

Crowdsourcing for Film-induced Tourism: An Approach to Geolocation

Daniel Zilio, Andrea Micheletti, and Nicola Orio

Department of Cultural Heritage – University of Padua
Piazza Capitaniato, 7 – 35139 Padua – Italy
{daniel.zilio, andrea.micheletti, nicola.orio}@unipd.it

Abstract. This paper presents an initial proposal for the development of a crowdsourcing platform aimed at gathering information about audiovisual productions. Our goal is to foster the task of identifying movie excerpts that are filmed in a given location and of manually geolocating these excerpts. In this way it is possible to link audiovisual content to a geographical area. The information is used to populate a database system that is the core component for designing novel *film-induced tourism* offers. The task of identifying the locations of movies, TV and web series, spots and music videos is carried out by a community of non experts, using a crowdsourcing approach where participation and engagement are promoted by a *reputation* mechanism. The paper presents a multilevel approach to geolocation and discusses its motivations and initial results.

Keywords: crowdsourcing, film-induced tourism, reputation, geolocation

1 Introduction

Film-induced tourism regards the interest towards destinations motivated by a connection with audiovisual content. Film-induced tourism has been defined by Sue Beeton as “visitation to sites where movies and TV programmes have been filmed as well as to tour to production studios, including film-related theme parks” [2]. It may be very helpful for destination management and destination marketing: it attracts new visitors and also tourists who have already seen the area for a different reason; it is largely independent of seasonal trends; it might convey tourists from overcrowded sites to new and less explored ones; and it can be suitable for a substantial re-branding of a certain area. A number of film-induced tourism related initiatives has been already undertaken at the international level by public and private bodies, which developed movie maps and movie tours or exploited the success of a particular movie as a tool for destination branding [5].

Visit Britain, the tourism office of the UK, was one of the first agencies to use movies for attracting tourists: since the nineties it designed movie tours and developed tools such as movie maps, that is maps representing locations of movies and TV series. Another interesting case is related to *The Lord of the*

Rings saga (Jackson, 2001-2003), which is still attracting masses of tourists in New Zealand, thanks to a thoughtful policy by the local government. Also Italy has a number of relevant cases, including *The Passion of the Christ* (Gibson, 2004) that doubled the number of tourists visiting Matera in Basilicata and a number of fictions such as *Inspector Montalbano* in Sicily, *Carabinieri* and *Don Matteo* in Umbria.

Tourism effectiveness can be improved through the access to a rich collection of structured data, which is the starting point for designing an exhaustive and attractive tourism offer and for developing ICT tourism tools. Unfortunately, the creation of a structured collection of multimedia items is a time-consuming task, which may require a huge quantity of individuals involved in a boring and unsatisfying activity. In this paper, we address this problem using crowdsourcing and reputation principles. Since these tools showed their effectiveness in a variety of fields, we believe that they can be a viable solution to the task of geolocating a large number of audiovisual excerpts.

2 Motivation

Our interest towards film-induced tourism applications dates back to 2014 [8]. The goal was the creation of a database to represent the relationship between movie excerpts and locations. The project had a double goal: promoting film-induced tourism and raising interest towards audiovisual productions, their language and their relation with the landscape. Although the focus was on locations in Padua and its Province, the in-depth description even for a relatively small geographical area showed to be a time-consuming task. One critical aspect was that a restricted number of users – e.g. scholars in film studies – were involved in providing all the information, from film analysis to location identification. Yet, only part of this information required the competences of trained experts while most of the time was spent in the less specialized task of annotating multimedia items with geographical information. These considerations suggested a crowdsourcing approach.

The term *crowdsourcing* was coined in 2006 [7] as a variant of the term (and the practice of) outsourcing: an institution asks that a given function is undertaken by an undefined, and possibly large, group of individuals. The term is applied to situations where individuals receive a payment for their work, but it is also used when participation is voluntary. Amazon Mechanical Turk¹ and CrowdFlower² are well-known crowdsourcing platforms, which have been exploited also for multimedia research. Relevant research projects that are based on crowdsourcing include the segmentation of ancient Maya hieroglyph-blocks [4], the assessment of Internet video quality [6] and the teaching in cinematography field [1].

In order to encourage participation to a crowdsourcing activity, we envisage an interaction with a simple and intuitive interface, which implements a reputa-

¹ <http://www.mturk.com>

² <http://crowdfower.com>

tion mechanism. The latter is likely to be part of a more complex process that is an investment in social capital. In this contest, “social capital is the actual or potential resources which are linked to a durable network of institutionalized relationships of mutual acquaintance or recognition” [3]. The quest for gaining a social capital requires a continuous search for consensus [9] and sharing with the group what we know is becoming a pleasure without the expectation of reciprocity.

This identification with a group, based on sharing just for the pleasure of it with no hidden agenda, is one major aspect of new digital media also because digital objects (i.e. recordings in a music digital library) can be replicated and shared an infinite number of times without losing quality. Moreover, sharing with the group both what we have and what we know is becoming a pleasure without the expectation of reciprocity, an openness based on selfless approach. This sense of belonging, which is shared also by on-line communities that have been called *emotion communities*, is balanced by an opposite drive: consumption and participation as a narcissistic act of exhibitionism in order to be socially recognized [?].

2.1 Crowdsourcing for Film-induced Tourism

Our goal is to develop a web portal where registered users can contribute to data gathering in all the forms that can be useful for film-induced tourism applications. The main functions that will be available are the following.

- Discover, and insert into the database system, metadata related to multimedia content – from now on called *MOVIE* – that is relevant for film-induced tourism: movies, TV and web series, advertisements, documentaries, music videos.
- Identify the excerpts of a given *MOVIE* that are relevant for a given geographical area or a particular location. We define these excerpts *MOI*, from Moment Of Interest.
- Geolocate the *MOIs*, that is associate to the excerpts one or more geographical locations; geolocation can have different resolutions, from a wide area to a precise place with given spatial coordinates. We define these locations *SOI*, from Site of Interest as an extension of the well-known concept of Point Of Interest (POI).
- Vote the activities carried out by the other users, in particular geolocation, in order to create a community of users based on the concept of reputation.

The next section focuses on the geolocation task, providing a first case study.

3 Geolocation Model

We define a *MOI* m as a portion of a *MOVIE* defined by the time interval $[t_1, t_2]$ with $t_1 < t_2$. Considering the set M of the *MOIs* of a *MOVIE*, geolocation

consists in the association of geographical information to each $m \in M$. Given U the set of users, the outcome of geolocation can be described by

$$\bar{v}_{u,m}^l = \langle \bar{s}_{u,m}^l, \bar{g}_{u,m}^l \rangle \quad (1)$$

which represents the operation carried out by user $u \in U$ for MOI $m \in M$; l is the index of the SOI. Although it is possible that two different geographical areas are relevant for the same movie excerpt, for instance because they are the result of film editing, for simplicity of notation we assume that a single SOI is associated to a MOI and hence we omit the index l in the following discussion. The outcome of geolocation can thus be represented in two, not exclusive, forms:

- $\bar{s}_{u,m}$ represents the symbolic description of a geographical area, such as the name of a region, a province, a city, and so on.
- $\bar{g}_{u,m}$ represents the geographical coordinates, usually in the form of latitude and longitude. Symbolic information is strictly required, at least defining a wide area such a country or a region, geographical information is not mandatory.

Let us define V as the set of all insertions for a given SOI, and $\bar{v}_{o,m}$ the ground truth. In this case, we assume that the ground truth can be obtained by an expert in film-induced tourism who has a perfect knowledge about both the MOVIE and the SOIs. In our context, the ground truth represents the optimal position for a tourist to fully appreciate the connection between the audiovisual content and the physical location. We assume that the ground truth is described with the maximum available accuracy for associating a SOI to a MOI. The goal is to find $\bar{v}^* \in V$ that minimizes the distance with the ground truth, providing that each user of the system can provide an incomplete, imprecise or even incorrect estimate of a SOI. Thus, a distance function needs to be defined between \bar{v} .

The choice of the distance function depends on how \bar{v} is represented. The computation of a distance between two SOIs that include geographical information \bar{g} is straightforward but, in a crowdsourcing environment, the requirement of providing precise geographical coordinates may discourage users, who may not be able to geolocate MOI at this level of detail. Yet, users may contribute some relevant information, because they are aware of the geographical area, the city, or even the neighbors of a MOI.

Distance between Symbolic Descriptions In order to overcome this problem, we propose to use the symbolic representation \bar{g} of a SOI. Formally, the insertion made by user u about MOI m is defined as

$$\bar{s}_{u,m} = \langle \alpha_{u,m}^0, \alpha_{u,m}^1, \dots, \alpha_{u,m}^V \rangle \quad (2)$$

while the ground truth for MOI m is defined as

$$\bar{s}_{o,m} = \langle \alpha_{o,m}^0, \alpha_{o,m}^1, \dots, \alpha_{o,m}^N \rangle \quad (3)$$

where elements have an increasing level of accuracy, and $V \leq N$ that is the information provided by a user is at most as accurate as the ground truth.

A practical example may help describing the approach. Let us consider as MOI m the first four minutes of the MOVIE *Galileo* (Cavani, 1968), that have been filmed inside the Anatomical Theatre of the Palazzo Bo, the main building of the University of Padua, located in Via VIII Febbraio. Two SOIs, with different level of detail, could be inserted by users u and w :

$$\bar{s}_{u,m} = \langle \text{Italy, Veneto, Province of Padua} \rangle$$

$$\bar{s}_{w,m} = \langle \text{Italy, Veneto, Province of Padua, Padua, Via VIII Febbraio} \rangle$$

where the second represents also the ground truth, which has the highest level of accuracy according to film scholars. It is important to note that the different levels need to be consistent, but this can be obtained through the user interface that computes whether the insertion corresponds to a real location or not.

We propose to measure the difference between two symbolic descriptions through a binary similarity. Starting from α_0 the two arrays are paired until a difference is found or it is reached the end of the shortest one: the longest the common subsequent the highest the similarity. Moreover, it is possible to give a weight ρ^k to each level of α^k , with $\rho^0 \leq \rho^1 \leq \dots \leq \rho^N$ in order to penalize very short common subsequences. A conceptual algorithm that outputs the difference between SOIs $\bar{s}_{u,m}$ and $\bar{s}_{w,m}$, with respective length Lu and Lw is the following:

```

 $\Delta \leftarrow MAX\_DIST$  // maximum distance between SOIs
 $L \leftarrow \min\{Lu, Lw\}$ 
for  $i \leftarrow 0$  to  $L$  do
  if  $\alpha_{u,m}^i = \alpha_{w,m}^i$  then
     $\Delta \leftarrow \frac{1}{\rho^m}$ 
     $i \leftarrow i+1$ 
  else
    return  $\Delta$ 
  end if
end for
return  $\Delta$ 

```

3.1 SOI Insertion

Users are completely free to insert new SOIs, and each new contribution is also managed as a reinforcement of similar records that are already present in the system. For instance in the case of MOI m , user u can insert the array $\bar{s}_{u,m} = \langle \alpha_{u,m}^0, \dots, \alpha_{u,m}^V \rangle$. If another user w has already inserted array $\bar{s}_w = \langle \alpha_{w,m}^0, \dots, \alpha_{w,j}^N \rangle$, with $N \geq M$ and with the same first M elements, his/her contribution automatically obtains a reinforcement. For instance, if for

the first four minutes of *Galileo* a user already inserted the array $\bar{s}_{w,m}$ that has an accuracy until the street number, the insertion of $\bar{s}_{u,m}$ with same location but a lower accuracy gives automatically a positive weight to $\bar{s}_{w,m}$. In case of SOIs having the geographical component \bar{g} they are considered equal only if the distance is within a given threshold. Since we want to encourage participation, the user is asked to give his/her contribution also by voting the quality of the existing contributions, without inserting new SOIs.

3.2 Evaluation of Users' Insertions

As described in Section 2, users need to be motivated to contribute to geolocation and in general to the different crowdsourcing tasks. This can be based on the fact that each contribution is evaluated by other users, providing an indirect judgment about user's knowledge about audiovisual content and the geographical areas. Users participate to the activity to gain reputation as experts in the subject. We propose the following evaluation mechanism.

Starting from the SOI $\bar{v}_{w,m} = \langle \bar{s}_{w,m}, \bar{g}_{w,m} \rangle$, each user $u \in U$ with $u \neq w$ can provide a negative, neutral, or positive judgment to all the elements that constitute $\bar{v}_{w,m}$. Being $\bar{g}_{w,m}$ the most accurate description, a positive judgement is inherited by all the elements of $\bar{s}_{w,m}$. On the contrary, a negative judgement to any of the α^i of $\bar{s}_{w,m}$ is propagated to all $\alpha^k \in \bar{s}$, with $k < s$, and to $\bar{g}_{w,m}$ if present. The possibility to provide a partial judgement to a SOI allows us to highlight the potentially correct symbolic components, encouraging the community to cooperate inserting new SOIs that refine existing and incomplete information. Using again the example of the movie *Galileo*, if $\bar{s}_{w,m} = \langle Italy, Veneto, Province of Padua, Padua, Via VIII Febbraio \rangle$ receives positive judgements to the component *Province of Padua* and negative judgements to *Padua*, it is likely that the real location is within the province but outside the city of Padua.

Judgements are weighted by user's reputation, which in turn depends on the number positive and negative judgements received by the user. It is likely that reputation can be split in *general reputation*, which takes into account the general reliability of a user, and in *local reputation*, which is related to the ability of a user to carry out correct and precise identification in a given geographical area. The way reputation is computed is part of our future work.

3.3 Experiments

We are carrying out a first test of the proposed approach, to verify the feasibility of a crowdsourcing approach. We developed a web portal that is currently used by 103 undergraduate students of *Design and Management of Cultural Tourism*, who do not have a training in film studies. We assigned the following tasks, to be completed in subsequent steps.

1. Identify and insert in the system audiovisual productions that have been filmed in Veneto, specifying the source of information and describing the reason of the insertion.

2. Confirm the insertions made by other students, also in this case describing the reason of the confirmation (to this end, the reason of the insertion was not available to other students).
3. Geolocate on GoogleMaps interface the MOIs of the already inserted movies, specifying start and end time of the MOI; geolocation was possible only for audiovisual content available on YouTube.

Figure 1 shows a screenshot of the portal homepage, while Figure 2 shows the web interface used to geolocation. The first two tasks started together, while the third one started three weeks later. Students inserted 125 audiovisual productions – including movies, TV and web series – and provided 108 confirmations of already inserted productions. After one week from the beginning of task three students inserted 60 SOIs.

The reputation mechanism has not been introduced yet, mainly because our primary goal was to test the effectiveness of the web interface. Nevertheless the initial results are encouraging and we believe that the crowdsourcing mechanism can be extended also to general users. Students feedback was very important for the improvement of the system interface and to refine the model described in Section 3.

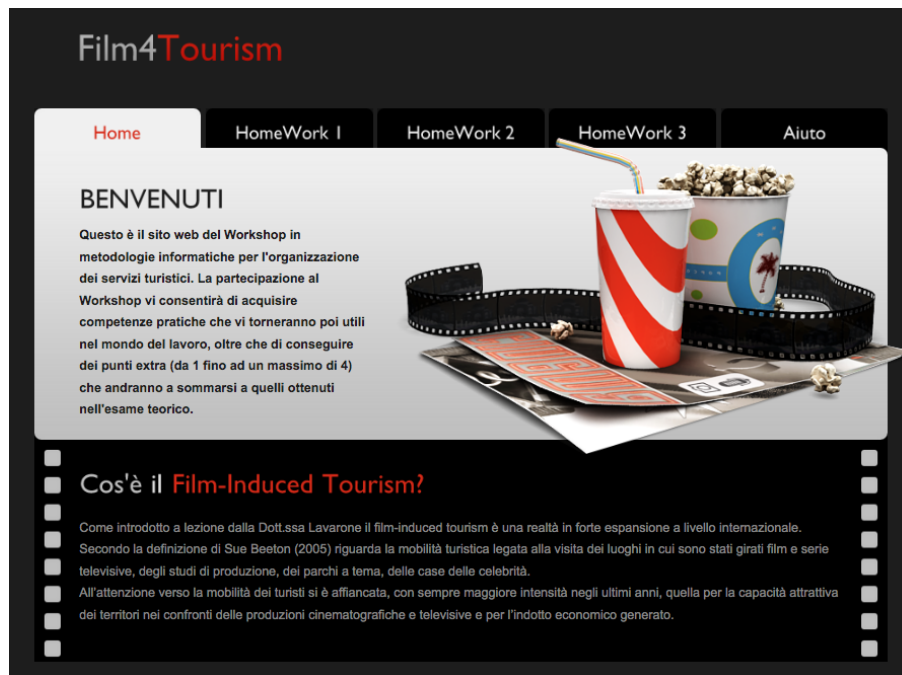


Fig. 1. Screenshot of the homepage of student's workshop on film-induced tourism

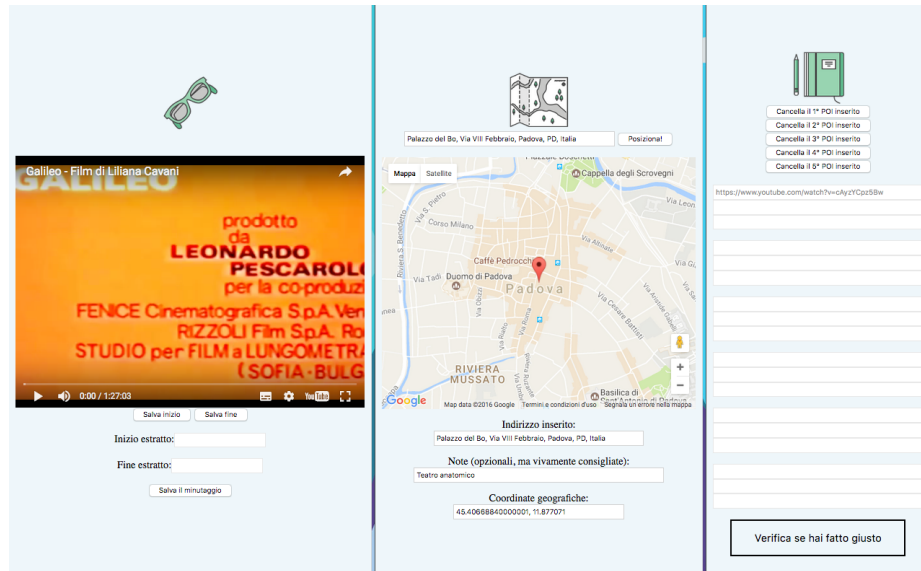


Fig. 2. Screenshot of the geolocation page of student's workshop on film-induced tourism

4 Conclusions and Future Developments

This paper described the main ideas behind the design and development of a crowdsourcing portal aimed at gathering information for film-induced tourism. Among the different tasks, the geolocation of audiovisual excerpts plays a central role for film-induced applications and it thus should be carried out reliably also when non expert users are involved. This goal can be achieved by reinforcing user engagement through the concept of reputation, that is the social status of being acknowledged as an expert by a social community. The case study presented in Section 3.3 showed that the approach is feasible: the information gathered during this initial step are the basis for the development of the crowdsourcing platform that will be starting point for the development of the crowdsourcing platform.

Our future activities will address the complete definition of the reputation mechanism. Moreover, we envisage to include a number of gamification tools to increment users' involvement in the crowdsourcing effort.

Acknowledgements

The research project presented in this paper is funded by the European Social Fund (FSE projects, Veneto Region). We thank our colleagues Farah Polato, Giulia Lavarone of the Department of Cultural Heritage and Sandro Savino of the Department of Information Engineering, both at the University of Padua.

References

1. Bares, W., Schwartz, D.: Film ties: Crowd-sourced teaching of cinematography using intelligent example galleries. In: Proceedings of the 19th ACM Conference on Computer Supported Cooperative Work and Social Computing Companion. pp. 229–232. CSCW '16 Companion, ACM, New York, NY, USA (2016)
2. Beeton, S.: Film-induced Tourism, p. 11. Channel View Publications, Clevedon(UK) (2005)
3. Bourdieu, P.: The Forms of Capital. In J.C. Richardson (Ed.) Handbook of Theory and Research for the Sociology of Education, New York: Greenwood (1985)
4. Can, G., Odobez, J.M., Gatica-Perez, D.: Is that a jaguar?: Segmenting ancient maya glyphs via crowdsourcing. In: Proceedings of the 2014 International ACM Workshop on Crowdsourcing for Multimedia. pp. 37–40. CrowdMM '14, ACM, New York, NY, USA (2014)
5. Di Cesare, F., Rech, G.: Le produzioni cinematografiche, il turismo, il territorio, pp. 45–88. Carrocci, Roma (2007)
6. Figuerola Salas, O., Adzic, V., Shah, A., Kalva, H.: Assessing internet video quality using crowdsourcing. In: Proceedings of the 2Nd ACM International Workshop on Crowdsourcing for Multimedia. pp. 23–28. CrowdMM '13, ACM, New York, NY, USA (2013)
7. Howe, J.: The rise of crowdsourcing. Wired Magazine (June 2006)
8. Lavarone, G., Orio, N., Polato, F., Savino, S.: Modeling the concept of movie in a software architecture for film-induced tourism. In: Digital Libraries on the Move - 11th Italian Research Conference on Digital Libraries, IRCDL 2015, Bolzano, Italy, January 29-30, 2015, Revised Selected Papers. pp. 116–125 (2015)
9. Wellman, B., Wortley, S.: Different strokes from different folks: Community ties and social support. *American Journal of Sociology* 96(3), 558–588 (1990)