

MemoryBook: Generating Narratives from Lifelogs

Heather S. Packer
University of Southampton
WAIS Group
Southampton, UK
hp3@ecs.soton.ac.uk

Ash Smith
University of Southampton
WAIS Group
Southampton, UK
ads04r@ecs.soton.ac.uk

Paul Lewis
University of Southampton
WAIS Group
Southampton, UK
phl@ecs.soton.ac.uk

ABSTRACT

Automatic capture of life logging data can be extremely information rich, large and varied. Extracting a narrative from this data can be difficult because not all of the data is conducive to producing interesting narratives. Life logging data can be enriched by linking to the Semantic Web and narratives can be enriched with data extracted from Semantic Web knowledge stores. In this paper, we present MemoryBook which is a web interface that automatically generates narratives from life logging data in RDF form and from Semantic Web knowledge stores, and highlights maps and images which are associated with events in the narrative.

Categories and Subject Descriptors

H.1 [Models and Principles]: General

General Terms

Human Factors, Theory

Keywords

Narrative, Life Logging, Semantic Web

1. INTRODUCTION

Recently life logging has become more popular because of the availability of devices, such as Microsoft's SenseCam [5], DejaView [3], smart phones, tablets and computers. Each device captures a range of data continuously, and can include time, geographical co-ordinates, images, sound level and the people which were present (with the use of networks). In order to make life log data more accessible to the average person, we present an interface that generates a narrative from a person's data, which highlights relevant information associated with an event, such as geographical locations and or images. In more detail, a life logger with RDF data can use our web tool to select a range of days, weeks or months to generate narratives. Our tool outputs a web page that

displays the narrative, related images, maps, weather and hyperlinks.

We contribute to the state-of-the-art by enriching personal narratives with information from the Semantic Web, namely DBpedia, GeoNames and YAGO2. We also highlight multimedia and hyperlinks associated with the events and entities within a narrative to allow the viewer to clearly see relevant information. For example, if the mouse cursor is hovered over a reference to a person it highlights any available links to the person's homepage and or photographs of them.

The next section describes related work on events, purpose built life logging devices and narrative generation for large scale data. Following that sections 3, 4 and 5 describe the data collection methods and format, how we generate the narrative with information from the Semantic Web, and MemoryBook's interface used to generate and display the narrative, respectively. In the final section we conclude and discuss future work.

2. RELATED WORK

Narratives are a sequence of events describing entities or characters, therefore there has been a focus on dividing life logs into events. For example, the Stories Ontology¹ [6] which categorises a story as containing a sequence of events, can be used to markup lifelog data. Doherty and Smeaton [4] also divide lifelogs into discrete units or events. Therefore, in our approach we aim to indicate relevant media specific to an event.

Small scale studies have been used to generate narratives using images from the SenseCam, which records images, their GPS co-ordinates and timestamp, every 30 seconds or when readings from the light or sound sensors dramatically change. The timestamp and location provide a 'picture' narrative based on a map. This narrative is limited because it does not offer an insight into its composition. However, we propose that it could be enriched with additional life logging devices or information from the Semantic Web. Similarly, DejaView also uses a camera and sensors but in addition detects and recognises people and places on the fly to aid memory.

The work of Appan et al. [1] explores the composition of digital narratives for everyday experiences using photos. They note that traditional narrative models, like those used in cinema or in the Agent Stories framework [2] seem to be unsuitable for expressing everyday experiences and that users do not want to edit or author their narratives. Therefore, they propose using an emergent story framework that

¹Stories Ontology: <http://cpmtextris.net/stories/>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

NHT'12 @ HT'12, June 25, 2012, Milwaukee, US.
Copyright 2012 ACM X-XXXXX-XX-X/XX/XX ...\$10.00.




Figure 1: Imouto showing a summary of an event.

evolves through user interaction and feedback, that uses low-sampling of media to address the problems with large data sets.

3. LIFE LOG DATA COLLECTION

Our primary lifelogging system is Imouto (see Figure 1). Its mobile collection system collects sensor data from a user's mobile phone or PDA. This includes GPS, nearby Bluetooth devices, wifi devices and other temporally indexable information stored on the user's device, such as calendar and message metadata. This is combined with information collected from the web, such as local weather observations and public event schedules. Finally, the raw data is processed by the viewer application, split into separate events and combined with linked open data to produce a rich visualisation of the events in the subject's day.

The system currently infers the locations of the subject and the people encountered. Location is determined using GPS where available, or known static Bluetooth devices where it is not. For example, if an office worker is at her desk, her phone may not be able to get a GPS signal, but will probably be able to discover her PC via Bluetooth. If she has previously told the system the location of her PC and its Bluetooth MAC address, the system knows where she is. Known people are determined in a similar way using the Bluetooth addresses on their mobile phones, but also using face recognition in the user's photos. If a photo is deemed to depict a contact, and that photo also contains time and date metadata, it is assumed that the user was with that particular person at that time.

4. EVERYDAY NARRATIVES

In the following two sections, we describe how we generate a narrative from the lifelog data and how we enrich it via the Semantic web.

4.1 Narrative Generation

Our current system generates two forms of narrative; textual and RDF. The textual representation is a simple prose generator which generates sentences based on boilerplates depending on how much information is available for a particular event. The RDF is based on the Event² and Time-

²Event Ontology: <http://motools.sourceforge.net/>

line³ ontologies, as well as FOAF⁴, Dublin Core and the more standard RDF ontologies such as RDFS and OWL. Firstly our data is split into significant events based on collected Bluetooth information; a sudden change in the set of nearby Bluetooth devices will trigger a new event [7]. The event thresholds are then refined slightly. As Bluetooth scans may only realistically take place once a minute or so, it is possible that the user has already left a location and is in motion when the device is deemed to be 'out of range'. We refine the Bluetooth information with GPS log data. If there is a significant change in speed within a minute of an event threshold, it is 'snapped' back to the point at which the GPS data changes. This ensures that an event during which the user travels from one place to another actually contains a reference to both places, something which is important when generating RDF descriptions of events, as will be explained later.

In order to generate a textual description of the event, the system contains a set of sentence templates. A suitable structure is selected based on variables such as the number of places visited, distance travelled, and number of people encountered. Examples include:

The 125 minutes between 19:35 and 21:40 were spent in or around Kai Mayfair.

Between 12:30 and 1:00 was spent in the company of Jon and Dave. The weather was fair.

In addition to a human-readable text description of the event, a much richer narrative is available as RDF data. This is generated using the Event and Timeline ontologies. Each event is given a URL with the domain imouto.org.uk, and an RDF description is generated for each. Temporal properties and relations to other events are described using properties from Timeline, and a better description of the event, including links to the people and places involved, are described using Event.

The Event ontology makes available the properties `event:agent`, `event:place` and `event:factor`. Imouto uses these to indicate people encountered, places visited and places passed, respectively. A place is considered 'visited' if it is the only place encountered within the timespan of the event, or it is the first or last entry in a string of locations, which indicates a journey to or from the place in question. An event may also have other `event:factor` properties, such as the UK postcodes of places visited.

Finally, each event exported as RDF is given an `rdfs:comment`, the textual description generated previously. Also, if a photo depicting the event exists, it is added to the description using the `foaf:img` property.

4.2 Semantic Narrative Expansion

In order to enrich the generated narrative, we expand it with information taken from the Semantic Web knowledge stores YAGO2, DBpedia and GeoNames. We identify the entities within the narrative and lookup whether they exist in YAGO2, which is derived from Wikipedia, WordNet and GeoNames. Most importantly, it contains a large number of entities (people, place names and places of interest) and

³Timeline Ontology: <http://motoools.sourceforge.net/timeline/timeline.html>

⁴FOAF: <http://www.foaf-project.org/>

property	mustache template
http://dbpedia.org/ontology/currency	<code> {{entity}} uses the currency of {{value}}.</code>
http://www.mpii.de/yago/resource/hasWikiCategory	<code> {{entity}} is a {{value}}.</code>
http://www.geonames.org/ontology#population	<code> {{entity}} has a population of {{value}}.</code>
http://dbpedia.org/property/location	<code> {{entity}} is located in {{value}}.</code>
http://dbpedia.org/ontology/largestCity	<code>The largest city in {{entity}} is {{value}}.</code>
http://www.mpii.de/yago/resource/hasPreferredName	<code> {{entity}} is also known as {{value}}.</code>
http://dbpedia.org/property/exhibits	<code> {{entity}} has the exhibits {{value}}.</code>
http://dbpedia.org/property/motto	<code> {{entity}} has the motto "{{value}}".</code>
http://dbpedia.org/ontology/areaTotal	<code> {{entity}} covers a total area of {{value}} square meters.</code>
http://dbpedia.org/property/region	<code> {{entity}} is located in the region of {{value}}.</code>

Table 1: 10 templates from the semanticnarrative ontology.

has sets of alternative names for the entities, which means that it is particularly suited for disambiguation purposes. Once we have identified whether the entities are within the narrative, which we verify against geographical co-ordinates and links collected in the lifelog when available. So far we aim to enrich an entity with a maximum of three sentences, the number of sentences added depends on the information available in the knowledge stores. We use the following three steps to generate these possible sentences:

1. We extract the entity’s abstract from DBPedia and if it contains the word ‘is’ then we use it to describe the entity in the narrative.

2. We extract facts from the entity’s properties in DBPedia, YAGO2 and GeoNames. These facts are used to generate a sentence, the number of clauses in a sentence depends on the number of facts available. If there are more than three facts then we select three facts randomly, and if there are fewer then we select the appropriate template from below.

- (a) $\{Entity/It\} \text{ clause}^1$.
- (b) $\{Entity/It\} \text{ clause}^1 \text{ and clause}^2$.
- (c) $\{Entity/It\} \text{ clause}^1, \text{ clause}^2 \text{ and clause}^3$.

where, the first variable is selected at random and can be either the name of the entity or ‘It’. The clause variables are generated using an ontology, where key properties in the semantic knowledge stores and clause templates have manually been defined. The ontology contains one class, *clause* which has two properties *uriproperty* and *template*. This class describes a template variable and where the URI represents a property in an ontology and provides the parameters for the template. This ontology is detailed below in RDF/XML form. We use mustache⁵ to create our templates (see Table 1 for 10 examples from our ontology). For example, the sentence “Bob went to London Zoo” contains the entities Bob and London Zoo, we lookup these entities in YAGO2 and find that Bob is a name and we can’t disambiguate him, therefore we only use properties from DBPedia, YAGO2 and GeoNames to enrich the narrative about London Zoo.

⁵mustache <http://mustache.github.com/>

```
<owl:Class rdf:about=
  http://users.ecs.soton.ac.uk/hp3/
  semanticnarrative.owl#clause"/>

<owl:DatatypeProperty rdf:about=
  http://users.ecs.soton.ac.uk/hp3/
  semanticnarrative.owl#template"/>

<owl:DatatypeProperty rdf:about=
  http://users.ecs.soton.ac.uk/hp3/
  semanticnarrative.owl#uriproperty"/>
```

3. We then add information into the narrative about personalised facts concerning the entities described in the lifelog (see Table 2). For example, if an entity is the tallest or the furthest away from home that you’ve travelled then it is added to the narrative. In order to calculate this we identify all the entities contained in the lifelog and use DBPedia to compare interesting properties about them. We also use GeoNames to calculate geographical statistics. In order to add facts into the narrative we either take one of three approaches:

- (a) If there are more than two facts then we random select two facts and generate a sentence using the template $\{Entity/It\} \text{ clause}^1 \text{ and clause}^2$.
- (b) If there is only one fact and there are less than two clauses in the sentence generated in step 2 we append the fact as the last clause using the last two templates in step 2.
- (c) If there are three clauses in step 2, then we generate a sentence using the template $\{Entity/It\} \text{ clause}^1$ where the first variable is random and the clause is the personal fact.

The sentences generated in the above steps are added the first time an entity is mentioned in the narrative. In addition to these sentences, we add a paragraph at the beginning of the narrative summarising the day by using templates and variables describing the area and the places that were visited, and we also calculate the distance travelled during the day. Specifically, we use the template *The day was spent in {area} {areaClause}*. In particular, the day was spent in {places}, and a distance of approximately {x} miles were travelled., where the variable *area* is describes the general

Fact	Source
is the tallest building you've been in	DBpedia
is the furthest from home you've been	GeoNames
is the furthest you've travelled in a day	GeoNames
is the most north/south/ west/east you've been	GeoNames
is the oldest building you've been in	DBpedia
is the highest/lowest elevation	GeoNames
is the most sparsely/densely populated city you've been in	GeoNames

Table 2: Personalised Facts

area (i.e. London, Hampshire, North England, etc), and *areaClause* is a template from our ontology, *places* is a list of places visited in the day, and *x* is the number of miles travelled.

In our London Museum example, our approach enriches it so that it reads (where the added narrative is in bold):

The day was spent in London the UK's capital city. In particular, the day was spent in Westminster and Kensington and Chelsea, and a distance of approximately 4 miles was travelled. Generally, the weather was mostly sunny.

Between 8:15 and 9:30 was spent in or around the Dorchester Hotel. The Dorchester Hotel is an Art Deco building in London.

Between 9:30 and 10:00, the 1.8 mile journey from the Dorchester Hotel to The Natural History Museum was made. The Natural History Museum is one of three large museums on Exhibition Road, South Kensington, London, England. It is a Grade I listed building in London. The 120 minutes between 10:00 and 12:00 were spent at the Natural History Museum. Between 12:00 and 12:35 was spent in or around the Restaurant at the Natural History Museum.

Between 12:40 and 12:45, the 0.1 mile journey from the Natural History Museum to the V&A Museum was made. The V&A Museum, also known as the Victoria and Albert Museum, is the world's largest museum of decorative arts and design, housing a permanent collection of over 4.5 million objects. It is a museum in Kensington, was established in 1852 and is the oldest building you have ever visited. Between 12:45 and 15:10 was spent in or around the V&A Museum. Between 15:15 and 16:00, the 1.8 mile journey between the V&A Museum and the Dorchester Hotel was made.

Between 19:05 and 19:20 the 0.3 mile journey to Kai Mayfair was made. The Kai Mayfair is a Chinese restaurant established in 1993, it has a Michelin Star. The 125 minutes between 19:35 and 21:40 were spent in or around Kai Mayfair.

5. INTERFACE

The user is able to select a group of dates from the first page (see Figure 2). Our website then generates the narrative using the approach in the previous section. This narrative is displayed alongside a map with all the entities that have geographical coordinates and images from the lifelog. It also retrieves the weather from a web service, and displays a weather symbol and the temperature (see Figure 3).




Figure 2: MemoryBook Interface

When a user hovers their mouse cursor over events or entities in the narrative then it is highlighted in blue and white, respectively. Once an event or entity is hovered over, all of the links, images and map markers that describe it are highlighted too (see Figure 4 where the Natural History Museum event and Dorchester Hotel entity are highlighted).

6. CONCLUSIONS AND FUTURE WORK

In this paper, we present a technique for semantic narrative expansion, situated within the context of automatically generated lifelogging narratives. We utilise the semantic knowledge stores YAGO2, DBpedia and GeoNames to expand narrative with additional facts about people, places and things. We also present an interface called MemoryBook for interacting with the semantically expanded narrative, which displays maps, images, and links.

In this work we have encountered the need to decide how much information to add. Now for we have decided to use a maximum of three sentences to describe an entity, for future work, we want to create a feature that can make the narrative more verbose, by including interesting information by determining peoples' interests from their lifelog.

We propose that this type of narrative generation lends itself to aiding memory, publishing blogs, and generating digital postcards. For future work, we would like to look at generating our narrative based on different audiences. For example, the narrative for a work and personal require different types of information but both would benefit from images and links.

7. ACKNOWLEDGMENT

The described work was funded by the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 270239 (ARCOMEM) and 231126 (Living-Knowledge) together with the LiveMemories project, graciously funded by the Autonomous Province of Trento (Italy).




Figure 3: MemoryBook: London museums example

The first author is also grateful for funding from EPSRC in the form of a Post Doctoral Award.

8. REFERENCES

- [1] P. Appan, H. Sundaram, and D. Birchfield. Communicating everyday experiences. In *Proceedings of the 1st ACM workshop on Story representation, mechanism and context*, SRMC '04, pages 17–24, New York, NY, USA, 2004. ACM.
- [2] K. M. Brooks. Do story agents use rocking chairs? the theory and implementation of one model for computational narrative. In *Proceedings of the fourth ACM international conference on Multimedia, MULTIMEDIA '96*, pages 317–328, New York, NY, USA, 1996. ACM.
- [3] D. de Jager, A. L. Wood, G. V. Merrett, B. M. Al-Hashimi, K. O'Hara, N. R. Shadbolt, and W. Hall. A low-power, distributed, pervasive healthcare system for supporting memory. In *Proceedings of the First ACM MobiHoc Workshop on Pervasive Wireless Healthcare*, MobileHealth '11, pages 5:1–5:7, New York, NY, USA, 2011. ACM.
- [4] A. R. Doherty and A. F. Smeaton. Automatically segmenting lifelog data into events. In *Proceedings of the 2008 Ninth International Workshop on Image Analysis for Multimedia Interactive Services*, WIAMIS




Figure 4: MemoryBook: Dorchester entity selected

'08, pages 20–23, Washington, DC, USA, 2008. IEEE Computer Society.

- [5] 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 P. Dourish and A. Friday, editors, *UbiComp 2006: Ubiquitous Computing*, volume 4206 of *Lecture Notes in Computer Science*, pages 177–193. Springer Berlin / Heidelberg, 2006.
- [6] M. O. Jewell and C. J. Hooper. Linking locations: Storytelling with pervasive technology. In *Narrative and Hypertext Workshop at Hypertext'11*, June 2011.
- [7] A. Smith, K. O'Hara, and P. Lewis. Visualising the past: Annotating a life with linked open data. In *Web Science Conference 2011*, Koblenz, Germany, June 2011.