

D6.5 FINAL VERSION OF SECOND SCREEN



Grant Agreement nr	611057
Project acronym	EUMSSI
Start date of project (dur.)	December 1st 2013 (36 months)
Document due Date :	Month 32
Actual date of delivery	21 October 2016
Leader	Eelco Herder (LUH)
Reply to	herder@L3S.de
Document status	Submitted

Project co-funded by ICT-7th Framework Programme from the European Commission



Project ref. no.	611057
Project acronym	EUMSSI
Project full title	Event Understanding through Multimodal Social Stream Interpretation
Document name	D6.5 Final Version of Second Screen
Security (distribution level)	PU – Public
Contractual date of delivery	August 31st, 2016
Actual date of delivery	October 21st, 2016
Deliverable name	D6.5 Final Version of Second Screen
Type	P – Prototype
Status	Final
Version number	1
Number of pages	22
WP /Task responsible	UPF / LUH
Author(s)	Eelco Herder
Other contributors	Mainul Quraishi, Jahid Kabir
EC Project Officer	Mrs. Alina Lupu Alina.LUPU@ec.europa.eu
Abstract	Design and development of the Second-screen demonstrator
Keywords	Second screen, use case, linked data, automatic questions
Circulated to partners	Yes
Peer review completed	Yes
Peer-reviewed by	Jean-Marc Odobez
Coordinator approval	Yes



Table of Contents

1. BACKGROUND	3
2. INTRODUCTION	4
3. FEEDBACK CYCLES AND MAIN IMPROVEMENTS	6
3.1 First Version of the Second Screen	6
3.2 Feedback from the User Day in Bonn (November 2015)	7
3.3 Feedback from User Day in Bonn (July 2016)	7
3.4 Technological Improvements and Extensions	8
4. CURRENT VERSION	9
4.1 Architecture	10
4.2 Generation of Information Boxes and Questions	11
4.2.1. Overview of the generation process	11
4.3 Evaluation and Improvement of Question Generation	13
4.4 Editing Second-Screen Content	14
5. FINAL USER EVALUATION IN BARCELONA (September 2016)	18
6. CONCLUSIONS AND FUTURE DIRECTIONS	20
7. REFERENCES	21



1. BACKGROUND

The goal of WP6 is to explore multimodal content-based recommendation techniques and to use them in two use-cases: a contextualising tool for news editors and a second-screen application for end-users. These use cases have been elaborated in WP2 and WP7, leading to a clear concept for the second-screen use-case. In D6.3, we explained how this concept has been translated into a practical and showcase the first implementation.

In this deliverable, we report the improvements, extensions and evaluations carried out in the past year, based on internal evaluation and external feedback from Deutsche Welle and the participants of the EUMSSI User Days.



2. INTRODUCTION

While watching television or online broadcasts, users often interact with other applications and websites on the same device or on a different device, such as a laptop, tablet or smartphone. In addition to unrelated activities like chatting or consuming social media streams, there are also plenty of user activities that are related to the current program. This includes social interaction (such as discussing a live soccer match), searching for biographies of actors and reading background information or definitions on Wikipedia.

Several media providers aim to support or manage these program-related online activities by providing so-called 'second-screen content' themselves, varying from live content during live television broadcasts (such as social media feeds) to program recommendations for streaming services (as provided by the BBC iPlayer and educational or informative content for online video media, such as available at YouTube and Vimeo, or in archives, such as the TED Talks¹.

In EUMSSI, we provide second screen content that falls in the category *edutainment*, which is defined by Wikipedia² as 'games explicitly designed with educational purposes, or which have incidental or secondary educational value'.

The second screen provides more background information without distracting the user from the program broadcast on the first, primary screen. Otherwise, it becomes too overwhelming for the user and the user loses track, or we simply lose the user. The second screen is targeted at the end user, enriching the user experience with content that is related to the video-on-demand that s/he is currently watching.

The second screen application provides both 'passive' content that provides additional information and 'active' content that invites the user to interact with it. The most basic form of additional information is the provision of text-based background information on politicians, locations, historical figures, events and locations that are featured in a news broadcast or other programs. In order to provide such content, we use the entities extracted and linked to DBpedia and enrich them using DBpedia links and relations, as

¹ <https://www.ted.com/talks>

² https://en.wikipedia.org/wiki/Educational_entertainment



well as other linked databases, such as Geonames³. The intuition is that each country has its currency, population and capital (among other things).

More formally, EUMSSI exploits the concept of Linked Data (Heath 2011). The foundation of Linked Data is that the Web is not only used for storing documents, but also for data. The data objects on the Web are identified by Web addresses (URIs) similar to the way documents are. The representation of the data - i.e. the information associated with a data object - is then represented by a Web link, which can itself be characterized by a Web address. This makes it possible to represent information in such a way that it is materialized as a graph, where nodes are Web addresses and the edges are Web links.

For each video, a large variety of metadata - as extracted by the EUMSSI framework - is available. These entities, to which a timestamp has been assigned based on the contents of the video transcript, form the basis of the content generation process. This process involves 'translating' the so-called subject-predicate-object triples of DBpedia into human-readable sentences or questions. For example "*Germany dbo:capital dbr:Berlin*" may be translated into the informative sentence "*The capital of Germany is Berlin*" or lead to the question "*What is the capital of Germany*". This approach works for the majority of cases, but depends on how clean the entries in DBpedia are.

In deliverable D6.3, we provided a motivation and design considerations for the second screen concepts that we target in EUMSSI. In that deliverable, we also presented the first preliminary version of the demonstrator. The second screen has been iteratively evaluated and improved, making use of qualitative and quantitative feedback obtained on two User Days.

The most important change is the implementation of a server-side architecture that allows one or more second screen clients (laptop, tablet or smartphone) to connect to a video running on another device. All clients are automatically synchronized with the running video. We also significantly enhanced the pool of generated content (information and pub-quiz like questions) and created an editor for editing the automatically generated content (to correct or optimize them, a feature requested by Deutsche Welle). Finally, the user interface has significantly improved.

In the next chapter, we summarize the earlier version of the second screen, as presented in D6.3, followed by a discussion of the feedback obtained on the three user days and the technological improvements that we carried out. We continue with a

³ <http://www.geonames.org/>



description of the current version of the second screen, including an evaluation of the correctness and interestingness of the automatically generated content. We conclude the deliverable with results and feedback from the third EUMSSI User Day, held in September 2016 in Barcelona, followed by conclusions and future perspectives.

3. FEEDBACK CYCLES AND MAIN IMPROVEMENTS

3.1 First Version of the Second Screen

Our first demonstrator of the second screen was implemented as a plugin of the Amalia video player⁴. Amalia.js is an extensible HTML5 multimedia player that allows users to view any type of metadata along with video or audio streams.

The video player loads an MP4 video file, of which the URL is given by the corresponding entry in Solr. Together with the actual video, a large variety of metadata is provided as well. In short, in Amalia, the video is loaded along with its corresponding entities. These entities are matched with entries in DBpedia in order to obtain additional information about these entities. This additional information is used for providing general information or trivia and for generating questions.

A screenshot of the first second screen prototype is shown in Figure 1. In D6.3, more technical details can be found. In this deliverable, we focus on the improvements made and the final, current version of the second screen.



Second Screen Demo

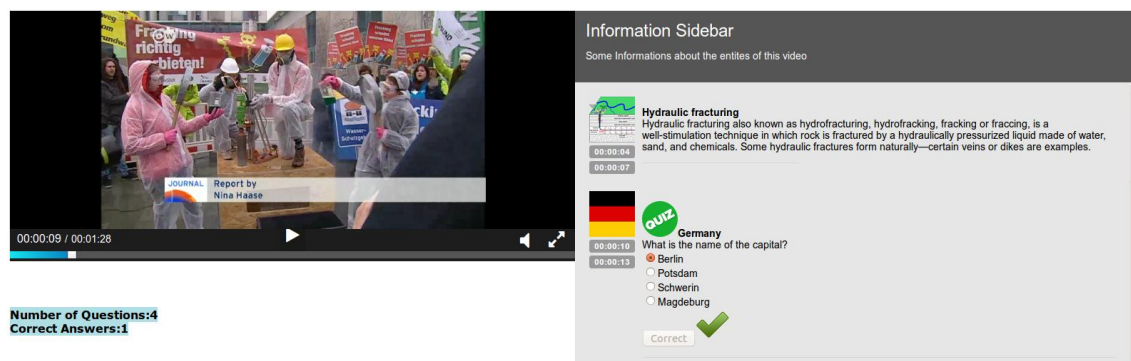


Figure 1: First second screen prototype

⁴ <https://ina-foss.github.io/amalia.js/>
EUMSSI D6.5 - Final Version of Second Screen



3.2 Feedback from the User Day in Bonn (November 2015)

The preliminary version of the First Screen was first presented during the User Day in Bonn as a novel, experimental application that exploits the building blocks provided by the EUMSSI Toolbox for *edutainment* applications targeted at the end users.

At the time of the User Day, the second screen application was work-in-progress and presented to the user as a sidebar of the Amalia player (see screenshot above). The informal presentation at the user day was aimed to obtain first feedback and inspiration.

Colleagues from Deutsche Welle appreciated the ability to generate information and interactive application to offer to the viewers without much additional effort. They also thought it would be hard to find interesting formats that would engage the end-user.

Harald Sack, member of the advisory board and one of the developers of WhoKnows (Waitelonis 2011), was interested in the question generation process. He offered several interesting points of research to pursue, such as using the quiz for detecting errors in DBpedia. He also proposed to use popularity and centrality measures for estimating the difficulty of questions (with the intuition that people probably know more about a popular/important person such as Barack Obama than about, say, the mayor of Hannover, Germany). These importance measures have not been exploited, but they are now being collected in the context of the metadata enrichment in WP5.

The users did not comment much upon the user interface, but it became clear that in order to be attractive for end-users, the second screen needed to work on a separate device. For this reason, it was decided to invest time and effort in the development of a server-side solution for synchronizing video and second screen content on different devices - as will be explained in the next chapter.

3.3 Feedback from User Day in Bonn (July 2016)

During a joined User Day with the EU project SAM in Bonn in July 2016, the second screen prototype was presented and evaluated. This was already the version with two or more synchronized devices and the iteratively improved user interface, as will be introduced in the next chapter. The majority of the users and testers found the prototype promising. Six user provided feedback with a questionnaire.

The feedback provided was used to optimise the second screen tool in many respects. One of the most important comments addressed the need to develop an editing layer.
EUMSSI D6.5 - Final Version of Second Screen



Authors would then be able to modify and edit the automatically derived items. This was considered important for two reasons. First and foremost, as explained in more detail in D6.3, the automatically assembled content may contain errors due to incorrect, inconsistent or outdated information in DBpedia. Second, even though the content may be technically *correct*, some pieces of knowledge or questions may be less relevant or interesting to the user - a project manager at Deutsche Welle specifically indicated that he would not want automatically generated content to be provided to the end-users without human approval or curation.

This specific feedback has led to the decision that infoboxes and questions would not be generated on-the-fly anymore, but automatically generated and stored just once - when first watching the video. An editor would then be needed to select which information or questions to provide for which entity and to allow for editing the phrasing or the choice of distractors (the wrong answers in multiple choice questions).

3.4 Technological Improvements and Extensions

To summarize, the main improvements that we implemented, based on our own experiences, discussions with stakeholders and feedback obtained during the User Days are as follows:

- We abandoned the Amalia video player with the information sidebar and, instead, implemented a server-side solution that allows a video to be connected with one or more second screen clients on different devices.
- We developed a communication protocol for synchronizing the videos with the second screen content.
- Instead of randomly generated content on the fly, we now create all possible second screen content upon the first time that a video is watched. An editor has been developed that allows for changing the type of content to be shown and for editing information boxes and questions.
- Various iterations on improving the user interaction (navigation infrastructure, labels, headers, look&feel) have been carried out.

4. CURRENT VERSION

The EUMSSI second screen application is implemented as a server-side application that connects a 'first screen' client, which shows the video in an HTML5 player, with one or more 'second screen' clients (which may be on the same machine, or on a separate laptop, tablet or smartphone). All possible information and questions that can be shown to the user, are stored in JSON format in a WebVTT file (WebVTT is a W3C standard for displaying timed text, such as subtitles, in connection with the HTML5 video⁵).

All second screen clients that are logged in using the same identifier as the video show the same content at the same time. The video client sends the relevant content to the server at the moment specified in the VTT file. The server-side agent forwards the content to the second screen clients. Figure 2 shows a typical usage scenario for the second screen.

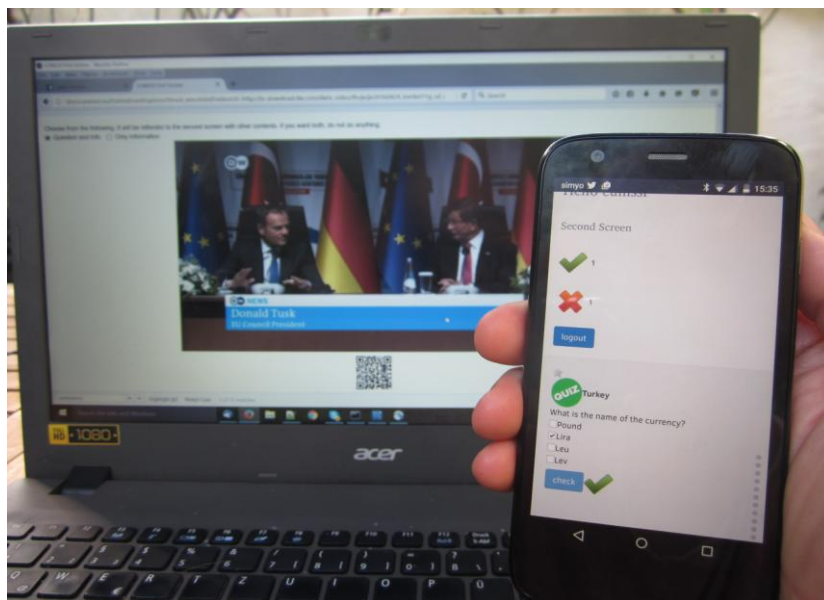


Figure 2: The second screen application running on a smartphone, showing a question related to the video running on the laptop.

⁵ <https://w3c.github.io/webvtt/>

4.1 Architecture

WebVTT⁶ (Video Track Text) is a W3C standard that is mainly used for connecting video with external text tracks, particularly for creating (multi-language) subtitles to be shown on the video. WebVTT is a simple text file with an ordered list of timestamps and content (like subtitles) to be shown when the video reaches this timestamp. The following example (taken from the W3C website) contains subtitles to be shown between the first and fourth second of a video:

```
00:01.000 --> 00:04.000  
Never drink liquid nitrogen.
```

A naive strategy would be to send the WebVTT file to the second screen clients, which show the second screen content with the given time intervals. However, this obviously does not work, as videos can be paused, users can move forward or backward in the video, or choose another video. For this reason, synchronisation between the video and the second screen client is needed.

Note that communication on the Web is stateless: once a connection is closed, all information on this connection is lost. There are two main techniques for synchronisation on the Web: pushing and polling⁷. *Pushing* means that the Web server leaves a connection to the client open, to immediately send information to the client once it becomes available. A *polling* technique would mean that the server waits for the clients to request new information. For pushing, it would be necessary that the server knows which (second screen) clients are connected to the video. For polling, all clients would need to request whether there is new information available on a regular basis. A variation on polling, which basically emulates push, is *long polling*: the client sends a request to the server and then waits until the server responds. This technique is popular among chat clients. We adopted a technique similar to long polls, as illustrated in Figure 3:

- Video clients and second screen register with an identical (self-chosen) identifier at the server (red arrows).
- The server stores lists of all registered videos and associated clients.
- Once a video passes a point with an associated entity (as specified in the client-side WebVTT file), this content associated with this entity is sent to the server. The server forwards the content to the second-screen clients.
- The clients display the associated content and wait for the server to send the next pieces of content.

⁶ <https://www.w3.org/TR/webvtt1/>

⁷ https://en.wikipedia.org/wiki/Push_technology

This way, we avoid unnecessary server load or network communication and still allow for synchronization without delays.

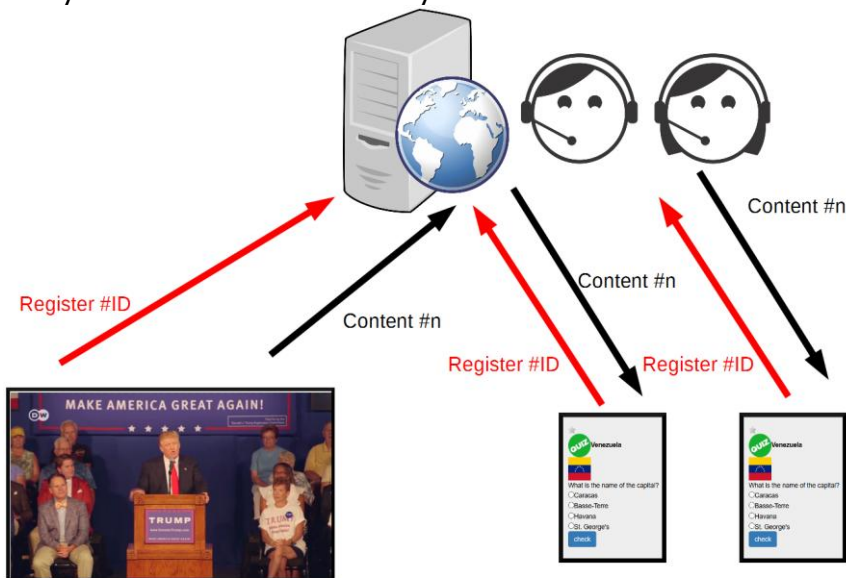


Figure 3: Synchronization and communication between video and second screen clients.

4.2 Generation of Information Boxes and Questions

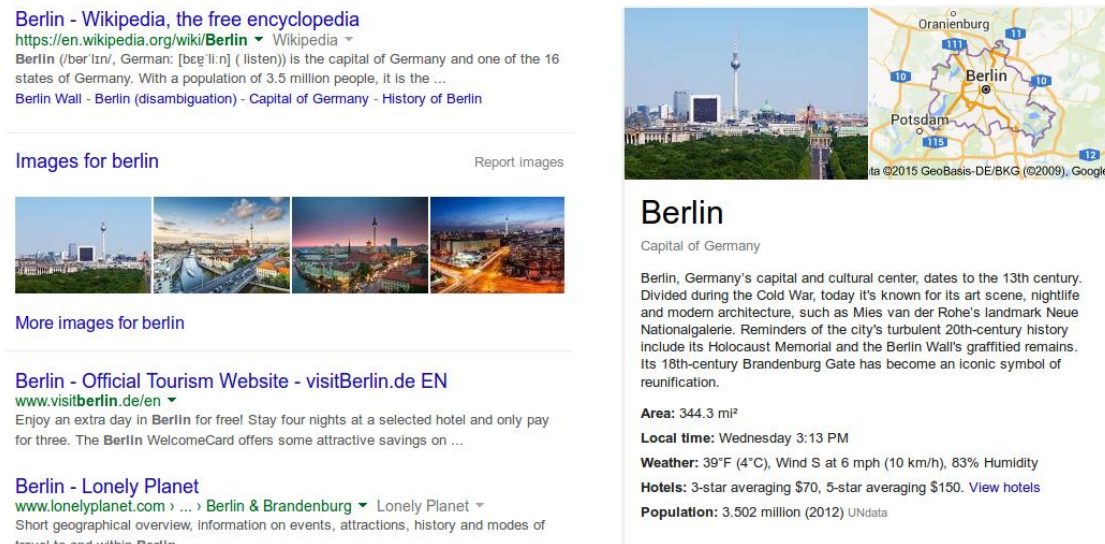
The generation of information boxes and question has not substantially changed from the process explained in D6.3, on the first version of the second screen. However, the set of questions has been substantially enlarged and the question generation process has been evaluated and improved.

Generating multiple choice questions does not only involve translating a pattern in human-readable content, but also finding distractors (wrong answers that are close enough to be considered as true). There is a small number of previous work on automatically creating questions (Seyler 2015, Waitelonis 2011) and our approach is in line with these works.

4.2.1. Overview of the generation process

For the sake of completeness, we summarize the process for generating second screen content. For more technical details, please refer to chapter 6 of D6.3.

Virtually all DBpedia entries contain an *abstract*, a short description that is also used by search engines, such as Google (see Figure 4), for displaying an information box next to the search results.



Berlin - Wikipedia, the free encyclopedia
<https://en.wikipedia.org/wiki/Berlin> - Wikipedia
 Berlin (/berˈlɪn/, German: [bɛʁˈliːn] (listen)) is the capital of Germany and one of the 16 states of Germany. With a population of 3.5 million people, it is the ...
[Berlin Wall - Berlin \(disambiguation\) - Capital of Germany - History of Berlin](#)

Images for berlin Report images

More images for berlin

Berlin - Official Tourism Website - visitBerlin.de EN
www.visitberlin.de/en
 Enjoy an extra day in Berlin for free! Stay four nights at a selected hotel and only pay for three. The Berlin WelcomeCard offers some attractive savings on ...

Berlin - Lonely Planet
www.lonelyplanet.com > ... > [Berlin & Brandenburg](#) - Lonely Planet
 Short geographical overview, information on events, attractions, history and modes of travel to and within Berlin.

Berlin
 Capital of Germany

Berlin, Germany's capital and cultural center, dates to the 13th century. Divided during the Cold War, today it's known for its art scene, nightlife and modern architecture, such as Mies van der Rohe's landmark Neue Nationalgalerie. Reminders of the city's turbulent 20th-century history include its Holocaust Memorial and the Berlin Wall's graffitied remains. Its 18th-century Brandenburg Gate has become an iconic symbol of reunification.

Area: 344.3 mi²
Local time: Wednesday 3:13 PM
Weather: 39°F (4°C), Wind S at 6 mph (10 km/h), 83% Humidity
Hotels: 3-star averaging \$70, 5-star averaging \$150. [View hotels](#)
Population: 3.502 million (2012) UNdata

Figure 4: information box in Google search results

In addition to showing the abstract, information boxes are generated that present concise facts or trivia. It does so by 'translating' the subject-predicate-object triples of DBpedia into human-readable sentences. For example:

Germany dbp:capital dbr:Berlin

becomes

The capital of Germany is Berlin

This approach works for the majority of cases, but depends on how clean the entries in DBpedia are. Unfortunately, despite continuous efforts from the Wikipedia and DBpedia



community and despite the availability of accepted metadata schemes, inconsistencies or partially redundant or missing information can be found in several places.

The generation of questions is similar to the generation of entity information, with the difference that false answers need to be generated on the fly. For performance reasons, we locally store information on all countries and major cities (about 3300 in total) in the world.

We illustrate the approach with an example: for the question "*which city is the capital of Germany*", we can either search for capitals of neighbouring countries or neighbouring cities in the same country – in both cases by finding the top- n cities or countries with the smallest distance (in terms of longitude and latitude) from the current country or city - for example Schwerin, Potsdam and Magdeburg. For German readers - or European readers in general - these are obviously false answers, but many Europeans would probably not know immediately whether the capital of Canada is Toronto, Ottawa, Montreal or Vancouver. In a similar manner, we generate incorrect answers on questions regarding the national currency or the official language of a country. For generating incorrect answers for persons, we rely on a database of about 2500 persons that appear in the collection of annotated DW videos.

4.3 Evaluation and Improvement of Question Generation

During the second User Day, we received the remark that it would be desirable to edit and improve the automatically generated content, in order to correct erroneous content or to improve the relevance or interestingness of the content for the user. As our initial goal was to generate second screen content completely automatically, without any human effort or organization resources needed (which, according to media partner Deutsche Welle, was also very desirable), we decided to first investigate to what extent the automatically generated questions needed improvement.

In a preliminary study, we generated 15 instances of each of the possible questions for locations and people, created for a manually selected set of entities that covered a broad range of persons, countries and cities (e.g. Sydney, Ankara, Mumbai, Manchester, Berlin, ...). These questions were evaluated on correctness and interestingness by two student assistants. We found that 77% of the questions were considered correct and 79% of the questions were considered interesting. The reason behind the higher percentage of interestingness than correctness is that not all the questions were generated properly, even though the evaluators appreciated the question as users. However, there is a high correlation between correctness and interestingness (0.89, $p < .05$). An example of a technically correct but uninteresting

EUMSSI D6.5 - Final Version of Second Screen



question would be 'In which country is Sydney located?' with the answer options 'USA, Nepal, Sri Lanka, Bangladesh'. First of all, users will probably think that the correct answer fails (apart from the famous city in Australia, there is also a Sydney in the USA). Second, the choice of distractors (wrong answers) is correct but suboptimal.

We also investigated the correctness and interestingness of the generated distractors (wrong answers), which were rated 70% and 65% respectively. Incorrect distractors were in almost all cases caused by errors in DBpedia. The most common causes are *deprecated* information (such as 'Weimar Republic' for the country that is currently known as Germany), *name conflicts* (the city of Sydney in Canada is used instead of the more famous city in Australia with the same name), *ambiguous answers*, *typo problems* (particularly related to underscores) and *empty fields*. In some cases, solutions to such problems can be found, but due to inconsistencies in DBpedia, these solutions also do not work in all cases. For examples, the 'Weimar Republic' has a predicate *dissolutionYear*, but that does not yield for all former countries (for example, Yugoslavia does not have this predicate).

In summary, the process of automatically generating information and questions for the second screen generally leads to fairly correct and interesting content. However, particularly for the distractors, post-hoc verification would be desirable. In order to facilitate this process, we developed an editor that allows for manual adaptation of generated infoboxes as well as questions and distractors.

4.4 Editing Second-Screen Content

In order to allow humans to choose between different choices of second-screen content and to edit the content itself, we developed a simple form-based editor. Using this editor, the content of the WebVTT file (which contains the content associated with the video), as stored on the server, is updated.

As explained earlier, the first version of the second screen randomly generated second-screen content on-the-fly while a video was displayed. In order to allow human editing, we changed the strategy into a one-time only generation of the second-screen content, with all variations for each entity stored in the WebVTT file. This information is stored in JSON format. The snippet below shows the content for the entity `United_States`, to be displayed between 2:03 and 2:06, which includes:

- A thumbnail (from Wikipedia) to be shown along with the entity
- A list of questions (on capital, currency, language, neighboring countries) with the correct answer and the distractors (such as San Francisco, Miami and New York for the capital question)

EUMSSI D6.5 - Final Version of Second Screen



- An abstract, as obtained from DBpedia
- A Google Map with the associated entity as a query
- Infoboxes
- A field that indicates which of the above is currently selected to be displayed (in this case 'question number 0', which is the 'capital question'.

United_States

00:02:03.000 --> 00:02:06.000

```
{ "thumbnail": "http://en.wikipedia.org/wiki/Special:FilePath/Flag_of_the_United_States.svg?width=300", "name": "United_States", "questions": [ { "question": "What is the name of the capital?", "correct": "Washington", "options": [ "San Francisco", "Washington", "Miami", "New York" ] }, { "question": "What is the name of the currency?", "correct": "Dollar", "options": [ "Dollar", "United States Dollar", "Gourde", "Peso" ] }, { "question": "What is the official language spoken here?", "correct": "English", "options": [ "Dominican Spanish", "French", "English", "Cuban Spanish" ] }, { "question": "Which countries are the neighbours?", "correct": "Canada, Mexico, Cuba", "options": [ "Canada", "Mexico", "Connecticut", "Cuba" ] } ], "abstract": "The United States of America , commonly referred to as the United States , America, and sometimes the States, is a federal republic consisting of 50 states, 16 territories, and a federal district. The 48 contiguous states and Washington, D.C., are in central North America between Canada and Mexico. ", "map": "<iframe width='300' height='300' frameborder='0' style='border:0' src='https://www.google.com/maps/embed/v1/place?q=United_States&key=AIzaSyAJCit6wJoDOEdy-PPb0L5_3SydfJbJ3Vg' allowfullscreen></iframe>", "infos": [ "The capital is Washington", "The local currency is Dollar", "The language spoken is English", "language spoken English, Spanish, Hawaiian, French" ], "default_content": { "number": "0", "type": "questions" } }
```

The editor provides a more human-readable representation of the WebVTT files and a simple way to adapt the second-screen content. On the overview screen (Figure 5), the human editor gets an overview of the currently selected content and has the opportunity to change the type of content to be shown or to edit the selected content (buttons on the right-hand side).

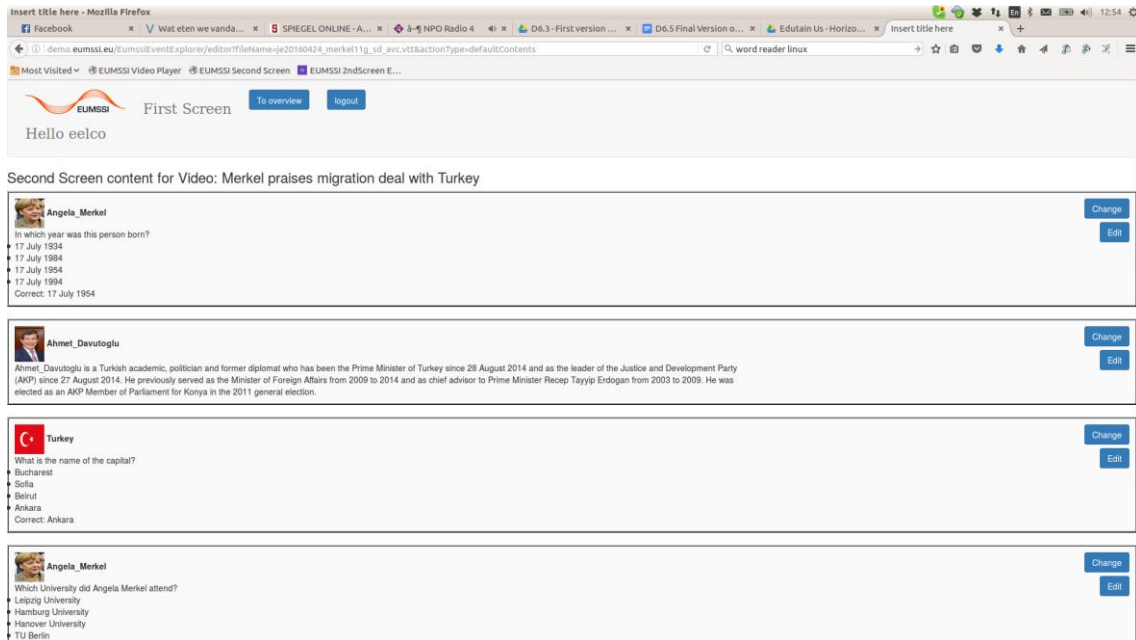


Figure 5: Overview of current second-screen content for a video

When the user presses the 'Change' button, the currently selected question can be changed (as displayed in Figure 6).

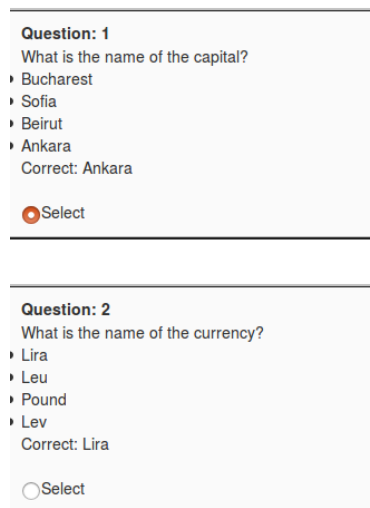


Figure 6: Choosing between different types of content

EUMSSI D6.5 - Final Version of Second Screen

When the user presses the 'Edit' button, a page appears on which the content is displayed as an editable Web form, in which the fields can be edited. Upon pressing 'Submit Query', the associated content of the WebVTT file is updated.

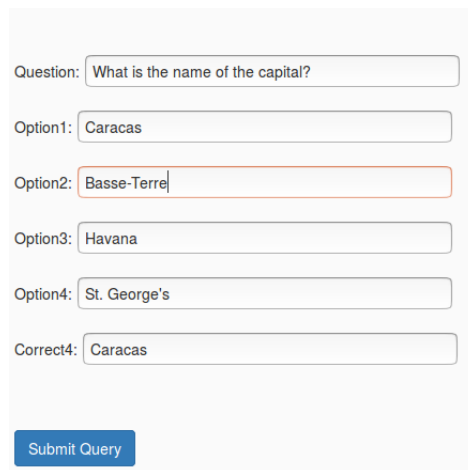
A screenshot of a web form for editing a question. The form consists of several input fields and a submit button. The fields are labeled as follows: "Question:" with the text "What is the name of the capital?"; "Option1:" with the text "Caracas"; "Option2:" with the text "Basse-Terre"; "Option3:" with the text "Havana"; "Option4:" with the text "St. George's"; and "Correct4:" with the text "Caracas". A blue button labeled "Submit Query" is located at the bottom of the form.

Figure 7: Editing a question

In this way, we can assure that the second-screen content is accurate and relevant, while at the same time still exploiting the time-saving method of automatically generated questions. As discussed in the previous section, about 70% of all generated content was considered correct and interesting.

5. FINAL USER EVALUATION IN BARCELONA (September 2016)

During the third User Day, held in Barcelona and visited mainly by people with an entrepreneurial background in different professions (more details to come in the associated WP7 deliverable), the second screen was tested by a good number of users, mainly in couples, of which ten filled out the associated questionnaire. In this chapter, a number of relevant observations will be discussed. As can be seen in Figure 8, all users agreed that the content of the second screen was interesting or inspiring (all rated this with 3 stars or more). The positive difference with the previous User Day can be attributed to the editor, which we used for manually selecting and optimizing the content. The results also show that there is still room for improvement - particularly the variety of questions could be extended with content for other entities than persons and locations (such as companies, events, observances).

I found the content (quiz / explained terms) shown in the 2nd screen interesting / inspiring.

(10 responses)

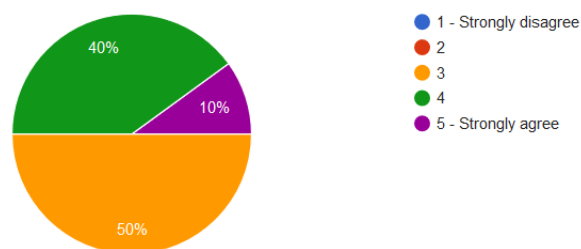


Figure 8: Rating of the content of the second screen.

The synchronization of the second screen with the video content received very good ratings, as shown in Figure 9. Some users found it 'magical' that the content was displayed on their own smartphone.

The 2nd screen content was well synchronized with the primary content. (10 responses)

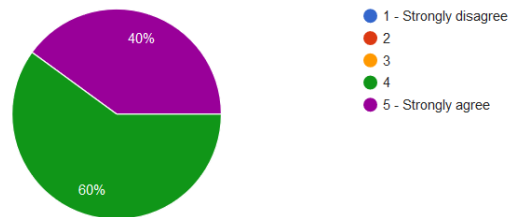


Figure 9: Synchronization of content with the video

From the comments and discussions it became clear that synchronization is of importance. Users tolerate it when content regarding an entity appears some seconds later (for example, in order to leave sufficient time for the content related to the previous entity), but are surprised when the content appears before an entity is mentioned (in one occasion, we showed a question on Angela Merkel only shortly before she appeared).

Some users found that there was too much second-screen content and that the intervals were too short. We deliberately showed content with 10-20 second intervals, for demo purposes (where one only has a couple of minutes to showcase a demo). In real settings, intervals of several minutes would probably be more appropriate.

We received a good number of other suggestions. For example, one smartphone user suggested that a vibration (haptic feedback) once new content appears would be useful. Another smartphone user reported that he already had this haptic feedback. It turned out that this is browser-dependent and that some browsers automatically support this. In addition, some users suggested to pause the video when users interact with the second screen (either automatically or by the users themselves). Overall, the second screen was much appreciated and particularly the concept was rated as very promising - see Figure 10.

The use of EUMSSI was overall a promising experience. (10 responses)

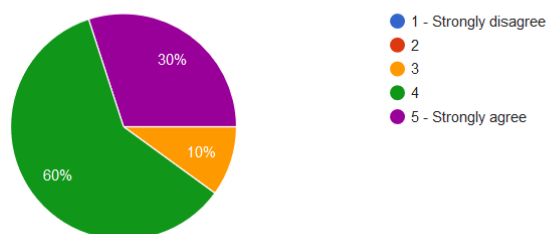


Figure 10: User ratings on the experience with the second screen.



6. CONCLUSIONS AND FUTURE DIRECTIONS

In the past year, the second screen has evolved from a simple prototype to a mature demonstrator. The current version synchronizes the content between the video client and one or more second screen clients, which may be smartphones, tablets or laptops. The content generation process has been improved and evaluated. In order to allow for manual corrections and improvement of the selected content, a simple and easy-to-use editor has been developed.

The second screen has been presented to the users during the third User Day in Barcelona and received positive response. The concept was rated as highly promising and the second screen content as interesting.

Despite the positive response, we are aware that the second screen demonstrator, as developed in EUMSSI, still would need extension and improvement in order to be fully mature. This falls outside of the scope and effort assigned to this topic within the overall project. Nevertheless, we think it is well worth to discuss two concrete future directions to explore.

In EUMSSI, most second screen content is related to persons and locations. As explained in D6.3, the main reason is that for these types of entities the information provided by DBpedia is most complete and correct (and even for this type of content, errors can be found - as we saw in this deliverable). Moving on from this 'low hanging fruit' to less structured and more varied entities (such as holidays, observances and companies) requires significantly more effort and experimentation. For example, for companies, we could generate question on their founding year, their director, home country or city or product ranges. However, this depends on the availability of this information in structured form. An informal investigation on tech companies (Apple, Microsoft, Facebook, Twitter, Google, Yahoo, ...) revealed that this would require a lot of experimentation and brute-force queries and data collection.

A particularly promising, but also labor-intensive, direction would be to provide users with polls, opinions and links to related articles and discussions in the social media. This would allow users to be informed about possible bias in a video and to investigate other voices and opinions, in order to break out of their so-called *filter bubble*⁸.

⁸ https://www.ted.com/talks/eli_pariser_beware_online_filter_bubbles?language=en
EUMSSI D6.5 - Final Version of Second Screen



7. REFERENCES

Tom Heath and Christian Bizer (2011) *Linked Data: Evolving the Web into a Global Data Space* (1st edition). *Synthesis Lectures on the Semantic Web: Theory and Technology*, 1:1, 1-136. Morgan & Claypool.

Seyler, D., Yahya, M., & Berberich, K. (2015). Generating Quiz Questions from Knowledge Graphs. In *Proceedings of the 24th International Conference on World Wide Web* (pp. 113-114). ACM.

Waitelonis, J., Ludwig, N., Knuth, M., & Sack, H. (2011). Whoknows? evaluating linked data heuristics with a quiz that cleans up dbpedia. *Interactive Technology and Smart Education*, 8(4), 236-248.