spatio-temporal changes in the data. Each sub-level of the context management layer will have access to this data store over time. The application logic layer implements the logic of the actual context-aware 'Ambient Assistant' application and will make calls for additional resources such as web-based services on hosted servers in the cloud. Based on the results of context reasoning and decision making mechanisms in the lower layers, the higher-level applications can adjust their behaviours and adapt their services for users. Service discovery and registration is also delivered at this level in the AMiCA architecture.

The key components within the multi-layered AMiCA architecture are:

– Sensing mechanisms/procedures
– Sensor integration layer
– Context management layer
    – Intelligent reasoning module
    – Inference module
    – Activity prediction module
    – Contextual shared data store
– Application logic layer
– Service registry
    – Cloud data store

A cloud-based assistive living application, namely Ambient Assistant (AA) will be developed to validate the efficiency of the framework. The application will be cloud-based thus promoting true ambient intelligence. In addition, the creation of such a highly adaptive and context-aware cloud-based assistive living application will have the potential to provide a wide range of support.

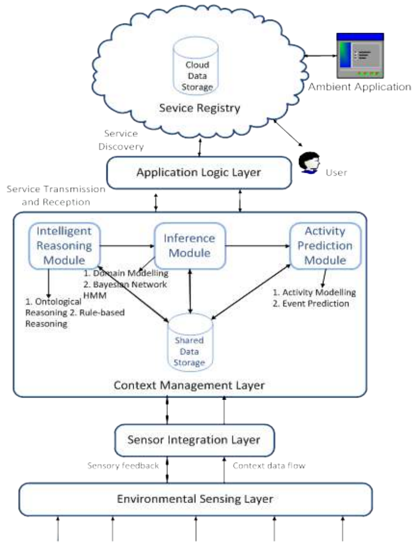

Figure 1: High-level overview of AMiCA framework

2.3 Sensing using the Vicon Revue

The majority of the context data collected in this research will be sensed using the Vicon Revue. The Revue is a small lightweight, wearable device that passively captures the persons' day-to-day activities as a series of images and sensor readings [4,6]. In our research the Revue image data is mapped against actual locations both inside and outside the home and is used to monitor the user's activities on a daily basis. The Revue's image mapping data although created inside and outside the home by somewhat different means uses key reference points and their associated tagged images for intelligent monitoring of the person. This background mapping data can now be used in two main ways. As the GPS-enabled Revue is not available at this time, we have used a Sony Ericsson k850i mobile device which is GPS enabled to capture images for real-time monitioring. This involves real-time monitoring of the persons' behaviour from the incoming images and using these images to determine if the person is coping sufficiently well or needs immediate prompting. If they are outdoors and are not following their expected daily routine (e.g. they are half a mile from home, when their 'meals on wheels is arriving') then they can be proactively prompted via a voice based prompt from a mobile device. Another way of using the image data is to use the data gathered over a period, say a month/year to determine the lifestyle behaviour of the person and evaluate if they are carrying out their activities to an acceptable standard – e.g. less time spent outside the house or shorter distances travelled may indicate deterioration in their mobility. This history profiling requires intelligent processing of the collected image data and sensor readings to enable meaningful results to be obtained. For example the persons' behaviour can be profiled with regard to expected daily/weekly/monthly behaviour and then using rule-based reasoning and intelligent techniques can be used to determine any deterioriation in health related activities recognised. The data acquired to date is integrated into the AMiCA framework and an intelligent user interface application and voiceXML based mobile prompting system is been developed. In addition a communication link to the family carer/helper via a text message will be incorportated into the design and if major deviations from the persons' expected behavior occur during monitoring the carer will receive communiction. This work provides an attractive characteristic in that it offers a supportive environment for both inside and outside the home without relying on any supporting sensor infrastructure, which means it can be easily deployed in a range of contexts.

2.3 'Ambient Assistant' Application

The proposed 'Ambient Application' (Figure 2) will showcase the unique contributions of the enhanced middleware platform. The application will offer the user a fully interactive touch-screen companion on a screen appropriately located inside the home and also on a mobile device. The application will prompt, remind and locate the person if they are lost and generally offer comfort and support.

Figure 2: Ambient Assistant GUI

3 Conclusion

Ambient frameworks for assistive living environments have particular requirements not present in other domains thus presenting many unique problems for the researcher. It is generally accepted that a middleware platform is necessary to provide context-aware support. This paper has presented AMiCA, an ambient middleware framework for supporting the development of context-aware applications. It is based on a multi-layered intelligent reasoning approach where dynamically sensed context can be represented, reasoned, adapted and utilised. By providing this level of intelligent decision-making, innovative and intuitive assisted living solutions are facilitated thus supporting and enhancing the quality of life of our aging population.

References

[1] Age UK, Department for Work and Pensions (DWP). Opportunity Age – Meeting the challenges of ageing in the 21st centuryDepartment for Work and Pensions (DWP). http://www.ageuk.org.uk/latest-news/fifth-will-celebrate-100th-birthday/

[2] Amigo. 2008. http://www.hitech-projects.com/euprojects/amigo/

[3] Alzheimer's Society, Statistics. http://www.alzheimers.org.uk/site/scripts/documents_info.php?documentID=341

[4] A.R. Doherty, and A.F. Smeaton, Automatically augmenting lifelog events using pervasively generated content from millions of people. Sensors, 10 (3). pp. 1423-1446, 2010.

[5] Fast UK, http://www.fastuk.org/about/definitionofat.php

[6] J. Gemmell, L. Williams, K. Wood, R. Lueder, and G. Bell, Passive Capture and Ensuing Issues for a Personal Lifetime Store, *Proceedings of the 1st ACM workshop on continuous archival and retrieval of personal experiences* New York, USA October 2004.

[7] K. Henricksen, J. Indulska, T. McFadden, and S. Balasubramaniam, 2005. Middleware for Distributed Context-Aware Systems, *Proc. of the International Symposium on Distributed Objects and Applications (DOA'05)*, Cyprus, October 2005, *Lecture Notes in Computer Science*, LNCS 3760, Page(s): 846-863

[8] S. Martin, G. Kelly, W.G. Kernohan, B. McCreight, C. Nugent, "Smart home technologies for health and social care support", *Cochrane Database of Systematic Reviews*, , no. 4. (2008).

[9] openAAL. http://openaal.org/

[10] P. Paganelli, E. Spinicci, and D. Giuli. 2008. "ERMHAN: a context-aware service platform to support continuous care networks for home-based assistance," International Journal of Telemedicine and Applications, Volume 2008, Issue 5 (January 2008) Article No. 4

[11] Royal Commission on Long Term Care 1999. With respect to old age: long term care - rights and responsibilities. http://www.dh.gov.uk/en/Publicationsandstatistics/Publications/PublicationsPolicyAndGuidance/DH_4008520

[12] Senior Project. 2008. Ethics, E-Inclusion and aging, "Preserving independence has been shown to have a positive impact on elderly people," http://seniorproject.eu/events/senior_discussion_paper_for_the_vienna_conference

[13] Soprano. 2008. http://www.soprano-ip.org/

[14] K. J. Turner, J. L Arnott, P. D. Gray, S. Renals. Grand Challenges in Assisted Living- Home Care Technologies. Match Consortium. *www.match-project.org.uk*

[15] Vicon Revue. http://viconrevue.com/home.html

[16] R. Wichert. Fraunhofer Ambient Assisted Living Alliance, Darmstadt, Germany. AmI'10 Proceedings of the First international joint conference on Ambient Intelligence, Germany. Springer-Verlag. Berlin, Heidelberg 2010. ISBN:3-642-16916-3 978

Distributed Personal Archives: Deriving Health Profiles from the Web

Jochen Meyer
OFFIS Institute for Information Technology
Oldenburg, Germany
meyer@offis.de

Benjamin Poppinga
OFFIS Institute for Information Technology
Oldenburg, Germany
poppinga@offis.de

ABSTRACT
An active online user nowadays stores and publishes more and more information in internet based systems. The overall amount of information is building up a *Distributed Personal Archive* (*DPA*). Our vision is to analyse the individual user's data contained in the *DPA* to conclude information about the user and to derive information about his health state. By analyzing the data in the *DPA*, health features can be retrieved which, when jointly assessed, allow to conclude on the user's health profile with respect to a given health state or disease. Taking into account specific characteristics of *DPA* derived health profiles, medically sensible applications can be developed.

Categories and Subject Descriptors
H.3.3 [**Information Systems**]: Information Storage and Retrieval—*Information Search and Retrieval*; J.3 [**Computer Applications**]: Life and Medical Sciences—*Health*

General Terms
Design, Human Factors

Keywords
Distributed Personal Archives, Personal Health, Wellness Informatics

1. INTRODUCTION
With the evolving use of cloud, online and Web2.0 applications, people publish and store large amounts of media and information online. People post their current activities on Facebook, twitter their feelings and emotions, publish videos on YouTube, share photos on Picasa and organize their family's schedule using Google Calender. They are building up a distributed and heterogeneous amount of data that is complementing and to some extent replacing the local personal media archives, thereby forming what we call a "Distributed Personal Archive" (*DPA*).

DPA therefore contain considerable amounts of information about the user, his profile, context and interests. Our vision is to derive a personal health profile, i.e. information about a person's health state from his *DPA*. For certain diseases and health states we expect that it is possible to draw medically significant conclusions based on the analysis of *DPA*.

The paper is structured in the following way: First we explain the concept of Distributed Personal Archives. Following that we describe how data in *DPA* can be analysed for health purposes by first deriving health features which then are jointly assessed to conclude on a user's health profile, thereby taking advantage of the specific properties of *DPA* over conventional personal archives. We describe which types of applications we see on top of that data and end with a conclusion.

2. DISTRIBUTED PERSONAL ARCHIVES
Nowadays an active online user is often using a plethora of online services such as social networks, blogs, communities, portals, mail and calendar services etc [6]. Since these online services virtually always require the creation of an account, the online information, data and activities can be attributed to the individual user. This data is often highly social, being interlinked and connected to the user's social network, related to other users' data, interactively commented and "liked", cited and re-tweeted. The overall amount of such online information and data, attributed to an individual user, distributed over the various heterogeneous services and sources, and connected to the user's social network, is what we call that individual user's *"Distributed Personal Archive"* (*DPA*).

The information that is found in a *DPA* is often strongly related to a user's behaviour, his or her recent activities or mental state, but also the location, context and environment. Therefore it is well-suited to draw conclusions related to the user's well-being and health [5]. *DPA* sources that we consider to be potentially interesting are:

- In **microblogs** (e.g. Facebook status message, Twitter tweeds), users share current thoughts and activities, and generally express themselves.
- **Media Sharing** platforms (e.g. Flickr, YouTube) frequently are used for sharing memories about past joint events with friends and relatives.

- **Social Activities** (e.g. commenting on or "liking" other activities, re-tweeting on Twitter) as one core functionality of Web2.0 communities are particularly used for communicating with others but also for expressing oneself by embracing other content.

- **Location Based Networks** (e.g. Foursquare, Gowalla, Google Latitude) add spatial context to social activities and allow to identify a user's and her or his friend's past and current position.

- More and more **sensor based systems** upload information to internet portals. The most popular examples are GPS based systems, including Google Latitude or Glympse to log and share one's current position, and various applications used for hiking, running, biking and other outdoor and fitness activities. Other sensors include the NikePlus system that integrates a shoe-worn step sensor to monitor a runner's pace, and the Wiscale that uploads the user's current weight and body fat percentage via WLAN to the Withings portal.

- **Online Applications** relocate virtually all types of previously local applications into "the cloud". Systems such as Google Calendar, Toodledo (to do list) or Evernote (notebook) particularly support the productivity and the personal organization.

- A health specific data source is the **Personal Health Record** (PHR) [4]. PHR are motivated by the need for patient empowerment and the related information requirements of the individual. They store all kind of health-related information, including information about diseases and medication but also general data such as weight, blood-pressure, fitness level and others. Unlike electronic patient records that are stored in a hospital or at a physician and focus on diseases, PHR are owned and maintained by the individual person and stress the aspect of life-long monitoring of the health status and positive health behaviour. Recent examples are Google Health and Microsoft HealthVault.

The level of privacy and visibility of the data in a *DPA* varies considerably among the different sources [8]. While some (e.g. Twitter Feeds, Blogs) are publicly accessible, others (e.g. Facebook status messages, one's location on Gowalla) may be visible only to the peer network, one's "friends" or "buddies". The access may also be restricted to only an explicitly defined group (e.g. photo sharing on Picasa). Finally, cloud-based applications are usually designed for being accessed only privately by the user itself.

Insights about the individual user can already be gained by just taking into account the publicly visible information (and systems such as PleaseRobMe.com aggregate such data in a pretty spectacular way). However, for a meaningful data analysis access to the user's accounts in the various *DPA* sources is needed. The user needs to grant the system access to his *DPA*. Consequently the resulting health analysis must be strictly private to the user, too. Clearly this raises important questions of privacy and security, particularly in the area of personal health, which any application using the personal health profiles will need to address.

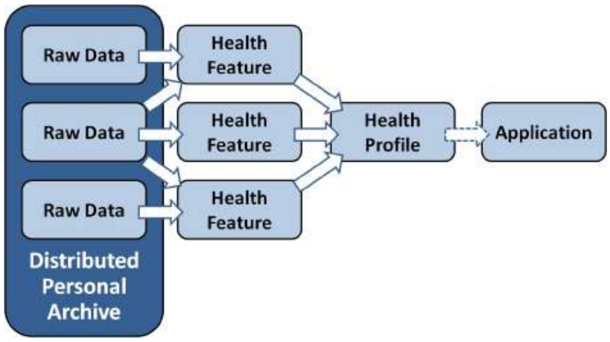

Figure 1: Process to exploit DPA health data

3. USING DPA IN HEALTH

With the aforementioned data sources, *DPA* contain a particular amout of data about a user's current behaviour and context. We expect that reasonable information about a person's health state can be concluded from *DPA*, and that such information can sensibly be used in a health related context. However, the deriveable information has limitations in (medical) quality, expressiveness and soundness, and not all information found in *DPA*s is useful for every disease. Therefore the user's information needs, addressing a specific health state or disease must be matched with possible *DPA* sources containing potentially relevant information that is related to that given area of interest.

We envisage a multi-step process to exploit the health data in *DPA* as indicated in figure 1. In the first step, the data in the *DPA* is being analysed to derive *health features*, single indicators that are relevant for a given health state. In the second step, these features are jointly assessed to identify the user's *health profile*, a medically sound description with respect to a given disease or health state. Depending on the health state and the intended use case, various *applications* can be realized or adopted.

Deriving health information from *DPA* is therefore always related to a specific need for information, according to a given health state or disease. The goal is not to do a general assessment of the user's overall health, thereby coming up with surprising and questionable diagnoses, but to use existing information in *DPA* to answer specific questions the user asks himself about his health.

3.1 Deriving Health Features from DPA

"Health features", as we call them, are single indicators of the personal health state. Examples of such features may be the level of activity, the weight, mood, nutrition, social connectedness etc. In itself, these features may be of limited expressive value or they may not be clearly quantifiable. For example, without extensive sensor integration it is probably not possible to estimate the amount of calories that a user has burnt by physical activities. However, based on information about the user's location throughout a given day, it may be possible to guess if that day has been more sedentary, or if the user has moved around a lot. Moreover, even if the absolute value of the individual feature is not meaningful, dependent on the envisaged application, the change of that value over a period of time may be valuable.

In order to analyse data for deriving a user's health features, we take advantage of the particular properties of a *DPA*. Since *DPA* include heterogeneous sources containing e.g. text, photos, videos, sensor and other types of data, they are inherently multimedia and multimodal data sources. Therefore all types of state-of-the-art analysis methods that are used on local content can be deployed. Beyond that, *DPA* analysis methods can take considerable advantage from the fact that many of the data sources are also "socially networked". Thus, various analysis methods are possible:

- The user's data, his content, messages, sensor data etc can **directly be analysed** using state of the art multimodal, multimedia analysis methods, e.g. image, text and video analysis. The concrete applicable methods depend on the specific type of data stored in the respective *DPA* source.
- **Social interaction data** is produced by the interaction of the user with his social network. It can be "outgoing" data that is originating from the user himself, e.g. by commenting, nudging or liking other activities, or "incoming" data by his social network, e.g. his buddies' comments or dashboard messages. Also, physical proximity in the real world may be identified by taking into account the user's and his buddies' spatial information from location based networks. Such social interaction data is a new type of data that is not available in local media archives but only in the online-context of *DPA*.
- Further context information can be gained by accessing publicly available **third-party sources** based on the user's own data. When the user's location at a given time is known from FourSquare, the weather or air quality can be derived by accessing an environment information service.

Based on these analysis methods, examples of health features that we envisage are:

- By analyzing location data from e.g. Google Latitude, Foursquare or Gowalla, it may be possible to conclude if the user had a more sedentary day, having been in the office troughout the day, or if he had been pretty active, moving around to different spots, shopping and using public transport.
- The number of both incoming and outgoing messages sent over social networks allows conclusions about the social connectedness of the user.
- Photos on social websites which the user is posting himself or on which he is marked, contain information about the "when" and "where" and "who", allowing different analyses e.g. on leisure and vacation, friends and family. Such approaches are already being used for the automatic generation of photobooks [7].
- Thematic and sensor-integrating portals allow for retrieval of quantified data, e.g. of weight (Withings), jogging and fitness (RunKeeper, NikePlus) or nutrition (DailyBurn).

There are multiple relations between *DPA* sources and features. The same feature may be derived from different *DPA* sources: Daily activity may be assessed quite precisely by using the data from a pedometer's internet portal (e.g. FitBit), or by analyzing GPS data (e.g. from Latitude), conclusions about the user's activities may be drawn, too, less precise, but also less obtrusive to the user. On the other hand, the same *DPA* source may be used to derive different features. Location position can be used to differentiate working (being in the office) from family time (being at home).

3.2 Deriving Health Profiles from Features

A "health state" is a medically relevant desription of a certain aspect of health and well-being, e.g. fitness or mental wellness, or a specific disease, e.g. obesity or back pain. An individual's "health profile" then describes to which level that specific person fulfils a given health state. Thus, health states and health profiles do not describe the general health and well-being of a given person, but rather focus on specific scenarios. They aim at fulfilling the user's information needs according to a specific and well-defined health question: Am I doing fine in losing weight? Does my daily lifestyle support prevention of cardio-vascular diseases?

The indicators or symptoms to be observed for such a health state must be defined according to medically sound models. Such models may for instance be derived from the International Classification of Functioning and Disabilities (ICF) [9]. The ICF, defined by the World Health Organization, is a classification that describes the functioning of a person on the level of the body or body part, on the person as a whole and within a social context. The ICF contains more than 1.400 classes, covering virtually every part of the body and every aspect of daily life. For certain diseases as diverse as depression, diabetes or breast cancer, subsets of the ICF, so called core sets, have been defined, consisting of as few as possible ICF classes to describe a person's level of functioning with respect to that health state [1]. Thus, by assessing the person's functioning in the classes of a given core set medically sound conclusions may be drawn on the person's health profile on that health state.

In order to derive a medically sound health profile from *DPA*, we aim to match the health features, as outlined above, to the classes of the given health model. For obesity, "walking" and "moving around" are classes of the ICF core set, hence physical activity is clearly a relevant feature. However, also "handling stress" and "immediate family" are classes in the core set, so both stress and social connectedness could be taken into acount as well. Frequently, the matching between derived features and classes will be unprecise and not perfect. For instance, social connectedness may be only vaguely related to immediate family. Nevertheless, it's an alternative and supplementary view on these aspect, going beyond the individual's self-assessment, third person's assessments and other measures.

Since we can particularly derive information about a person's behaviour and environment, the approach is well-suited for *behavioural diseases* that can be influenced by personal behaviour, and for *environmental diseases* that are influenced by environmental state. Examples include:

- **Healthy lifestyle, physical fitness and well-being**: Important indicators are the level of physical acitivity and the nutrition [3].
- **Mental well-being, stress**: The user's mood, his social connectedness, his level of communication with others or the time spent for work vs. leisure activities are considered to be some of the indicators for mental well-being.
- **Dermatitis**: Factors for acute aggravations include sweating, the exposition to certain allergene substances and stress. Hence, the surrounding's temperature and the concentration of the allergenes, but also the user's level of stress, as outlined above, are relevant health features.

3.3 Potential Applications

A health profile that is derived from *DPA*s has different characteristics compared to conventionally assessed health data. Clearly, the quality, expressive value and medical soundness are much lower. However, in a trade-off between data quality on the one hand and acquistion effort on the other, the approach is interesting: Clinical data that is acquired by professional and trained personal is of highest quality and reliability but very expensive in terms of time and money. Consequently, unless in case of a (chronic) disease or in certain phases of life such as pregnancy such data is often *not* being assessed. Hence, this high quality data has a very low availabiliy. The *DPA* analysis, in contrast, has a very low, often practically zero effort for acquisition. To some extent, it is a by-product of the digital lifestyle. Therefore, it can be made available for long periods of time. Moreover, since the data is acquired withouth further prerequisites, the assessment is very unobtrusive from the individual's point of view.

Personal health profiles are of lower data quality, they must not be confused with a general medical assessment and they cannot draw any new medical conclusions. Rather, their goal is to answer a user's specific question by re-using and analyzing existing data thereby "squeezing out" as much value-added information as possible. Applications that use the derived health profiles need to take into account these properties. Ideas that we identified include:

- The continuous long-term monitoring of the health state helps to identify slow and long-term changes, such as a decrease in physical activities and social interaction over several months or years. Such information can be useful to indicate the need for a health behaviour change. On an even longer perspective, data may be stored in Personal Health Records, enabling a life-long view on one's health history.
- *DPA* analysis enables a different view on one's personal health state. Thus, it supports the reflection about personal health and can be an aid in health behaviour change [2].
- The unobtrusive monitoring of one's health state may be a considerable relief for managing chronic diseases, where the feeling of being stigmatized by conventional monitoring methods is an issue.

4. CONCLUSIONS

The concept of *DPA* and the approach to assess a person's health profile is based on the deduction of meaningful information from online sources. However, not every person is extensively using online social media, and even for those who do, the data in their *DPA* is reflecting only a selective aspect of their life, being far from representative. The main challenge thus lies in the matching between the limitations of the data – what meaningful information can be derived from *DPA*s – and the potentials of of the approach – how certain health states can sensibly be supported by new applications.

The approach is still a vision. The idea will be proved by developing a prototype and concluding appropriate studies. This requires the definition of a use case by identifying a suitable target user group and a specific health state. An evaluation will then need to show that a *DPA* derived health profile for that use case is in fact as meaningful as a conventionally assessed one. We think the idea is promising and that it is possible to derive enough data from *DPA* so that for certain diseases or health states medically significant conclusions can be drawn and sensible applications be developed.

5. REFERENCES

[1] G. S. Alarcos Cieza, Szilvia Geyh. Entwicklung der ICF Core Sets. In *Diagnostische Verfahren in der Rehabilitation*, volume 5 of *Diagnostik für Klinik und Praxis*. Hogrefe, 2008.

[2] C. J. Armitage and M. Conner. Social cognition models and health behaviour: A structured review. *Psychology and Health*, 15:173–189, 2000.

[3] W. W. Hoeger. *Lifetime physical fitness and wellness*. Morton Publishing Company, 2 edition, 1989.

[4] D. C. Kaelber, S. Shah, A. Vincent, E. Pan, J. M. Hook, D. Johnston, and D. W. B. B. Middleton. *The Value of Personal Health Records*. Healthcare Information and Management System Society (HIMSS), 2008.

[5] N. Kamal and S. Fels. Online social networks for personal informatics to promote positive health behavior. In *Proceedings of acm MM 2010 - WSM Workshop*, 2010.

[6] I. Li, A. Dey, and J. Forlizzi. A stage-based model of personal informatics systems. In *Proceedings of CHI 2010*, 2010.

[7] P. Sandhaus and S. Boll. Semantic analysis and retrieval in personal and social photo collections. *Multimedia Tools and Applications*, 51:5–33, 2011.

[8] F. Stutzman and J. Kramer-Duffield. Friends only: Examining a privacy-enhancing behavior in facebook. In *Proceedings of CHI 2010*, 2010.

[9] World Health Organization. Towards a common language for functioning, disability and health - ICF. published on the internet, 2002.