Detecting attractive locations using publicly available user-generated video content – central Serbia case study

Milan Mirković, *Member, IEEE*, Dubravko Ćulibrk, *Member, IEEE,* Stevan Milisavljević and Vladimir Crnojević, *Member, IEEE*

***Abstract* — Emergence of various user-generated content sharing platforms led to tremendeous amounts of data becoming publicly available. Some of that data is geo-referenced, that enables anyone interested in it to access, analyze and visualize it in different ways. Research has already been conducted using user-contributed images (Flickr, Panoramino) and videos (YouTube) to reveal possibly attractive regions in an area. This paper presents results of analysis based on geo-visualization tecniques that used data obtained from publicly available geo-referenced videos available on Youtube, tagged as recorded in Central Serbia region. Density analysis is performed to detect and visualize the areas of interest in the region, using state-of-the-art visualization tools. Results presented indicate that YouTube videos can be used to identify places tourists and people in general find interesting, as well as analyze the temporal dynamics of their activity. In addition, the category of the videos provided on YouTube, can be used to discover what kind of attractions are of primary interest to the people visiting central Serbia, as well as those living in this area.**

***Key words* — Clustering, Geo-Information Systems, Geo-Visualization, Publicly available data, Video, Youtube.**

I. Introduction

Nowadays, many web platforms exist that provide users with easy, fast and convenient means of sharing various information with the online community. Social networks (e.g. Facebook, Myspace and LinkedIn), image and video-sharing services (e.g. Flickr, Youtube, Picasa) and blogs are just some of them. What all of those applications and services have in common is that they store a large amount of user-generated data, that is just a query away from anyone interested in taking a look at it. In addition, advancements in mobile technologies enable owners of various mobile devices (ranging from phones to netbooks) to create and share new content "on the go", often tagging that content with current date, time and location – thus adding a temporal and/or a spatial dimension to it. This temporal dimension is especially useful for visual analysis of geo-referenced data, since it can reveal some interesting patterns in human behaviour. Studies on this matter have already been undertaken, using geo-referenced images publicly available on Flickr and Panoramio [1][2]. A recent paper considered using YouTube videos for these purposes [3]. Youtube was chosen as a source of videos used in our analysis, because it is the most popular platform for sharing user generated video content [4], and offers an option to geographically reference uploaded videos. An automated tool developed in-house was used to gather video data that was geo-referenced as recorded in central part of Serbia. Obtained data was then processed and presented using several geo-visualization techniques, and we report results of those analysis in this paper.

The rest of this paper is organized as follows: Section II presents a brief overview of related work in the field, Section III describes the proposed approach and methods used in analysis of data extracted from Youtube videos, Section IV presents the results obtained and Section V holds the conclusions and some possible directions for future work.

II. Related work

It should not come as a surprise that a lot of research effort has recently been put into (visual)analysis of published user location data, since it can complement frequentation data gathered in traditional ways - via surveys and questionares - and enable urban planners, local authorities and traffic engineers to gain a more thorough understanding of urban environments. Kisilevich et al. [1] try to map what they refer to as attractive areas in a region, based on cluster analysis of images uploaded by Flickr and Panoramino users. Gennady et al. [5] introduce a temporal dimension to geo-referenced image analysis by extrapolating trajectories of people who took pictures in an area, using image timestamps alongside their spatial coordinates. Girardin et al. [6] combine geo-spatial and temporal image data with mobile phone usage data to separate tourists from residents in an area, and then reveal interesting patterns in tourist movement and behavior. In another study, Girardin et al. [7] use only publicly disclosed data (i.e. Flickr images) to analyze tourist dynamics in the region of Florence, Italy – besides uncovering hot spots of tourist concentration, they also reveal patterns of flow from and to the city of Florence. Building on the approach of Girardin et al., Mirkovic et al. [3], showed that YouTube videos can be used to analyze sites attractive to tourists. That approach is what was used

Milan Mirković, Faculty of Technical Sciences, Trg Dositeja Obradovića 6, 21000 Novi Sad, Serbia (phone: 381-21-4852117, e-mail: mmirkov@uns.ac.rs)

Dubravko Ćulibrk, Faculty of Technical Sciences, Trg Dositeja Obradovića 6, 21000 Novi Sad, Serbia (phone: 381-21-4852169, e-mail: dculibrk@uns.ac.rs)

Stevan Milisavljević, Faculty of Technical Sciences, Trg Dositeja Obradovića 6, 21000 Novi Sad, Serbia (phone: 381-21-4852177, e-mail: steva@uns.ac.rs)

Vladimir Crnojević, Faculty of Technical Sciences, Trg Dositeja Obradovića 6, 21000 Novi Sad, Serbia (phone: 381-21-4852520, e-mail: crnojevic@uns.ac.rs)

as basis for the study presented here, aimed at understanding what information about the attractions in central Serbia may be gained from geo-referenced publicly available videos.

III. Approach and methods

An automated tool that relies on a Youtube crawler [8] was developed in-house to obtain video location data. It was then used to query for videos that were tagged as recorded in wider central Serbia area. The result were 23,869 unique records containing video id, latitude and longitude, date and time published, category, duration, number of views and username of the uploader. These data were then fed to CommonGIS[1] (Visual Analytics ToolKit) and Open Heatmap[2] to create maps of uploaded videos' spatial density that were then used for further analysis.

In order to create density maps for different levels of detail (e.g. province, city, parts of city), a matrix of square cells covering the area had to be created. The following method was applied: first the size of the cell is determined depending on the level of details needed (2km, 100m and 10m respectively) then the number of videos recorded in each cell is counted and finally latitude and longitude coordinates are assigned to the center of the cell. Thus, latitude, longitude and number of videos for each cell is obtained, which are the parameters needed for constructing a heat map (or choropleth map). In addition, density-based cluster analysis was performed using the OPTICS (Ordering Points To Identify the Clustering Structure) algorithm [9]. Density-clusters are built using only two parameters: the neighborhood radius and a minimum number of points in the neighborhood. The clusters can have different shapes and there is no need to pre-define their number.

A. Limitations and challenges

There were two potential problems we had to take into consideration when analyzing data: a) there is no guarantee on the degree of accuracy of coordinates provided by the uploaders of videos and b) Youtube doesn't allow for pure spatial queries (e.g. entering desired latitude, longitude and radius, but without a search term). First possible pitfall was overcome to a certain degree by restricting search only to videos that did contain complete geospatial data (i.e. latitude and longitude) – since tech savvy users who go to trouble of pinpointing their location when recording a video are likely to use those precise coordinates to describe its location when they publish it online. The latter issue was addressed by repeatedly querying for videos, using an Serbian and an English frequently used words list as a basis for search terms. This yielded a large dataset, from which duplicate entries were removed prior to further analysis.

Since no Serbian words frequentation lists were readily available, one (comprising 2000 words) was constructed using the following method: first a list of popular Serbian on-line news sources was made (comprising local as well as national newspapers websites, news portals and agencies, etc.), then available RSS feeds from those websites were downloaded and converted to pure text format so words could be extracted and their frequency counted. Each of the 2000 words that appeared most frequently in those feeds was in turn used as a search criteria. It is worth mentioning that visual skimming of the list confirmed that words which appeared in the list are indeed used very often in everyday communication.

IV. Analysis and results

Out of 23,869 videos used in our analysis, less than 0.1% were published in 2006, just under 4% in 2007, 19% in 2008, 40% in 2009 and 37% in 2010. This overall rising trend is not surprising, considering that Youtube was founded in early 2005 and officially launched in November 2005 – afterwards, it has only gained popularity.

It is also interesting to note that search using Serbian words frequentation list returned 18,838 results, while Serbian-English list returned 23,869 results, suggesting that about 20% of videos geo-referenced as recorded in Serbia contain non-Serbian words (or at least words that are not so commonly used in Serbian news-providing services) in their description.

Table I shows some basic statistics for the dataset used. The queries were restricted to a 150km radius around 44.0000, 21.0000 (lon, lat) coordinate, since it encompasses most of the area of interest.

Table I: Basic data statistics

Number of videos	***23,869***
Time frame	2006-2010
Mean number of views	6,708
Median number of views	375
Mean duration	214 s
Median duration	190 s
Unique users	11,245

As can be seen from Fig. 1, uploaders seem to be most active during the summer. This doesn't come as a surprise, considering that during those months (July – September) a lot of working-class people, as well as students, usually take their time off and go to vacations.

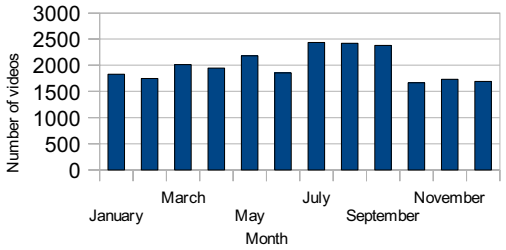

Fig. 1. Number of videos uploaded by month of the year

Incidentally, it is also the time when „Exit" festival is held (during the first half of July), which is known to attract many visitors from abroad – many of whom come from the United Kingdom, so it might be one of the factors that contributes to the high number of results

[1]http://geoanalytics.net/VAtoolkit

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

returned when English word list was used for the search.

As for the time of day when videos were uploaded the most, results don't deviate from what one might intuitively expect; majority of videos were uploaded during the 10 a.m. - 12 p.m. period, as shown in Fig. 2.

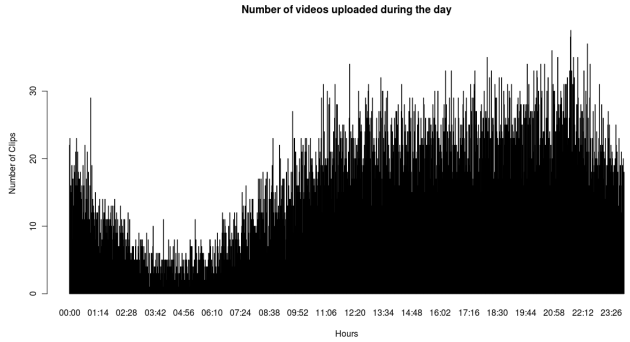

Fig. 2. Number of videos uploaded by time of the day

To visually represent concentration of uploaded videos on different scales the Open Heatmap tool was used. Figure 3 shows heat maps for: northern part of central Serbia region (a), Novi Sad city area (b) and Novi Sad wider-city centre (c). Yellow, orange and red colors represent low, medium and high amount of video concentration (number of videos per unit of area), respectively. Scrollable and zoomable versions of maps presented here can be accessed online[3].

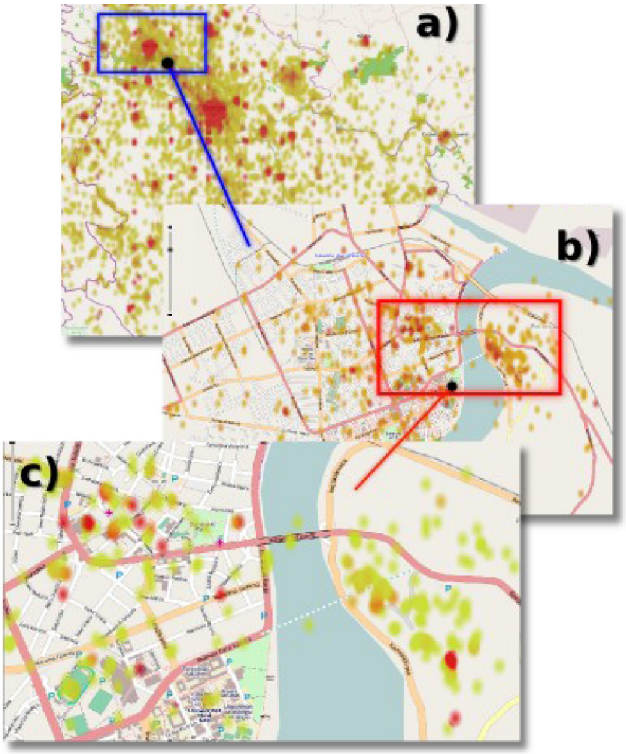

Fig. 3. Heatmaps generated using Open Heatmap tool

As can be seen from Fig. 3(a), most of the videos were taken (or tagged as taken) in, or close to major Serbian cities, namely Novi Sad, Belgrade, Nis and Kragujevac. Other hotspots include smaller cities such as Zrenjanin, Smederevo, Pozarevac, Sabac, Valjevo, Cacak, etc. Areas farther from settlements rarely have a high concentration of videos taken there, with an exception of Kopaonik and Fruska Gora mountains; which is understandable since Kopaonik is a popular winter resort, and Fruska Gora hill has many monasteries, monuments and other attractions related to national heritage. In Novi Sad City area Fig. 3(b), only two locations seem to be attractive – first one being the city centre, and second one the Petrovaradin fortress (where, among other things, „Exit“ festival is held annualy). Closer inspection of Fig. 3(c) reveals that the most attractive locations in the city of Novi Sad appear to be the National Theatre building, the city centre square and several main streets leading to and from it. As for the Petrovaradin fortress, results were a bit unclear at first – it looked like most of the videos were reffering to an area without a monument, known tourist attraction or even a building in it. But this phenomenon made a perfect sense after a few videos from that location were viewed; it is the place where one of the main stages during the „Exit“ festival (the so called „DJ Arena“) is located, that seems to attract a very specific profile of people (usually young people who listen to modern music – techno, house, etc. - and appear to be technically savvy and keen to use their mobile devices to record atmosphere at the concerts, and share them later on with the community).

Visual Analytics Toolkit for Spatial and Temporal Data was used to explore data further, and Fig. 4 shows a map plotted using data obtained for the wider-centre area of Belgrade city.

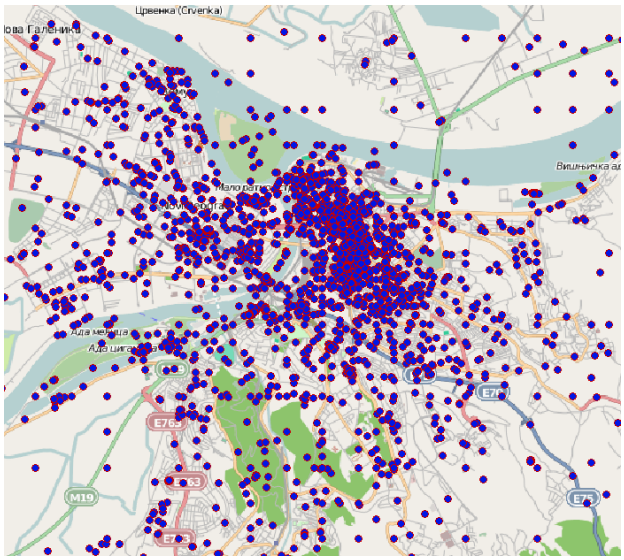

Fig. 4. Videos recorded in wider-centre area of Belgrade city

Each dot represents one video from the dataset. While it is obvious that dots tend to be tightly-grouped near the centre of the city where the majority of tourist attractions are located, in order to further - and more in-depth - investigate potentially attractive locations, a cluster analysis using OPTICS algorithm was performed. Results of this analysis can be seen on Fig. 5.

[3]Http://www.iim.ftn.uns.ac.rs/iks

Fig. 5 Results of OPTICS clustering for the wider-centre area of Belgrade city

Besides Kalemegdan Fortress, City centre and nearby parks („Tašmajdan park“ and „Botanic garden“) – which were all intuitively expected to be attractive to residents as well as tourists – it appears that entertainment-related locations seem to rate pretty high when it comes to uploading videos. Two soccer stadiums (those of „Partizan“ and „Crvena Zvezda“ football clubs) and two well-known concert and show venues („Belgrade Arena“ and „Ušće“) were marked significant when 100m and 5 neighborhood objects parameters were used for OPTICS analysis. This observation was further supported by filtering plotted videos according to the category they were assigned to (upon uploading a video, a user has to choose among one of 15 pre-defined categories), and Fig. 6. shows a map with only Music and Sports categories

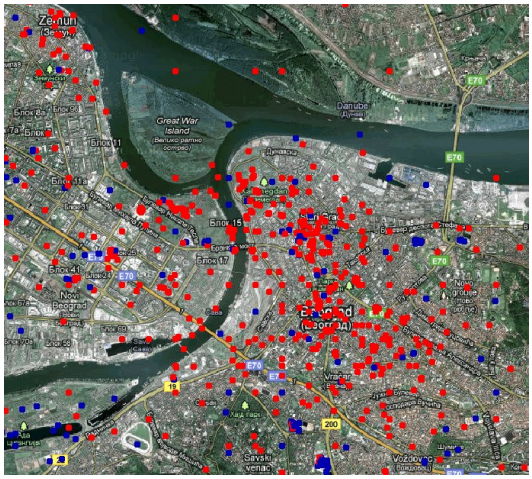

Fig. 6. Videos tagged as related to (containing) Music (red) or Sports (blue) content

plotted (it can be seen pretty clearly that blue color, representing Sports catgory, dominates „Partizan“ and „Crvena Zvezda“ stadiums, and that red color takes up the area just across the „Great War Island“, where „Ušće“ park is located, in which many concerts were held).

V. Conclusion

Contemporary user-generated content sharing platforms, along with the explosion of mobile devices market created a suitable climate for abundance of geo-referenced material to become publicly available. This material can then be used for data extraction that makes various urban and mobility analysis possible.

This paper presents a case study where data obtained from publicly-available videos was used to identify possible attractive locations in northern part of Serbia. The results presented show that a link between the detected clusters / generated heatmaps and locations already known to be visited by a large number of people exists. The main difficulty in the work presented lies in the scarcity of data available – number of videos obtained for the part of Serbia where majority of its population lives hardly reached 25,000; whereas a similar study (still ongoing as of time of writing this paper) yielded over 70,000 results for the northern part of central Italy alone. Such a large difference might mean that Italy has more to offer (regarding tourism, since tourist attractions usually generate a noticable footprint in analysis such as the one presented in this paper) than Serbia, that people living in (or visiting) Northern Italy tend to be more tech savvy and/or willing to share the content they recorded. However, as the number of videos uploaded each year increases, one should expect that analysis such as that presented in this paper will be able to provide more insight into the true attractions of Serbia, in the near future.

Acknowledgment

This paper was supported by the Ministry of Science and Technological Development of Republic of Serbia, under the grant number TR-11022.

References

[1] S. Kisilevich, F. Mansmann, P. Bak, D. Keim, and A. Tchaikin, “Where Would You Go on Your Next Vacation? A Framework for Visual Exploration of Attractive Places,” in 2010 Second International Conference on Advanced Geographic Information Systems, Applications, and Services. IEEE, 2010, pp. 21–26.

[2] F. Girardin, F. Dal Fiore, J. Blat, and C. Ratti, “Understanding of tourist dynamics from explicitly disclosed location information,” in 4th International Symposium on LBS and Telecartography, Hong-Kong, China, 2007.

[3] M. Mirkovic, D. Culibrk, D. Stefanovic, V. Crnojevic, “Visual Exploration of Attractive Spaces and Urban Dynamics Based on Community-Contributed Videos”, Joint Urban Sensing Event, submitted for publication

[4] M. Cha, H. Kwak, P. Rodriguez, Y. Ahn, and S. Moon, “I tube, you tube, everybody tubes: analyzing the world’s largest user generated content video system,” in Proceedings of the 7th ACM SIGCOMM conf. on Internet measurement. ACM, 2007, p. 14.

[5] G. Andrienko, N. Andrienko, P. Bak, S. Kisilevich, and D. Keim, “Analysis of community-contributed space-and time-referenced data (example of Panoramio photos),” in Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems. ACM, 2009, pp. 540–541.

[6] F. Girardin, F. Calabrese, F. Fiore, C. Ratti, and J. Blat, “Digital footprinting: Uncovering tourists with user-generated content,” Pervasive Computing, IEEE, vol. 7, no. 4, pp. 36–43, 2008.

[7] F. Girardin, F. Dal Fioreb, C. Rattib, and J. Blata, “Leveraging explicitly disclosed location information to understand tourist dynamics: a case study,” Journal of Location Based Services, vol. 2, no. 1, pp. 41–56, 2008.

[8] C. Shah, “Tubekit - a query-based youtube crawling toolkit,” Bulletin of IEEE Technical Committee on Digital Libraries, Spring 2009.

[9] M. Ankerst, M. Breunig, H. Kriegel, and J. Sander, “OPTICS: Ordering points to identify the clustering structure,” ACM SIGMOD Record, vol. 28, no. 2, p. 60, 1999.