

# 3D georeferencing of historical photos by volunteers

Timothée Produit and Jens Ingensand

**Abstract** Historical photographs are a very rich source of information that can be useful in a variety of different contexts such as environmental analyses, land planning and illustration of landscape evolution. However to reach this goal the images must be accurately georeferenced. In this paper we propose to use the crowd to perform the 3D georeferencing of collections of historical images. To this goal we implemented a web 3D georeferencer that offers volunteers the possibility to semi-automatically identify 1. the location of the point from where a picture has been taken, 2. the three angles: tilt, roll and yaw and 3. the field of view. A virtual web-based globe is the central element in this tool that allows both for the georeferencing in three dimensions by volunteers and for the visualization of georeferenced images to assess the landscape variation through time. In this paper we evaluate the method and the georeferencer and give suggestions for further developments and exploitation of the database.

## 1 Introduction

Historical photographs contain information about both natural and man-made processes that have occurred in the past and that have an influence on today's world and future development. These photographs are often stored in archives that are difficult to access for researchers. Moreover, historical photographs are rarely scanned

---

Timothée Produit

Territorial Engineering Institute (insit) University of Applied Sciences and Arts Western Switzerland (HES-SO) Route de Cheseaux 1, CH-1401 Yverdon-les-Bains, e-mail: [timothee.produit@heig-vd.ch](mailto:timothee.produit@heig-vd.ch)

Jens Ingensand

Territorial Engineering Institute (insit) University of Applied Sciences and Arts Western Switzerland (HES-SO) Route de Cheseaux 1, CH-1401 Yverdon-les-Bains, e-mail: [jens.ingensand@heig-vd.ch](mailto:jens.ingensand@heig-vd.ch)

and stored in digital files. Another problem is the precise georeferencing of these photographs. Indeed if the exact location and orientation of a picture is known, the picture can be visualized as a 3D object for instance in a virtual globe. Moreover the georeferencing of a picture in three dimensions has an important value since it enables an exact identification and indexing of places that are visible in these images. Furthermore the georeferencing also has a value for scientists who can extract geographic data from the pictures using photogrammetric methods and thereby for instance recreate three-dimensional models of destroyed buildings or calculate the volume of a glacier at different times.

The rise of web-based 3D solutions opens opportunities for volunteers to generate and interact with three-dimensional data. Today, the public is increasingly familiar with virtual globes and 3D software such as video games. In the project that we present in this paper we focus on the georeferencing and visualization of historical images using a 3D virtual web-based globe (Produit and Ingensand, 2016). In order to compute the location and orientation of a picture, a user needs to manually provide 3D coordinates of several visible locations in the picture. This task can for example be achieved by measuring 3D points in the real world using surveying instruments, by extracting the 3D locations in a GIS or by digitizing 3D points in a virtual globe. Hence, in the method that we present in this paper we want to use the familiarity with 3D software and to offer a tool for volunteers for the georeferencing of historical images. In this tool the goal is to identify and to digitize similar locations that are visible both in the picture and in a virtual globe. These correspondences allows us to compute the exact location and orientation of the picture and to visualize the picture as a 3D object in a virtual globe.

Our first case study was a set of 1500 postcards shared with us by the *Archive de la Construction Moderne (ACM) - EPFL* representing the swiss Alps shot in 1960.

This paper is structured as follows. First we will identify different ways of collecting 3D information by volunteers and georeferencing images. Thereafter we will explain our idea and the tools that we have implemented. Then, we will describe the first case study. In the following paragraphs we will analyze and evaluate how real-world users use the tools in order to georeference historical photos. Finally ideas for future work and for an exploitation of the collected information are presented.

## **2 State of the art**

### ***2.1 Approaches to collecting 3D data using volunteers***

The digitization of 3D objects can be regarded as a challenging task: most existing software that can be used for this task is generally dedicated to professionals. Several researchers have addressed the challenge to involve volunteers for the digitization of 3D objects and to design this task as easy as possible.

A first approach for volunteers to provide 3D data is to add attributes to 2D objects such as a buildings height or the type of a roof. These attributes can then be used for the rendering of a 3D object (Goetz and Zipf, 2013). A second approach is to offer tools for volunteers to perform measurements in 3D. For instance a user can use the GPS of a smartphone to record a position in 3D (Brovelli et al., 2013). Pictures are a specific kind of data. Associated with a GPS tag, pictures can be used to compute 3D models. Some authors (Snavely et al., 2006; Strecha et al., 2010; Hartmann et al., 2016) show that pictures downloaded from photosharing platforms can be used to generate 3D point clouds. A similar technology is used in the project Mapillary. In this web platform volunteers can provide smartphone pictures that have been collected along a path. Using computer vision algorithms Mapillary generates several outputs based on the added pictures such as a dense 3D point cloud, however it is not possible yet to derive 3D GIS objects such as 3D buildings from the point cloud.

By providing attributes or measurements, a volunteer passively provides 3D data and does not have to navigate a 3D interface to digitize 3D data. Indeed, the digitization of 3D data is a challenging task for a volunteer. In the project *Building Maker* that has been developed until 2013, Google asked volunteers to draw buildings in oblique aerial images. The users were required to align and to reshape a 3D building model using 2D pictures. In this project, the georeferencing task remained accessible to volunteers. Uden and Zipf (2012) suggest to involve volunteers with 3D digitization skills in the project OpenStreetMap. To reach this goal, the authors provided a repository of 3D buildings models which can be associated with OpenStreetMap objects.

## 2.2 3D Georeferencing of (historical) photos

The 3D georeferencing of a picture requires that the 3D location (X, Y, Z) and 3D orientation (three angles) as well as some camera parameters such as the field of view are computed. For a similar task Google Earth requires the volunteer to manually provide the picture parameters *location, orientation and field of view*. The volunteer navigates the virtual globe and refines both the orientation and location parameters with a slider. This method however has the drawback that it has the potential to be time-consuming and less accurate.


The regular approach applied in the computer vision and photogrammetry communities is 2D-3D image orientation (camera orientation or space resection). The picture parameters are calculated using Ground Control Points (GCP) which are 3D points that have been measured in the field or in a map and that have corresponding 2D points in the image (see Figure 1). In photogrammetry this concept is generally applied to compute the location of a collection of overlapping images in a bundle adjustment. 2D-3D image orientation is also applied in several projects dedicated to the extraction of geographic data from single pictures (Bozzini et al., 2012; Produit and Tuia, 2012) or defined in the project Viewfinder (University of Southern Cali-

fornia, 2008) to insert an image in a 3D model. In this way, similarly to the project *Building Maker*, photogrammetric relations between the reality and a picture simplify the task of finding the image georeferencing. By using these relations the users does not need 3D modelling skills as the task of clicking points in a map or a virtual globe remains less complicated.

### 3 Methodology

#### 3.1 Georeferencing

Our concept, similarly to the project Viewfinder (University of Southern California, 2008) relies on the idea to establish GCP of a single image using a virtual representation of the reality rather than measurements in the field. This implies that a virtual globe needs to be positioned roughly at the same location as a corresponding photo. Thereafter a user can digitize GCP both in the virtual globe and in the 2D photo using locations that are both visible in the photography and in the virtual globe. The key in this concept is to use the workforce as well as the place knowledge of volunteers in order to accomplish this task.



**Fig. 1** Ground Control Points: 2D locations in the image and corresponding 3D locations in the virtual globe.

A photo often has a location name associated with it. This location name can be used for a rough georeferencing and positioning within a virtual globe. However, the most difficult task of the georeferencing process is to navigate the virtual globe in order to reach the location represented in the picture. To facilitate this process, we have identified four steps. All steps imply that the user is familiar with the scene represented in the picture and also approximately with the topography of the area.

- First, the volunteer navigates a 2D map. The images are marked in the map according to their *a priori* location which has been geocoded from a place name. The volunteer chooses a picture.
- Second, the volunteer provides an approximate location of the picture by indicating the 2D location of the camera position on a 2D map. If the original photo has

been associated with a location name, this location name is used to automatically position the 2D map.

- Third, the user provides the orientation of the picture.
- Fourth, the provided location and direction are used to approximately position a virtual globe. At this point, if the two previous steps have been completed successfully, the volunteer should be able to see a similar scene in the 3D globe and in the 2D photograph. Now the task of the volunteer is to identify points that are visible both in the virtual globe and in the photograph. These points must also be stable from the time when the picture had been taken until today. Examples for such points that are easy to identify are mountain peaks, churches, crossroads and shapes of water-bodies. Once that at least four GCP have been provided, the picture location can be computed.
- Finally, the user has to provide two more GCP to improve the accuracy. Functionalities to move and delete inaccurate GCP are also provided.

### ***3.2 Validation of the georeferencing***

In the presented project, the alignment of the picture with the virtual landscape is a good indicator of the picture location and orientation accuracy. Moreover, the visualization of 3D images is central since it allows the visitors to compare the current virtual state of the landscape with the historical state shown in the historical picture. However, in order to have an optimal experience during this comparison, the 3D image must have a good alignment with the virtual globe. To ensure this alignment accuracy we have implemented two validations steps.

We implemented automatic checks to avoid that the volunteers provide incorrect locations. First, at least six GCP must be provided for each image. With six correspondences, more GCP are available than required to solve the equations. Hence, an error can be computed for each GCP. Namely, we compare the 2D coordinates digitized by the volunteer with the corresponding 2D coordinates computed by our algorithm. If these errors are above a certain threshold, the picture location cannot be saved.

A second check is based on the computed picture parameters. If the computed field of view is impossible (too small or too large) the picture location is not accepted.

Finally, a manual visual validation is performed by our team of advanced users. They visualize each picture and check its alignment with the virtual landscape. If the alignment is not satisfactory, the validator can improve the correspondences or reject the georeferencing.


## 4 Case study

### 4.1 Implementation

The web-based prototype has been created using the open source virtual globe API [www.cesium.com](http://www.cesium.com). The virtual globe renders data (aerial images and a digital terrain model (DTM)) provided as web services by the Swiss Federal Office of Topography Swisstopo. The aerial images have a resolution ranging from 25 to 50 cm and the DTM has a resolution of 1m. The virtual globe is used for the georeferencing but also for the 3D visualization of historical images. Historical images thereby can be compared with the current virtual landscape.

The 2D-3D image orientation algorithm uses the GCP provided by the volunteer to compute the camera position from which the 3D locations of the image corners can be derived. The camera parameters are stored in a PostgreSQL/PostGIS database. The image is stored as a 3D entity in the *glTF* 3D format which is the format used by Cesium.

The website can be accessed at: <https://smaphot.heig-vd.ch>.



**Fig. 2** Screenshot of the 3D interface. A picture of the Swiss alps shot in 1910 by E. Spelterini from an hot-air balloon. The picture is visualized as an overlay on top of the virtual landscape.

## 4.2 Case study setting

The prototype was officially made public on February 1<sup>st</sup> 2017. An exhibition was organized by the Archives de la Construction Moderne (ACM), EPFL. The ACM also provided the first photo collection; 1200 postcards shot between 1960 and 1970 showing the swiss Alps. These photos were scanned and delivered with the corresponding metadata which typically contained the photographer's name, a precise or approximate date and a place name. The geocoded place name allowed us to provide an *a priori* image location.

A key issue was the acquisition of visitors and volunteers. In order to attract volunteers, we provided a set of 50 already georeferenced photos that could be visualized in the virtual globe. Hence, the first visitors were able to see the results and to understand the goal of the project. Second, to reach the volunteers a media campaign was organized by both universities (EPFL and HES-SO) communication teams. Hence, the project's aim was published in newspapers, radio, TV and web media. This first 1200 photos were georeferenced within 10 weeks.

After this first collection had been published, several other collections were retrieved from open data repositories. The collections contain both terrestrial and aerial pictures taken between 1850 and 1970. Currently, the prototype stores 2500 images among which 2300 have been georeferenced by 90 registered volunteers.

## 4.3 Hypotheses

In the presented project volunteers are asked to provide the exact position and orientation of a historical picture. To reach this goal, the volunteers have to navigate a 3D virtual globe, click on points which are both visible on the surface of the DTM and in the picture. In this context we have established two hypotheses which are about user participation and spatial data quality:

1. Volunteers are interested in the georeferencing task: they are attracted by the problematic of historical images metadata acquisition and visualization. Moreover the volunteers like spending time executing these tasks.
2. Volunteers are able to perform the task: the proposed tasks are sufficiently easy for volunteers and the volunteers provide a correct georeferencing.

These two hypotheses are critical - if either the users do not show any interest or the collected data shows a low quality, the outcome would be less valuable. Another important point in the analysis of these hypotheses is the understanding of user habits in order to improve the platform at subsequent stages.

#### 4.4 Data collection


The platform records information about the visits and the participation of volunteers. For instance, for each georeferenced image the time it takes for a specified user to accomplish the georeferencing task is recorded. We also used Google Analytics in order to collect and analyze statistics about the users.

### 5 Results

#### 5.1 Volunteers' activity

The first major result from the analyzed data is that some platform visitors acquired during the media campaign were converted into volunteers who were able to use the 3D georeferencer to provide the location and orientation of images.


The number of visits and the number of new volunteers is strongly related to the publications in the media. Each publication was followed by a peak of georeferenced pictures (up to 100 pictures during one single day, see Figure 3).



**Fig. 3** Number of georeferenced images per day. Peaks occur after a publication in the media.

As in other crowdsourcing projects (e.g. Neis et al. (2013); Sauermann and Franzoni (2015)) there are few volunteers who provide most of the georeferencing and there are many volunteers who provide the georeferencing of only a few images (see Figure 4). The most active users were using the platform every day. Some users stopped after a time, either because there were no more images in their area of interest or the images became too challenging (e.g. presumably in regions they are less familiar with or images with few prominent landmarks) or because they get tired.





**Fig. 4** Ranking of the volunteers. The y-axis shows the number of images per volunteer. Few volunteers provide the georeferencing of the largest portion of images. Many volunteers tested the platform and provided the georeferencing of a small number of images.

## 5.2 Georeferencing assessment


Every image's georeferencing has been checked by our team. About 50% of the images were correctly georeferenced, the remaining images were slightly improved. This validation has the goal to ensure the quality of the alignment with the virtual globe but also to understand the main sources for errors:

- **Incorrect correspondences** : The indicated control point is not the same in the picture and in the virtual globe. The landscape variation over time, the similarity of the land-cover and the quality of the rendering may have caused this problem. A typical error is a confusion of buildings which might be caused by the fact that the virtual globe does not provide 3D buildings, but only an aerial image draped over a DTM. Moreover the buildings in the Alps were often more scattered in the past while the density has drastically risen during the last 50 years. Another example of this type of error is a lake shore which moved a hundred meters due to a change of its level.
- **Roof tops** : As aforementioned, 3D buildings are not available yet in the virtual globe. The volunteers who are less accustomed to geographic data often click on a roof top in the picture and on a point at ground level in the virtual globe. This error generates a GCP height error of several meters which can have a high impact for pictures that have been taken close to the buildings (Figure 5).
- **Correspondences scattering** : To improve the accuracy and the alignment, the correspondences should be scattered across the entire picture. Sometimes the volunteers provide correspondences only in a portion of the picture where the correspondences are easy to find.

- **Connected pointer.** Once that a camera position is computed with the four initial GCP, the mouse cursor in the virtual globe is connected with a pointer in the picture. Namely, if the user moves the mouse in the virtual globe, the mouse pointer position in the image is automatically calculated. We implemented this functionality to ease the task of the volunteer to find additional correspondences. However, we also noticed that this feature can mislead the volunteer: if the initial position is not exact due to the fact that some initial GCP are inaccurate, the picture pointer indicates an incorrect location. It appears that some volunteers trust the picture pointer too much and provide incorrect correspondences (Figure 6). This situation occurs mainly when the correspondences are difficult to define; for instance in areas with an undiscriminating land cover.



**Fig. 5** Incorrect GCP provided on a building. The roof top is clicked in the image, while the same location is provided at ground level in the virtual globe.




**Fig. 6** Due to some inaccurate initial GCP, the mouse pointer on the left indicates a location which is slightly shifted.

In general, few pictures were rejected. Indeed, the automatic checking discussed in section 3.2 prevents the submission of completely incorrect locations.

Some pictures are difficult or even impossible to locate. The most difficult images are the ones having a very rough or incorrect *a priori* location (due to different places having similar names or metadata errors). In this scenario, only a person who recognizes the landscape in the picture can perform the georeferencing. Images that are impossible to georeference are the ones that do not match the settings of our georeferencing algorithm. Currently our implementation of the collinearity equations is not adapted to:

- Vertical aerial images
- Cropped images
- Photomontages

In these situations, the algorithm fails or computes an incorrect location. In the future, our goal is to improve the georeferencing algorithm in order to cope with the first scenario. The cropped images issue can also be solved if the volunteer provides more GCP. However, photomontages will remain an issue due to the fact that they look like real photos and the volunteer can spend a considerable amount of time trying to identify their location. For instance in the published collections some collages of two different pictures had been identified (Figure 7).




**Fig. 7** Photomontage: The foreground and the background pictures were cut from two different pictures and pasted together. This image was georeferenced with GCP marked only in the background (Source: Library of Congress, Views of Switzerland).

### 5.3 Volunteers habits

#### 5.3.1 Area of interest

Our assumption is that volunteers provide the georeferencing of images in their regions of interest. For each volunteer we extracted the exact location of *their* pictures and the date of the georeferencing. Hence we were able to visualize the evolution of the volunteer's area of interest.

In Figure 8 we compared for each volunteer the dispersion (standard deviation) of his five initial images with the dispersion of all images that the same user had georeferenced. This result concerns only the most active volunteers who have provided more than 10 images. We computed the dispersion histogram for every volunteer. The first images (green) generally have a dispersion smaller than 20km and few dispersion values are larger than 20km. It means that generally the five first images provided were located in a area of 20km in diameter. The dispersion of all images is more linear, but small dispersions still occur more often. This implies that the volunteers who continue to provide picture georeferencing move to other regions either because there are no more images available in the initial region of interest or because the users want to explore other areas. Interestingly we also noted that the most active and talented volunteers were able to provide the georeferencing of pictures everywhere in Switzerland as long as the *a priori* location provided by the place name was accurate.



**Fig. 8** The dispersion in [m] is represented on the x-axis. The histogram of the dispersion of the five first images is compared to the dispersion of every image. The dispersion of the first five images is generally smaller which indicates that the volunteers start to work on a specific region and tend to extend their region of interest afterwards.

### 5.3.2 Browser and device

The web-based platform was optimized for the Chrome web browser which provides the best performance for the rendering of the virtual globe. The georeferencing task is only comfortable with a desktop or laptop computer (smartphone and tablet screens are too small to perform the georeferencing task). Indeed 75% of the visitors used a desktop or laptop computer and 44% of the users used the Chrome web-browser. The sessions of the users who used the Chrome web browser also lasted longer. This result suggests that the rendering speed might refrain visitors. In general, the information about the browser and used device will help us to specify priorities regarding the improvement of the platform's user experience.

### 5.3.3 Target audience

Another interesting point is the audience of the prototype. This analysis allows for an optimized targeting of potential volunteers. 75% of the visitors are male and the most represented age group is between 35 and 44 years old which represents 30% of the visitors. A typical visitor is thus mainly in a group which can be described as "technology-friendly". This finding is in contradiction to our initial guess that the platform would mainly attract retired citizen (who could be more interested in visualizing historical landscapes and would have more time to perform the georeferencing task).

## 6 Discussion

### 6.1 *Volunteers' interest for georeferencing*

The launch of this project can be considered as successful. This success was helped by the coverage in regular media such as radio and TV. The media's interest in the project was risen by the media campaign created by two universities at the same time. Furthermore, this project can be considered as relatively visual and easy to explain. Finally, in the context of climate change and urbanization assessment, the project demonstrates the value of historical images as a witness for slow change.

At this stage, we do not exactly know the motivation of our volunteers. Their motivation could be related to the will to help archive managers to gain information about the images. A second possible motivation can be the usage and sharing of their geographic knowledge. Finally, among the volunteers who give us feedbacks, some of them told us that the georeferencing is a way to spend time in a meaningful way.

## ***6.2 Ability of the volunteers to perform the georeferencing task***

The first sets of image collections show that the volunteers are able to provide accurate location of pictures. The visual validation by our trained team is the best way to ensure the quality of the provided georeferencing. Moreover, the validation also helped us to understand the volunteers' common mistakes. In the next version of the platform, the validators will be able to provide the reason of the correction and rejection to the volunteers. In this way, one important goal will be to train the volunteers in order to improve their skills and to reduce the validation workload. We will also provide a tutorial explaining common mistakes. This will help the volunteers to understand how to improve the georeferencing accuracy.

The validation remains both time consuming and expensive and it can be considered as the main limitation of the project. Hence, in the future we aim at implementing a validation by the crowd itself (e.g. similarly to the Wikipedia project) and thereby offering tools to expert volunteers to improve or reject entries.

We also aim at recording the time spent by users to perform the georeferencing. This indicator will help us computing the value of the volunteers' activity (sum of the time spent), but also to understand if the volunteers improve their skills and if modifications of the georeferencing screen ease the georeferencing task.

## ***6.3 Volunteers habits***

The logging of the volunteers' activity is a way of understanding the volunteers' interests and habits. A deep comprehension of the volunteers is a key to keeping them active, and to collect information about the most efficient acquisition channels. In the future we will also analyze other parameters such as time habits, which can indicate if the volunteers mostly work at home or at their work place.

The sum of all information about the volunteers can be used to target specific audiences for instance in a social media campaign.

## **7 Conclusions and perspectives**

Our project is one of the first crowdsourcing project which actively involves volunteers in generating 3D data. Due to the fact that the georeferencing task has been separated into smaller and easier tasks, we can demonstrate that volunteers are able to generate quality 3D data. A major issue is the final validation of the 3D data which is currently manually performed by our team. In other projects such as Wikipedia, the validation by the crowd itself proves to be a working concept. This feature would be a major improvement of our platform which already includes an efficient automated validation system that minimizes the submission of completely incorrect data. The georeferencer is also at an early stage. Our objective is to make it more robust and

available for every place on earth. The semi-automatic 3D georeferencing also has a considerable potential. Indeed, the similarity between pictures can be computed with state of the art computer vision algorithms. Thereby an *a priori* 3D location can be provided that is more accurate than the geocoding of place names that is currently used.

Another current challenge is to keep the platform active. With the documented results, we draw the interest of archivists who can: 1) improve the visibility of their collections and 2) improve the metadata of their photos with the georeferencing. Moreover archivists can be provided with the exact 3D location of the picture, the exact extent of the area visible in the image and the place names that are visible in each image. The crowd can also be asked to correct or improve metadata such as the image title or caption. During the georeferencing task, a volunteer has a focused look at the entire image. He may therefore notice interesting elements that he could share. Hence, in the next stage of the project, we will improve the utility of the platform for archivists and continuously add new collections to keep the crowd active.

In the future, we aim at improving the use of the georeferenced pictures. Indeed, single historical images are already used by geoscientists and land planners to understand the evolution of the landscape. Our goal will be to draw their attention to the project and demonstrate how they can benefit from it. Regarding citizen science projects we have proved that the presented project already solves two problems regarding Volunteered Geographic Information for scientific projects (Ingensand et al., 2016): we have found volunteers and they are able to use the platform. The remaining problem is the data quality in terms of the accuracy of the generated geographic data. The visual alignment of the image with the virtual globe is a good starting point but it is not sufficient for a scientific purpose. We will have to work on two different aspects. First the improvement of the georeferencing quality provided by the users and second, the development of advanced photogrammetric functionalities for scientists who can improve the georeferencing accuracy themselves.

Our project has the goal to create a time machine that visualizes the past derived from historical images. Regarding this goal we imagine several improvements. 1) We can add other layers showing historical data such as maps or aerial and satellite images as base data for the virtual globe. 2) In some famous areas, we notice that historical pictures have a very high time resolution: many pictures with similar viewpoints are shot every year. Such densely photographed regions open opportunities regarding the creation of time lapses. Finally another objective is to provide a comparison of a historical picture with the real world rather than the virtual globe. Current advances in Augmented Reality (AR) technologies open the possibility to use our database in order to compare the reality with the past for instance during a city tour or a hike.

## References

- C. Bozzini, M. Conedera, and P. Krebs. A new monoplotting tool to extract georeferenced vector data and orthorectified raster data from oblique non-metric photographs. *International Journal of Heritage in the Digital Era*, 2012.
- M. Brovelli, M. Minghini, and G. Zamboni. Participatory gis: experimentations for a 3d social virtual globe. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 2013.
- M. Goetz and A. Zipf. The evolution of geo-crowdsourcing: bringing volunteered geographic information to the third dimension. In *Crowdsourcing geographic knowledge*. Springer, 2013.
- W. Hartmann, M. Havlena, and K. Schindler. Towards complete, geo-referenced 3d models from crowd-sourced amateur images. *ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences*, 2016.
- J. Ingensand, S. Composto, M. Nappez, T. Produit, O. Ertz, M. Oberson, and D. Rappo. Challenges in vgi for scientific projects. *PeerJ Preprints*, 2016.
- P. Neis, D. Zielstra, and A. Zipf. Comparison of volunteered geographic information data contributions and community development for selected world regions. *Future Internet*, 2013.
- T. Produit and J. Ingensand. A 3d georeferencer and viewer to relate landscape pictures with vgi. In *AGILE International Conference on Geographic Information, LINK-VGI Workshop*, 2016.
- T. Produit and D. Tuia. An open tool to register landscape oblique images and generate their synthetic model. In *Open Source Geospatial Research and Education Symposium (OGRS)*, 2012.
- H. Sauerermann and C. Franzoni. Crowd science user contribution patterns and their implications. *Proceedings of the National Academy of Sciences*, 2015.
- N. Snavely, S. M. Seitz, and R. Szeliski. Photo tourism: exploring photo collections in 3d. *ACM Transactions on Graphics*, 2006.
- C. Strecha, T. Pylvninen, and P. Fua. Dynamic and scalable large scale image reconstruction. In *Conference on Computer Vision and Pattern Recognition*, 2010.
- M. Uden and A. Zipf. Openbuildingmodels - towards a platform for crowdsourcing virtual 3d cities. In *7th 3D GeoInfo Conference*, 2012.
- University of Southern California. Viewfinder: How to seamlessly flickrize google earth. 2008. URL <http://interactive.usc.edu/projects/viewfinder/>.