“Let’s share a story”: Socially-enhanced multimedia storytelling

Mu Mu, Steven Simpson, Nicholas Race (Lancaster University, United Kingdom)
Omar Niamut, Gijs Koot, Rianne Kaptein (TNO, the Netherlands)
Jacco Taal (Bitnomica, the Netherlands)
Luigi Mori (ADB Broadband, Italy)

Abstract

User-generated audio-visual content is becoming the most popular medium for information sharing and social storytelling around a live event. This paper introduces an online multimedia storytelling eco-system comprised of purpose-built user applications, a collaborative story authoring engine, social context integration, and social-aware media services. The system enables online collaborative story co-authoring and provides an ideal platform to study the synergy between social networks and networked media in enhancing the user experience of storytelling as highlighted in our event-based user experiments.

Keywords

Creative media, storytelling, social analysis, multimedia authoring, community, live events, user experiment, caching.

Introduction

With their increasing user populations, social networks have evolved from pure text streams to photo albums and timelines towards richer mediums for social sharing using audio-visual content. The growing popularity of capturing and sharing user-generated video content can largely be attributed to the increased availability of video recording capabilities on consumer devices and their integration with online video sharing platforms. Amateur video capturing has also evolved from personal chronicles to citizen journalism, collaborative creation, and storytelling of live social events [Macro2009]. Unlike video mash-up systems [Saini2012], multimedia storytelling engages a much higher level of interaction during content capturing, sharing and editing. Recent studies also see trends of exploiting shared content (from other users) in providing additional perspectives to improve composite stories [Guimar2011]. Recent years have also seen related work on event detection and media annotation to improve content discovery and storytelling. The EventMedia project adopts an event-based approach for users to discover events through different media using Linked Data [Khrouf2014]. Seen is an online application that promotes trending events detected from multiple social media in a shared space [Seen]. Steiner

introduced his work on Wikipedia Live Monitor that automatically creates media galleries to illustrate live events [Steiner2014]. Becker et al. investigated how to associate social media documents with planned events automatically and effectively [Becker2012]. Furthermore, Chi et al. structure user stories around a text chat, augmented by an agent that continuously interprets the text to suggest appropriate media elements to illustrate the story [Chi2011]. Smits et al. also outlined a concept based on multimedia content analysis and crowdsourced tagging to enable intuitive, nonlinear access to large video collections [Smits2010].

Most of the aforementioned work exploits computer algorithms or systems to synthesize social stories from linked social medias such as Twitter and Wikipedia where humans take the roles of crowdsourcing annotator, providers of raw materials, or simply the audiences of stories. However, narrative properties are often discussed through their affective correlates and the corresponding experiences and are not constrained by any structure story models. Reflecting the notion of MM-hard, which refers to multimedia problems that require human-level insights and perception that can't be realized with a single algorithmic approach [Xie2014], Kelliher envisages the departure from the human's role as primarily increasing an algorithm's efficiency or facilitating a transaction, and considers the human also as an active and subversive force [Kelliher2014]. There have been a growing number of online storytelling and video editing systems ([Wevideo], [Dale2011], [Mu2013a]) that enable creative individuals or communities to stay at the center of the storytelling process and to become the editors of creative and collaborative storytelling using audio-visual content. The result is a more personalized experience being encapsulated in a dynamic and open story format.

This paper emphasizes the system aspect of storytelling, especially the designs and innovations that facilitate organic story authoring conducted directly by human. We introduce an online multimedia storytelling system designed to facilitate creative story authoring and sharing using user-generated audio-visual content in communities of social events. The genre or format of the story is not limited by the design of the system, though we use narratives of personal experience in social events as the reference scenario. Other applicable use scenarios include citizen journalism, self-promotion, interactive exhibition, collective awareness, etc. The system integrates purpose-built user applications, media services, a story authoring engine (based on an early prototype introduced in [Mu2013a]), and content-distribution mechanisms, which allow us to study and experiment with the entire eco-system of online storytelling. The system also incorporates social context derived from the analysis of social networks to improve user experience in creative editing. The design of the storytelling system employs the synergy between social context information and multimedia delivery applications. Media applications and services assist audio-visual social interaction, whilst the underlying system exploits social context to improve media retrieval and delivery.

Use case and requirement analysis

Designing an online multimedia storytelling system to support social interactions related to a live event is nontrivial. To better understand the user requirements in storytelling and how they can be met by an eco-system, we developed a use case based around a real-world event, the *Schladming Nightrace.* The nightrace is one of the classic events of the International Ski Federation World Cup annually hosted

in the ski-town Schladming at the foot of Alps. The event is famous not only for its popularity (over 50,000 visitors) but also the big party hosted at the finish line. Our use case follows two participants who attend the Nightrace and use their mobiles to capture moments of the trip to share with friends at home. At the hotel, they collaborate with these friends to add their daily highlights to a story, including footage from other participants and sources connected with the event. The system they use also allows them to make simple edits such as slicing and reordering to make the story more compact and interesting (Figure 1). This way, their joint story covers the experience of a group, and progresses along with the live event. Although the use case may sound simple and natural, its subtleties lie in a complete eco-system assisted by technology advancement. We summarize the requirements/challenges of an online multimedia storytelling system:

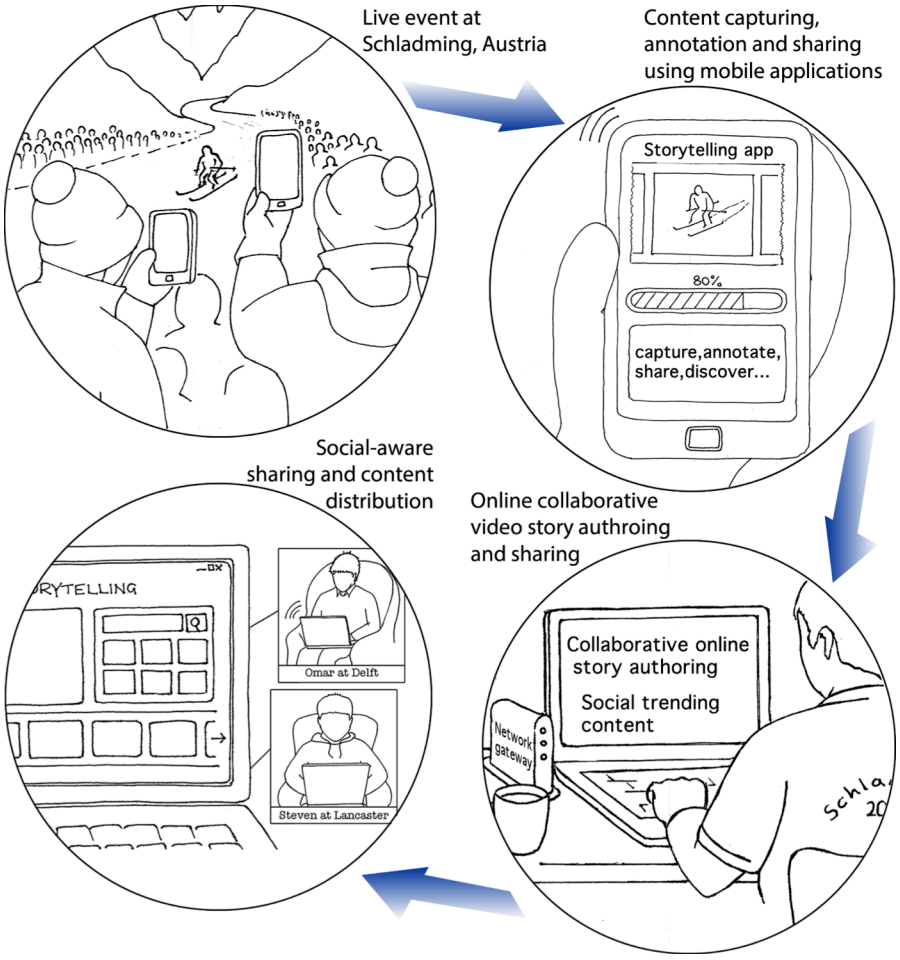

Figure 1 Online community storytelling use case

- A mobile application to streamline the capturing, tagging, sharing and browsing of user content. The application should also capture metadata such as the geographical location of the user device at the time of capturing to improve creative editing and social sharing.

- An online editing tool for creative story authoring and sharing. Conventional video editing tools such as Adobe Premiere provide rich editing tools and effects for offline professional editing. For online storytelling, story editors can benefit more from direct referencing and editing of videos online (without having to download them first).
- A "lightweight" multimedia story-authoring engine that has a small footprint on user devices and networks. An online story-authoring system should aim to make online video editing as easy as editing a shared text document and minimize the processing and network load at user devices.
- Integration of social context. Widely adopted social networks greatly influence user preferences and activities. The rich context information and social atmosphere embedded in social media is crucial in improving user experience within social applications such as online storytelling, particularly when the stories involve live events.

Online multimedia storytelling system

System overview

Figure 2 gives an overview of our social-aware online multimedia storytelling system. Two user applications, a web portal and a mobile application provide the user-facing elements including video capture, sharing, searching and story authoring. The user applications also capture live metadata such as location information, which is exploited to provide personalised media experiences. The multimedia processing function encompasses operations such as transcoding, chunking, content analysis, and indexing so that user-generated content can be better shared and discovered for story authoring. The story-authoring engine uses a unique and lightweight "edit by reference" design to enable collaborative online story editing. It uses edit-decision lists to define time-addressable references of user-generated content so that user stories can be made by manipulating hierarchical references to static video objects without any complex video rendering.

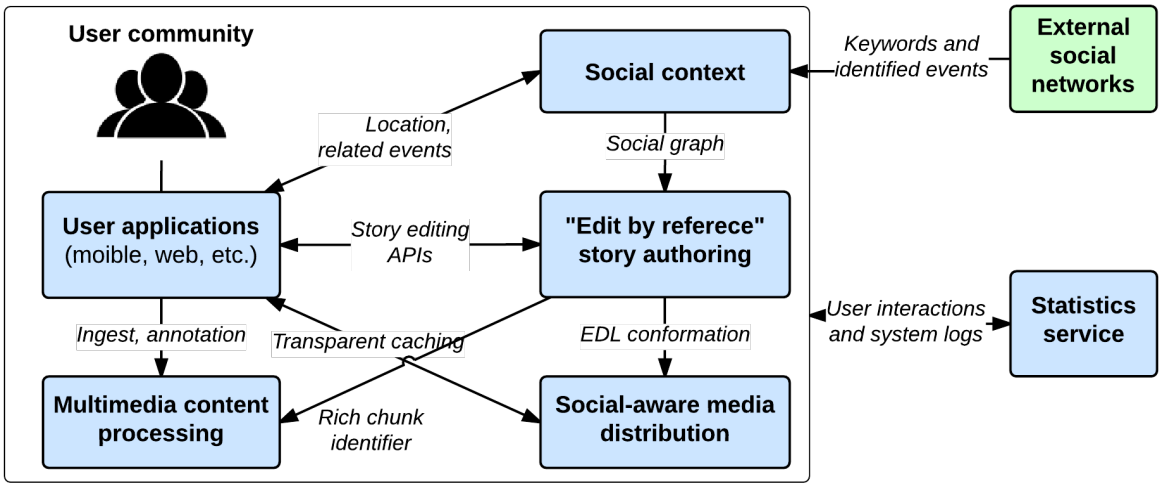

Figure 2 Online multimedia storytelling system

The social-context integration function monitors related events and trends in social networks in order to 1) improve the user experience in tagging user creations and 2) discover shared stories related to trending topics nearby. The social-aware media distribution function is introduced to improve the network efficiency and user experience of media distribution during story editing and sharing. As a research platform, the storytelling system is also equipped with a bespoke statistics service, which captures time-coded user activities as well as service status.

Multimedia content processing

Compared with conventional social sharing services such as Twitter, multimedia storytelling requires a fast and resilient audio-visual processing backend to facilitate sharing and delivery of user-generated content. We have developed a set of multimedia processing tools (known collectively as *Mediaplex*), to prepare multimedia content and associated metadata for other storytelling functions. It automates transcoding, meta-data extraction and management, audio-visual content analysis and content chunking to process raw media objects received by the system. The media objects prior to any story editing are also known as the *rushes*.

Mediaplex presents an ingest service for users to upload content via user applications. Uploads are transcoded using five pre-defined profiles (various bitrate and resolution), each targeting a specific playback scenario. Transcoded content is then split into a series of chunk files for quality adaptation when network bandwidth fluctuates (a mechanism similar to MPEG DASH [DASH2012]) and for partial object caching.

A wide range of metadata is also maintained by Mediaplex to help social sharing and to enable flexible definition of media assets within a media object (such as a 10 second scene in a 2-hour long sport event). Most metadata is provided explicitly by users (such as video tags) or user devices (such as geo-location of the captured video), while other metadata is derived using video analysis functions such as scene detection, which separates interesting shots from others. Currently supported metadata

includes: title, description, tag, and copyright information applied to the rushes as a whole, as well as the time-coded social annotation, GPS coordinates, and scenes.

"Edit by reference" story authoring engine

One unique feature of the online storytelling system, designed to facilitate community online co-authoring, is the "edit by reference" design. Instead of physically editing and rendering a video object, which consumes computational and network resources for every edit, we allow a user story to be composed by creating a list of references to existing video objects. The idea is similar to creating a playlist where each line of the list refers to a section of a song stored externally in the cloud. Editing a story simply maps to editing a small "playlist".

The Media Asset Referencing System (MARS) is a media management system that recognizes time-addressable media assets within chunked media objects, and generates new media objects on-the-fly by combining just the segments that represent time periods that fall within or overlap with those assets. The core of MARS is the management of edit-decision lists (EDLs). An EDL is a script-based expression of composition of a media object. It creates a time-addressable descriptive view on the internal structure of media objects using frame-accurate presentation and navigation of time-codes, so that virtual operations visually assisted by keyframes such as cut and fast-forward can be carried out without physically examining the corresponding media objects.

When a chunk is created from a media object, an EDL generator defines a rich chunk identifier (RCI) to index the chunk and uses it to generate a rushes EDL for the media object. Parameters within the RCI describe the media encoded in the chunk. A rushes EDL (which again has its own identifier) then combines multiple tracks as shown below. Each element identifies a source object (*src*), one or more time-coded periods of that source (*select*), which channels of the source are to be used, and which ones they map to in the described object (*channels*), and which single time period of the described object the selected periods contribute to (*range*).

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<edl xmlns="http://scc.lancs.ac.uk/ns/edl/2012">
<channel name="1"/>
<channel name="0"/>
<!-- video track -->
&lt;segment channels="0:0" range="0-31000/2997" select="200-31200"
src="chunk_670b216b16eed5d44d1c49edd2a7b8642bda89a9_v40-B_24+1f6d51-
S_1+310-R_2997-D_209-Z_+acb0-T_2b4+0-C_2cc+0-Y_a99c_10"/>
&lt;segment channels="0:0" range="31000-40900/2997" select="200-10100"
src="chunk_670b216b16eed5d44d1c49edd2a7b8642bda89a9_v40-B_1f6d75+2475b3-
S_311+99-R_2997-D_209-Z_+b188-T_2b4+310-C_b34+1-Y_a9c4_3"/>
…
<!-- audio track -->
&lt;segment channels="0:1" range="0-1536/175" select="0-387072"
src="chunk_670b216b16eed5d44d1c49edd2a7b8642bda89a9_a30-B_24+10085-
S_1+378-R_44100-D_1eb-Z_+28a-T_246+0"/>
&lt;segment channels="0:1" range="32256-40960/3675" select="0-104448"
src="chunk_670b216b16eed5d44d1c49edd2a7b8642bda89a9_a30-B_100a9+100c2-
S_379+102-R_44100-D_1eb-Z_+872-T_246+378"/>
…
</edl>
```

Using the story composer on a storytelling application, a video story can then be expressed as a story EDL that references defined segments in one or more rushes EDLs, or even the EDLs of other derived content; story editing then boils down to the editing of text-based scripts. At the time of playback, the story EDLs are converted into a directly playable media manifest on-the-fly. An EDL can be converted to MPEG-4 first by resolving all segments so that only segments referring to chunks remain (*src* contains an RCI). An MPEG-4 track is created for each EDL channel. Each chunk is allocated a non-overlapping time period within the track, then a description of it is derived from its RCI parameters, and encoded as part of the track description. An MPEG-4 edit list is included in the track to indicate which parts of the track are actually to be played and when. The URIs of the chunks are also derived from their RCIs, and listed in the track for the player to fetch. The resultant manifest file is an MPEG-4 header that describes the audio-visual file structure and includes URIs to locate the corresponding chunks in the network. The process requires no transcoding, and a virtual layer in the EDL hierarchy allows all profiles of a rush to be referenced together, one being chosen only during resolution; the results of editing are thus available at once, and at all quality levels.

Social context integration

When recording and sharing video items, the generation of metadata can significantly support the editor task of relating video times to one another. For online collaborative storytelling, the annotation of user content is often an onerous task, despite it being crucial for social sharing and content. The storytelling system adopts a bespoke social analysis module, the Adaptive Event Profiler (AEP), which analyses social medias and exploits relationships between recognized keywords from social context, to assist content annotation and content discovery by event detection.

AEP generates a list of related keywords, given one or more main keywords about a particular event. The social network Twitter is used as a source of textual context to find related keywords for video tagging. Tweets provide a valuable source of real-time information for tag suggestion. The AEP starts with the name of the location of the event, or the first tags assigned by users. A list of related keywords is created, including normal terms and Twitter hashtags. The AEP distinguishes internally three types of keywords:

- Main keyword(s), which are always followed;
- Related keywords, which have high relevancy to the main keywords;
- Candidate keywords, which have potential relevancy but not shown to the users.

Then the following process is repeated in regular time intervals t:

1. In period *t,* follow the main keywords and related keywords to retrieve a set of tweets. The Twitter Streaming API provides a feed of tweets that match one or more of the tracked keywords.
2. From those tweets, extract the top *n* most frequent terms and hashtags. These terms are the new candidate keywords.
3. In period *t* + 1, follow the main keywords, the related keywords and the top *n* candidate keywords to retrieve a set of tweets.
4. For each keyword, generate a Maximum Likelihood Estimation language model based on the tweets that contain the keyword.
5. For each candidate and related keyword, assess the similarity of its language model and that of the main keyword using Kullback-Leibler divergence.
6. Select the *k* keywords (either candidate or related) most similar to the main keyword, and designate them as related keywords for the upcoming period *t + 2*.
7. In period *t + 2* follow the main keyword and the *k* related keywords to retrieve a new set of relevant tweets.

A web service is also implemented to serve requests from the rest of the storytelling system for a ranking of events/keywords, based on an initial main keyword set, a geo-location or any combination of those. Only main and related keywords are ranked and exploited by storytelling functions. Examples of related system can be found in [Khrouf2014], [Seen], [Steiner2014], and [Becker2012].

User applications

Users interact with the storytelling system through two user applications: the Vidacle iOS application for video capturing, tagging and exploring shared content on mobile devices during the live event; and the Storiboard web application designed for creative and collaborative story editing and sharing beyond the live event.

Storiboard web application

The Storiboard web portal allows users to upload video clips or images via a standard web interface, and describe the uploaded content using pre-defined metadata such as title, description, tags, dates, and geographical location, which are analyzed and indexed so that shared content can be easily located for story authoring.

Users can easily start a new story, and allow other users to edit it collaboratively. A story is edited in the story building area, consisting of preview, search and storyboard panes. The search pane hosts a faceted search function with multiple filters. Users can search metadata entries for media assets by title, description, or tags. According to suggestions from user interviews, a diverse result set (including not only keyword-matching content but also content related to users' interests) could be beneficial to creative work. The search function adopts the social events and their correlation derived by the adaptive event profiler, so that search results are enriched by additional content (highlighted using green borders in Figure 3) that are socially related to the original search term, even if not an exact match.

Adding any content to the storyline is as easy as dragging it from the search results pane and placing it onto the "movie reel" in the story editor pane (Figure 3). The order of content on the reel can also be changed using the same operation. Cropping the added content can be easily done by adjusting the start and end handles in the editing pane. The pane also gives preview images for quick navigation within the video object, and includes preview images generated by Mediaplex during content ingest. Thanks to the *edit-by-reference* design, only references within the story's EDL are updated upon any edit, so the network only needs to transmit a few lightweight XML messages and preview images for a story to be edited.

Figure 3 Storiboard web interface

Audio narratives can be added to a user story via the HTML5 audio recording function, which is supported natively by many modern web browsers. Once a user finalizes a story (i.e., makes it public), sharing options are also available to encourage community forming and collaborative editing of a shared experience.

Vidacle mobile application

The Vidacle mobile application is designed with an intuitive user interface and content-management system to streamline capturing, tagging, sharing and exploring of community stories on the same mobile device during the live event. Once a capture is made, users are given the opportunity to type in a title,

description, and tags. Vidacle also manages the user videos, their associated annotations and other metadata such as geo-location in the background without further intervention.

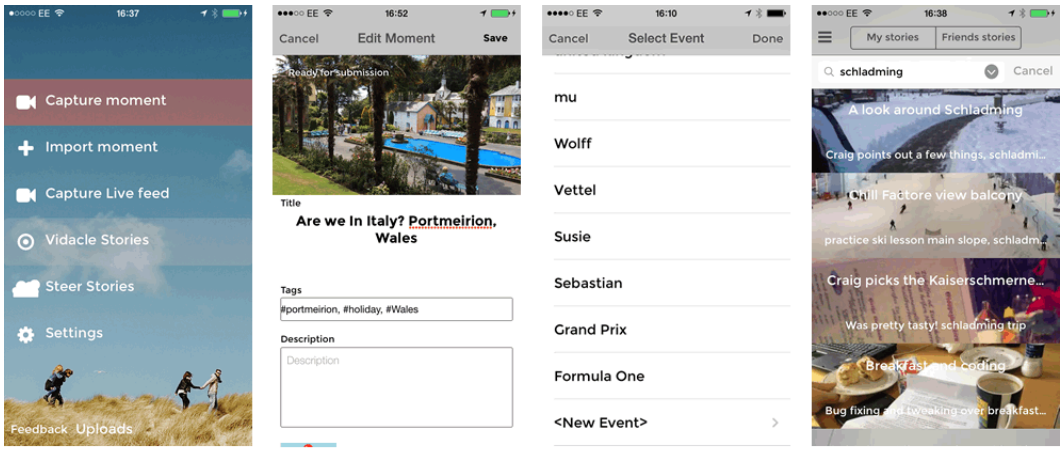

Figure 4 User-generate content capturing, sharing, annotation, and retrieval on Vidacle mobile application

User feedback from our previous experiments suggests that typing in a text annotation (i.e., title, description, tags, etc.) to accompany a video can be "*annoying especially when you are standing on a ski slope with many other things happening around you*". This illustrates the need for automatically gathering as much relevant information as possible and presenting the user with suggestions, making their life easier. Vidacle exploits the trending events derived by AEP in conjunction with the current location of the device to offer a number of suggested tags for users to select (Figure 4). This feature greatly reduces the overhead of video annotation between video captures, allowing users to focus on content creation during a live event and enabling quick social sharing. Stories made by a user's friends are also "pushed" to Vidacle in the "Friends' stories" tab, so storytelling users are always up-to-date with stories in their ego networks.

Social-aware media distribution

The distribution of media content is key in the overall user experience "in the wild". Cachability of content is a natural feature of the online multimedia storytelling design and the preparation of content for it. Having cached the chunks required to view one preview, a device will re-use many of those chunks on a subsequent preview. Furthermore, if a user is expected to require a certain piece of content soon (based on social relationships), a cache node near the user can preemptively obtain the necessary chunks.

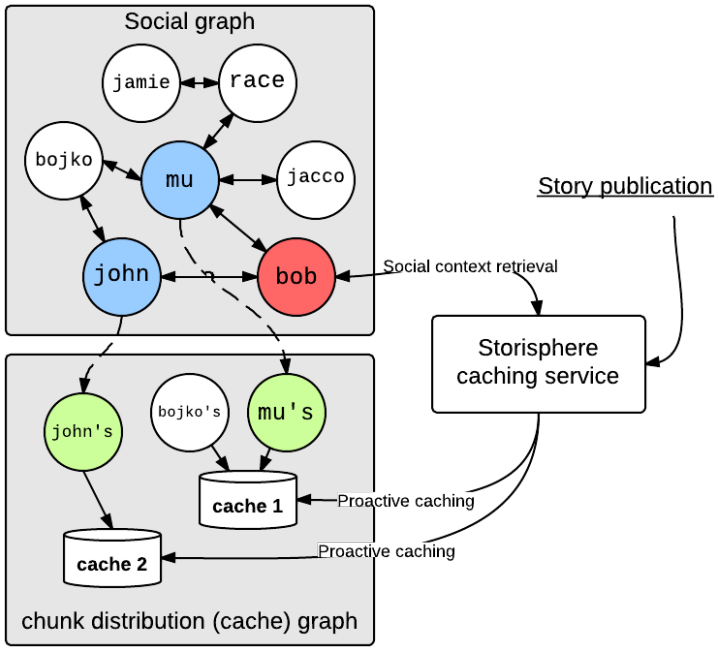

Figure 5 Social enhanced proactive caching

The caching function resides within networked caching nodes and operates as a proxy that resolves the story EDL of a user request down to required media chunks and fetches the cached chunks directly regardless of their previously associated stories. The social-aware caching function is enhanced by integrating and cross-referencing social and media distribution information. The caching service fetches the socially linked users of the story author by querying the *social graph*, a database that captures non-directional social connections (i.e., *friends*) (Figure 5). The caching service then queries the statistics service in order to identify the last known user device(s) and the network locations where users have last been seen. Using the IP address of the requesting client, the caching service employs the *chunk distribution graph* to name the most probable caching nodes to be exploited for the given user clients. Figure 5 illustrates how Bob's story is preemptively pushed to two cache nodes near his friends Mu and John. The cached content is also open to other users such as Bojko who may not be directly connected to Bob.

Storytelling experiments

An experiment was held at Silverstone, UK during the British Grand Prix Formula 1 racing event (http://www.formula1.com/races/in_detail/great_britain_924) on 6th July 2014. For this experiment, we selected two groups of members of the public as the participants. Group 1 is a family of three (plus one friend), who are long-term Formula 1 fans, on their first trip to the Silverstone GP, invited for Vauxhall

"VXR Power Events" which allows exclusive access to certain areas of the Silverstone circuit during the event. Group 2 is a family of two (plus one friend), who are frequent Formula 1 visitors. The two user groups do not know each other.

Both groups used the storytelling applications extensively to capture the highlights of their Silverstone experience. Group 1 shared 33 rushes (900MB) from one iPhone and 47 rushes (2GB) from a second iPhone. Group 2 returned with 89 rushes (4GB) and 104 rushes (4.6GB) from two iPhones respectively. This suggests that participants of live events are very likely to capture a large amount of audio-visual content (not only still images) for personal archiving or social sharing, especially when a group of friends and family travel together. The rushes cover the entire trip to Silverstone. We noticed that only a small portion of the rushes are about the actual racing, while most others are related to other parts of the Silverstone experience, including travelling, camping, and auxiliary entertainment.

The social context integration proved to be effective in improving video annotation and in enhancing the search function during the experiments. Using AEP to recognize trending events, the integration of social context ideally solves the classical "cold start" problem in content recommendation. Given user location, the storytelling system is able to suggest a number of popular user stories related to socially trending keywords nearby (for example, as shown in Figure 4). The main keyword used as input for the AEP in this case is "Silverstone". Event and participant names (such as *#britishgp* and *Hamilton*) are among the most popular suggestions during the event. For the enhanced search function, AEP provides a list of related keywords and a metric to quantify the relevance to the search request.

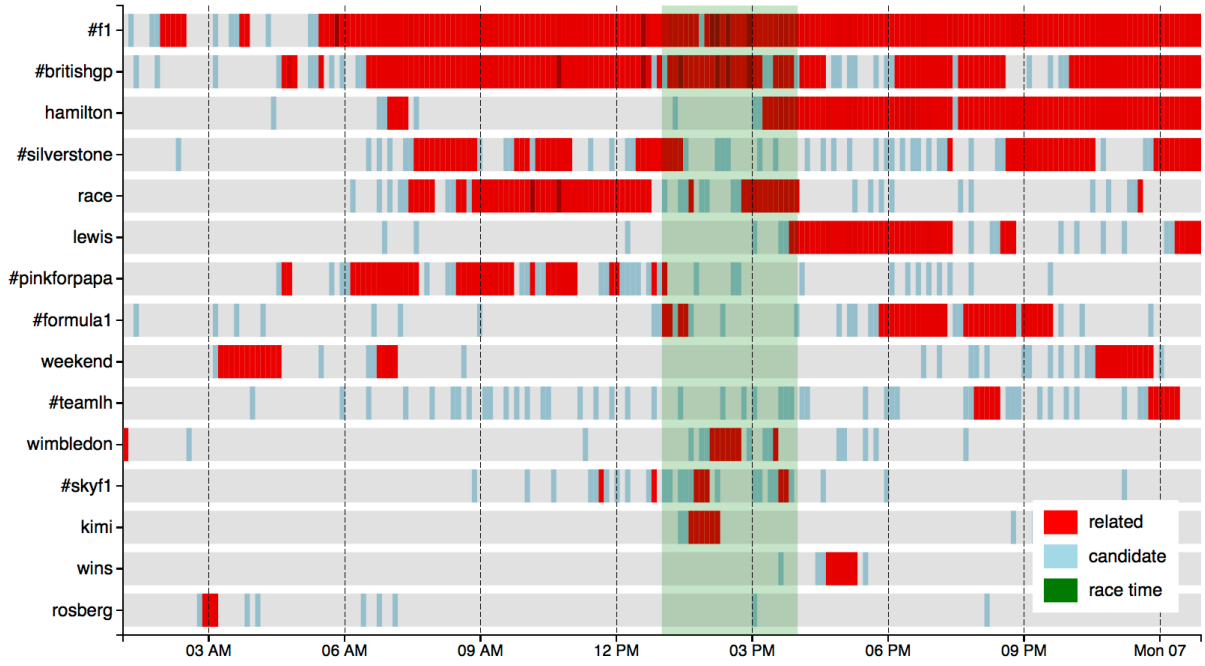

Figure 6 Heatmap of related social terms on the Silverstone F1 race day

Figure 6 shows a heatmap of related terms in a social network generated from data output by AEP on the F1 race day (6th July 2014). The darker the episode of the related keywords is, the higher its relevance to the main keyword "Silverstone". It illustrates how hot topics related to live events in social media evolve over time as influenced by how the event develops. For instance, *#pinkforpapa,* a story related to the British driver Jenson Button and his father, was a very hot topic of the Silverstone F1 prior to the actual race. This demonstrates how AEP identifies social discussions on background stories of the race. The hotness of the term '*kimi*' in the heatmap reflects our observations that the crash on the first lap quickly became a topic in social media but did not last further in the race when topics evolved to subjects around specific battles and the winner of the race ('*hamilton*','*#teamlh*',etc.). The related keywords were proven to be effective in helping participants discover and create personal stories in our system.

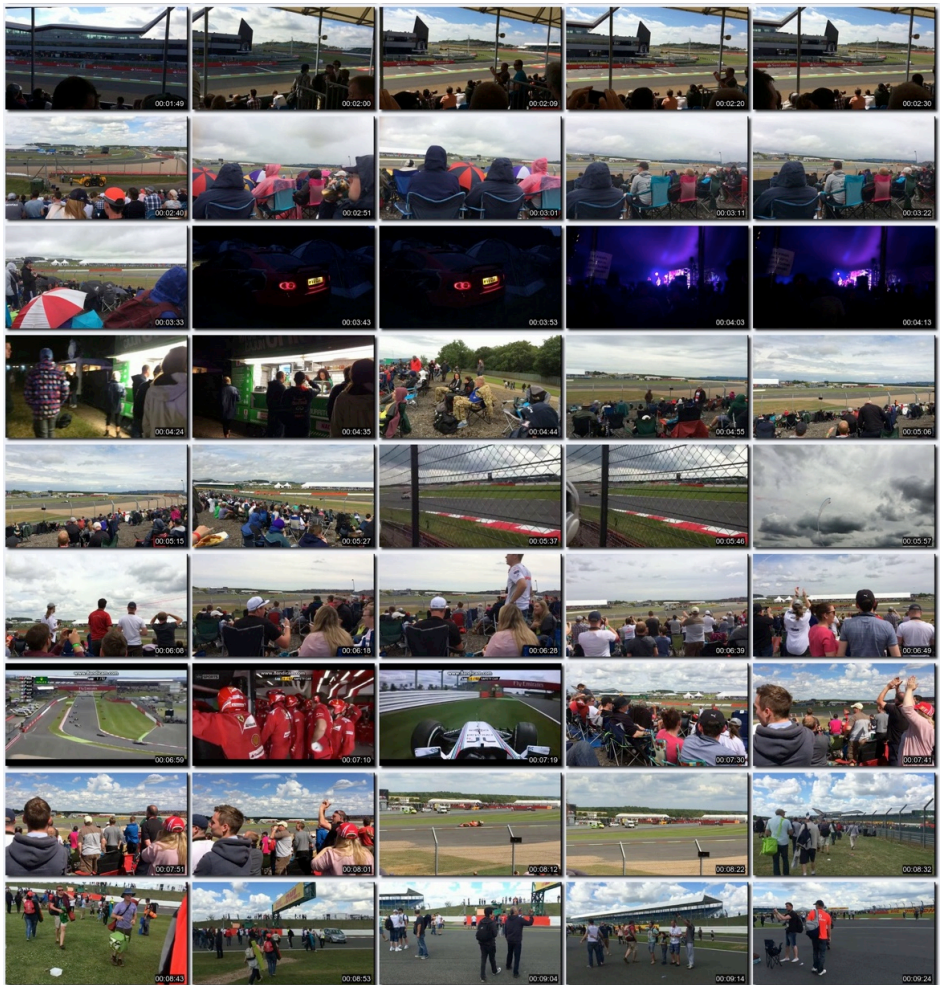

Figure 7 Silverstone F1 Storytelling experiments

During the course of the experiments, composite stories were created and repeatedly edited by many individuals. One example is the Silverstone F1 story (available on Youtube via

http://youtu.be/Zo83I1MPQcI), which demonstrates the experience of a group of participants through their personal and unique perspectives of the trip (Figure 7). The story was made by group 1 using 37 media assets and involved more than 100 revisions. It demonstrates how the storytelling platform assists users in creating a compact and engaging story from a large amount of user shared content. Most of the content used for the story was originally captured by the same user group, though the storyteller did adopt footage shot by user group 2 with a great viewing angle about an evening event, which both user groups attended coincidentally.

In the user interviews participants suggested "*the true best way of watching the F1 race is to 'sit at home and watch it on television'*". They continued by saying that "*the official broadcasters have the best access to all viewing angles so that viewers can keep track of incidents and accidents during the race as well as the background stories from reporters, while people at the Silverstone circuit normally have only the view of the race at a corner*". Our participants believe that "*the true F1 racing experience lies in the F1 atmosphere which gives you the experience of being with the crowd, enjoying the live sound of the F1 engine, walking on the F1 track, going to the evening events, etc*." The user feedback proves right the design principles of our storytelling platform, which is not made to replace or challenge conventional broadcasters but to assist individuals or small communities in recreating their personal experience by assembling pieces of highlights at a live event. Moreover, one consensus among participants of both groups is that storytelling of personal/group experience of an event is "*a very natural thing to do*". Most found that using the storytelling system for capturing and sharing their own creations made them feel that they were "*telling a live story to their friends*". They were mostly adding the narratives directly while recording. Sometimes a group member spontaneously acted like a reporter, letting other group members talk about what had just happened.

Conclusions

Recent years have seen the medium of social sharing evolving from text streams to user-generated audio-visual content. Following this trend, collaboratively authoring a story of live events by joining shared multimedia content with different perspectives is becoming popular. This paper introduces an online multimedia storytelling eco-system that enables the capturing, sharing, and authoring of user stories online using a unique edit-by-referencing design. The user experiment highlighted the contribution of social features in enhancing the user experience in content annotation, retrieval and media distribution as part of a storytelling eco-system. Through waves of user engagement and public experiments, we learnt that users generally found fabricating stories using video as the medium very intriguing compared with the conventional message- and image-based social media. We have also witnessed the unique role of user creativity in the entire development cycle of engaging social stories. Furthermore, social context derived from social media, location-based services and emerging mobile technologies can also greatly improve the creative story capturing and authoring process. Future work will look into addressing social challenges and collective awareness using the creative storytelling platform.

Acknowledgement

The work presented is supported by the European Commission within the FP7 Project STEER (grant no 318343). The authors would like to thank Arjen Veenhuizen, Hans Stokking, Craig Bojko, Jamie Jellicoe, and Ross Wilson who provided insight and expertise that greatly assisted the research.

References

[Khrouf2014] Houda Khrouf, Vuk Milicic and Raphaël Troncy. Mining Events Connections on the Social Web: Real-Time Instance Matching and Data Analysis in EventMedia. In Journal of Web Semantics, 24(1), pages 3-10, 2014

[Seen] Organizing the next generation of content, http://www.seen.co/

[Steiner2014] Thomas Steiner, Telling Breaking News Stories from Wikipedia with Social Multimedia: A Case Study of the 2014 Winter Olympics, In Proceedings of 1st workshop on Social Multimedia and Storytelling (SoMuS 2014). 2014

[Becker2012] Hila Becker, Dan Iter, Mor Naaman, Luis Gravano, "Identifying Content for Planned Event Across Social Media Sites", in Proceedings of the Fifth ACM International Conference on Web Search and Data Mining (WSDM '12).

[Chi2011] Pei-Yu Chi, Henry Lieberman: Raconteur: Integrating authored and real-time social media. ACM CHI Conference on Human Factors in Computing Systems, 2011

[Smits2010] Ewine Smits, Alan Hanjalic, A System Concept for Socially Enriched Access to Soccer Video Collections. IEEE Multimedia 17(4): 26-35. 2010

[Kelliher2014] Aisling Kelliher. Critical Multimedia. IEEE Multimedia, 21(4). 2014

[Xie2014] Lexing Xie, David A. Shamma, Cees Snoek. Content is Dead ... Long Live Content: The New Age of Multimedia-Hard Problems. IEEE Multimedia, 12(1). 2014

[Saini2012] Saini, Mukesh Kumar and Gadde, Raghudeep and Yan, Shuicheng and Ooi, Wei Tsang. MoViMash: Online Mobile Video Mashup. In Proceedings of the 20th ACM International Conference on Multimedia. Nara, Japan. 2012

[Macro2009] Marco A. Figueiredo, Paola Prado, Mauro A. Camara, and Ana M. Albuquerque. Empowering Rural Citizen Journalism via Web 2.0 Technologies, In Proceedings of the fourth international conference on Communities and technologies. 2009

[Guimar2011] Laiola Guimaraes, Rodrigo and Cesar, Pablo and Bulterman, Dick C.A. and Zsombori, Vilmos and Kegel, Ian. Creating personalized memories from social events: community-based support for multi-camera recordings of school concerts, Proceedings of the ACM international conference on Multimedia, 2011

[Dale2011] Michael Dale. Collaborative video editing for Wikipedia. In Proceedings of the 7th International Symposium on Wikis and Open Collaboration (WikiSym '11). ACM, New York, NY, USA, 2011

[WeVideo] Free Online Video Editor & Maker, https://www.wevideo.com/

[Mu2013a] Mu, M., Simpson, S., Bojko, C., Broadbent, M., Brown, J., Mauthe, A., Race, N. & Hutchison, D., Storisphere: from TV watching to community storytelling, In IEEE Communications Magazine. 08/2013

[DASH2012] "Information technology - Dynamic adaptive streaming over HTTP (DASH) -- Part 1: Media presentation description and segment formats", ISO/IEC 23009-1:2012, 2012

Mu Mu is a Senior Research Associate in the School of Computing and Communications at Lancaster University, UK where he received his Ph.D in Computer Science. His research covers the topics of QoE-aware content distribution, SDN applications, media orchestration, creative media, and big data analysis. For more information, visit www.drmu.net or contact him at m.mu@lancaster.ac.uk.

Steven Simpson is a Senior Research Associate in the School of Computing and Communications at Lancaster University, where he received his Ph.D. His research covers many aspects of networking and multimedia, including group communication, content delivery, network programmability and network resilience. Contact him at ss@comp.lancs.ac.uk.

Nicholas Race is a Senior Lecturer within the School of Computing and Communications at Lancaster University. His research interests lie principally within the areas of Wireless Mesh Networks, P2P and Content Distribution Networks, and Software Defined Networks (SDN). Contact him at n.race@lancaster.ac.uk

Omar Aziz Niamut is a senior research scientist at the Netherlands Organization for Applied Scientific Research (TNO). He received his MSc and PhD degrees in Electrical Engineering in 2001 and 2006, respectively, from Delft University of Technology, in the fields of signal processing and audio coding. His technical interests include social TV, media orchestration and immersive media. Contact him at omar.niamut@tno.nl.

Gijs Koot is a researcher at TNO in the Media Mining group. After finishing a MSc. in Stochastics and Financial Mathematics, he is now working on text mining, social physics, human modelling and web science. Contact him at gijs.koot@tno.nl.

Rianne Kaptein has obtained a PhD in Information Retrieval at the University of Amsterdam. In her PhD research she worked on entity ranking, searching in Wikipedia, and generating word/tag clouds. Currently, at TNO she works as a researcher in the Media Mining group in various projects on text mining, search, sentiment analysis and social media analytics. Contact her at rianne.kaptein@tno.nl.

Jacco Taal is Co-Founder and CEO of Bitnomica, based in the Netherlands. His research interests include video compression, video streaming, social network analysis and multimedia analysis and data mining. Taal received a Ph.D in Video Coding from Delft University of Technology. Contact him at jacco@bitnomica.com.