

Interactive person path analysis in reconstructed public buildings

Martin Brunnhuber*
VRVis GmbH

Gerd Hesina†
VRVis GmbH

Ropot F. Tobler‡
VRVis GmbH

Stephan Mantler+
VRVis GmbH

Abstract

The behavior of a large number of persons in big buildings is pretty difficult to predict. Long time observations have to be made and many things have to be tested in the real environment, for example how to route people around locked escalators, how to place vending or ticket machines, or how to place various kinds of advertisements. A lot of experience is needed to make the right decisions and mistakes and omissions can easily happen.

Models of buildings and simulations can improve the analysis of such environments where many people move about. A good visualization of human behavior in these models can be a great help in planning and improve the quality of necessary changes. It is also possible to analyze the focus of persons who are wandering around in the building, before performing actual changes and thus judge the impact of various alternatives.

A major problem is that for most buildings no adequate 3D models to perform such visualizations exist. Current building models need modeling experts for manipulation, and current simulation software needs simulation experts for analyzing movement of people within a building.

Ideally, simulated movement of people should be presented in a recognizable 3D model in such a way that it can be easily identified and interpreted without further explanation. It should also be possible to manipulate the building and automatically trigger the changed simulation in the background, without being an expert. This should make it possible to explore alternative planning scenarios and improve the applicability of an envisioned planning tool.

This paper presents a way to create an indoor person behavior visualization shown in its application to railway stations. The models are based on a combination of image based reconstruction and existing models of public buildings. The possibility to easily change the environment on the fly creates many possibilities to analyze alternative places for advertisement, vending and ticket machines. Therefore experiments on the real environment can be avoided to save a lot of time and money and increase customer satisfaction.

CR Categories: I.3.6 [Methodology and Techniques]: Interaction Techniques, I.3.7 [Three-Dimensional Graphics and Realism]: Color, shading, shadowing, and texture,

Keywords: visibility map, person path, public building,

*e-mail: brunn@vrvis.at

†e-mail: hesina@vrvis.at

‡e-mail: rft@vrvis.at

+e-mail: mantler@vrvis.at

1 Introduction

In planning railway stations, many design decisions depend on how people navigate through the building. For example, sign posts and advertisements are best placed in locations where most people may be looking, and vending machines should be close to the main pathways without being obstructive.

The objective of the presented application is to aid this design process by creating indoor 3D models of railway stations and to visualize person path trajectories in the created model. This way it should be possible to navigate the model of a 3D environment and see in a simple way where a large number of people can be found in a building. Another major point of interest is, where simulated people are looking when they are walking through the model.

In this project, 3D models of building interiors are created from arbitrary photos. To improve this data, additional CAD plans are manually merged with the created data. Another source of data are time depending person paths. They are displayed with the reconstructed 3D model in an interactive application, allowing the user to explore the given data. The person paths are animated to see the direction and moving speed of the represented persons. The focus of the visualized persons is shown by projecting color coded visibility points onto the model. The user is also able to change the model interactively by importing small 3D models, which leads to an automatic recalculation of the simulated person paths. Subsequently it is also possible to recalculate the shown zones of visibility in the model.

The given visualization can be possibly used to analyze human behavior based on the person path data and the focused points in the environment. The analysis can be used as a basis for the price of advertising space, for a good way to place vending and ticket machines and to analyze how obstacles change the walking behavior of people in buildings.

In the following section, we will describe the related work and afterwards the data acquisition and processing pipeline. Section 4 will present the rendering application and measures for interactive changes, and Section 5 will show results based on test data from two Austrian railway stations. In Section 6 the results are shown and some descriptions can be found, how a user of the presented application may benefit from the application presented.

2 Related Work

Photogrammetric reconstruction of 3D models has become an intensively studied research area, with many recent advances. For example, Sinha et al. [2008] have proposed an interactive system where users assist a semi-automatic reconstruction process of piecewise planar models from an unordered set of photographs which is quite similar to the system proposed in Debevec et al. [1996]. These systems employ user assistance to directly create planar models from the input images. Fully automatic processes have also been published, for example in Irschara et al. [2007]. The immediate result of these processes is typically a cloud of reconstructed feature points, which then needs to be further

processed to acquire a piecewise planar model. This is complicated by reconstruction errors due to noise, feature point mismatches, and partial occlusions.

The quality of photogrammetric reconstruction also depends on the knowledge of intrinsic camera parameters such as focal length and lens distortion characteristics. These parameters are typically computed from a single image of a known target (the most popular being a simple checkerboard image); various camera models exist, one of the most popular being the one proposed by Tsai [1986]. However, some algorithms avoid this requirement and integrate the estimation of camera calibration parameters into the reconstruction process ([Hartley 1994], [Pollefeys et al. 1996]).

Another approach for architectural reconstruction is based on laser range scanning, as in Wolfart et al. [1999] or Sun et al. [2002]. Again, the final goal of reconstruction is the creation of a smooth, accurate mesh from the point cloud as well as the recovery of texture information where available.

Both of these approaches work well for diffuse surfaces, but in the interior of public buildings such as railway stations a lot of glass panes and reflecting surfaces lead to relatively poor reconstruction results.

In general passenger flow visualization has been based either on depicting individual people, or on accumulated streams that represent large numbers of people. Both approaches have their advantages and drawbacks, depending on whether the behavior of individuals should be analyzed (for example, to study collision avoidance behavior) or the focus is on the general behavior of all passengers. Nakamura et al. [2005] have presented a combination of both, showing overall behavior with accumulated color coded lines and marking individuals of interest (stationary people or collision avoidance) with dots. Katabira et al. [2008] process the available passenger trajectories further to produce probability density maps, which can be interpreted as crowd density. Very high densities therefore represent areas that are prone to congestion. However, the walking speed of passengers, a parameter that is of significance for the presented application, is not encoded in these representations.

3 Overview

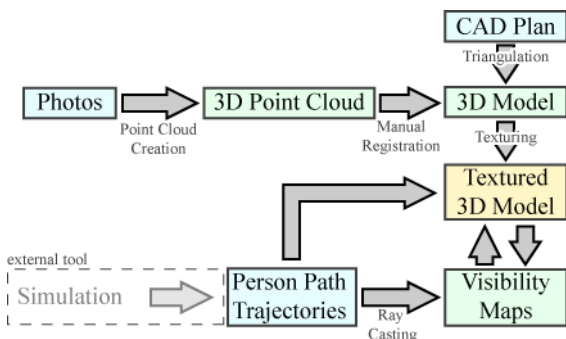


Figure 1: Incoming Data Pipeline

Figure 1 shows the data flow of the implemented application.

The photos of a building are used to create a 3D point cloud. This point cloud is combined with a triangulated CAD Plan resulting in an improved 3D model. After texturing this

model, the person path trajectories from the third party simulation are included, the visibility maps are computed using ray casting, and integrated into the textured 3D model.

4 Data acquisition and processing

The data for the construction of the 3D model is taken from several sources to ensure that the reconstruction is an adequate representation of the real building. A good combination had to be found to get the best solution out of the data. The person paths are taken from third party simulations and are used directly without significant adaptation. The focus of the simulated persons can be shown in order to highlight regions that get a lot of attention by the visitors of a building.

4.1 Picture Acquisition

The images used for the reconstruction of the real buildings as models can be made by standard, handheld cameras without further aids or registration. The only constraint is a certain amount of overlap between pictures to allow for correct registration.

4.2 Point Cloud creation by image processing

The geometry of the interior environment is created by photogrammetric reconstruction from the mentioned images. The used method, dense city modeling, is exactly described by Irschara et al. [2007]. It is necessary to calibrate the images first to rectify them and remove any lens distortion that would introduce reconstruction errors. The user has to take several pictures of calibration markers with the lens used for the construction and use them with a program described by Irschara et al. where the used focal length has to be entered too. With a SIFT-based method pairs of points are found, and with the triangulation of these points the 3d positions of the geometry can be calculated as explained by Irschara et al. [2007]. The manual work for the user is to create the photos and load it into the program with the already camera calibration matrix created with the markers.

The mentioned method was developed for reconstructing facades of houses. Therefore it is limited in its application to interior geometry. Glass and reflective surfaces are especially problematic and do not create useable results. The calculated point clouds include positions far behind the glass surface, where the mirror effect appears. However, some of the data is still usable, since the sections without glass surfaces are created in a correct way and are useful for the 3D model.

4.3 Optional inclusion of CAD plans

In many cases the CAD plans of a railway station are readily available. However, this data is created during the planning of the station and therefore it is usually not an exact representation of reality, due to changes during the construction and maintenance of the buildings. Nevertheless, the main dimensions of the data are precisely available in these plans and can be used for an exact 3D reconstruction. The usage of CAD plans can be complicated because it has to be converted to readable file formats. With the right tools it is possibly made in short time.

4.4 Optional comparison and adjustment of given data

The possibilities of the mentioned algorithm for 3D reconstructions and also the data of a CAD model are limited, as

mentioned. The combination of these two sources leads to improved results, but also to additional, manual work, which might need several days depending on the accuracy level needed. This is done in a common 3D modeling package. To create a good interpretation of the real world the point cloud of the reconstruction program is preferred since the CAD model may not represent the current state of the building. On the other hand, it sometimes shows sections where the reconstruction is wrong and several point clouds can be matched to the CAD model in order to create one complete model. The glass face problem can be avoided this way. The reconstruction with several sources leads to an improved output of the whole reconstruction process.

After the creation of the 3D model it has to be filled with details. Correct sections of the planes in the railway station for textures are taken from the photos which were already rectified and used for the reconstruction. Afterwards they have to be manually put onto the correct parts of the model to create a building model that resembles reality even to non-expert users. It might need some days of work to set the correct textures to the right parts of the model but it is still faster than a complete 3D design process of such a railway station.

4.5 Person path trajectories and lines of sight

The person path trajectories are held in a simple data format. They are defined by a person's position vector at regular time intervals. Additionally a view vector is necessary for every point of the person path.

The currently used third party solution does not include a realistic simulation of view directions. Therefore, in the current application the view vector is simple computed as the difference of two consecutive points of the person path, thus representing the walking direction of the simulated persons. With this simple model it is still possible to demonstrate the visualization without dependency on additional third party data.

4.6 Creation of the basis for visibility maps

The focus of the simulated persons has to be visualized to create good images for effective visual exploration. Our technique is based on the principle of shadow maps as described by Akenine-Möller et al [2008]. The main difference is that instead of illumination, the view focus of the simulated persons is used for updating the "shadow map" (which therefore is called visibility map for clarity, the definition of the word differs from the one used by Goradia et al. [2007]).

In order to produce such "visibility maps", textures have to be placed on all surfaces of a model. Since the input models are of varying quality, it is necessary to implement an algorithm for finding planar parts of the model and compute suitable texture coordinates for these parts. The employed algorithm groups the triangles into classes based on their position and their normal vector with selectable tolerances in both cases. The triangles that are in the same class, and thus approximately on the same plane are then assigned texture coordinates based on their position within a 2D bounding box that encompasses all of them. Each triangle class then receives a single visibility map, thus large numbers of small coplanar triangles are represented by single visibility maps. These empty visibility maps are subsequently filled with data based on view simulation.

4.7 Filling visibility maps by ray shooting

A human view is usually not punctual, but encompasses a certain field of view. To simulate this, ray bundles are used. The number of rays and the view angle to be used per position can be set by the user before the calculation is started. Every ray is shot from the position of the current person path trajectory point. Thus the person path also implicitly defines the eye level for each simulated person. The first ray of the bundle is shot exactly in the direction of the view vector calculated for the person path. The other rays are shot in circles around the center view vector. The first circle uses eight rays and every further circle doubles the number of rays (i.e. 8, 16, 32 rays). If it is not possible to fill a circle with rays with the remaining ones of the number the user entered, the last rays are shot equally distributed on the imagined circle. The deviation angle of every ray in a circle is simply the user defined angle divided through the total number of circles in the current calculation.

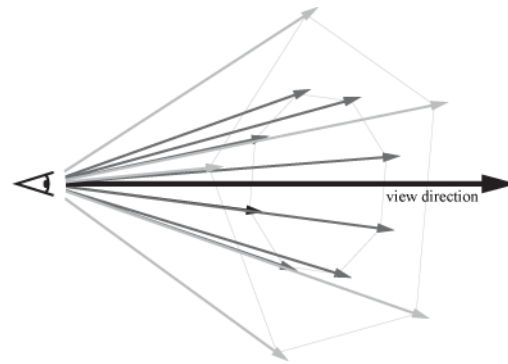


Figure 2: ray bundles as circles around the view direction

In Figure 2 an example is shown, where 14 rays are shot from one point. The first eight rays are shot in the middle of the set angle and the rest of the rays is shot regularly on the set angle away from the view direction.

The ray shooting process is optimized by computing a K-d tree for the 3D model as described by Havran and Bittner [2002]. Every ray on a surface hit is reflected in the corresponding visibility map. The value at this point is increased depending on the priority of the ray. Rays closer to the exact view direction are assigned a higher priority than peripheral ones. The exact value change is also adjusted for the total number of person path points and rays per point to create a solution that is independent of the specific input parameters. This approach simulates the fact that peripheral vision receives less attention than a person's central field of view.

5 Rendering and interaction

The rendering of the application combines simple techniques to get good pictures for exploration of the model. The interaction with the data is pretty important to increase the worth of the person path simulation information.

5.1 Model with visibility map rendering

The model is imported and rendered with ambient illumination only. The large number of textures, caused by the different parts of a railway station, leads to the importance of texture atlases to ensure real-time rendering performance.

The created visibility maps are still not directly usable for the visualization of the focus of the simulated persons. The computed values are changed to colors with a transfer function to be able to represent the degree of interest on the railway station model of the simulated persons.

Since the numerous individual visibility maps are relatively small, they can be efficiently combined into larger maps, i.e. texture atlases, which greatly improves rendering performance by avoiding a large number of context switches.

The visibility maps are used like light maps. The base texture and visibility map of a surface are combined in a pixel shader. This approach allows the visibility map to be changed at runtime without affecting the basic textures. Also, base textures can be reused throughout the model without consuming large amounts of texture memory.

5.2 Display of person paths

Person path trajectories are rendered as line strips. If solid lines were used directly, the directions of moving persons could not be discerned. Therefore, the individual vertices are assigned alternating one-dimensional texture coordinates of zero and one. In the vertex shader, these texture coordinates are increased by a small value for every frame. Using a small, two color texture, this creates moving stripes along the pathways. Additionally, since vertices are located at regular time intervals, the length of individual line sections varies with the simulated person's walking speed, which can thus be directly discerned from the animation.

5.3 Importing objects

Another important point of the presented application is the interaction of the user with the environment. The standard repertoire for moving is not enough to create satisfying solutions for the exploration of the imported data.

Our application supports the interactive placement and removal of additional 3D Models in the interior of the shown building. Also, the base textures of objects can be changed at runtime, for example to explore the placement of signposts and advertisements.

5.4 Recalculation of person paths and visibility maps

The newly imported objects are also integrated into the simulation of person paths and the creation of visibility maps. The third party simulation tool works with a 2D plan of the model environment with information about the position of changing levels on stairways. It is integrated in the tool with an xml-interface where the border information of the newly integrated objects can be found to recalculate the person paths. The newly created paths are directly integrated in the application and replace the old person path trajectory. Afterwards the recalculation of the visibility maps starts in a newly generated thread to keep the standard application running.

6 Results

Two Austrian railway stations, Vienna West and the main station of Graz, were reconstructed for use with the visualization application. The combination of CAD plans and point clouds lead to models which are easily recognizable. Figure 3 shows an

example where the point cloud represented geometry which was not found in the CAD plans. On the left side of the figure you see the point cloud in top view and on the right one the created geometry based on this data.




Figure 3: by point cloud reconstructed geometry

Simulated person path trajectories were also available for the mentioned datasets. The application was tested on a PC with an Intel Core2Quad CPU Q9550 @ 2,83 GHz with 8 GB RAM and a GeForce 280GTX with 1 GB graphics memory on Windows Vista Business with Service Pack 1.

The test cases shown on the figures were always calculated with a maximum deviation angle of five degrees and 25 rays per person path point.

| Rays | Main Mesh | Plain triangle classes | K-d tree | Intersection calculation | Rendering optimization | Texture packing | Main K-d tree | Total time |
|---------------------|-----------|------------------------|----------|--------------------------|------------------------|-----------------|---------------|------------|
| Station Vienna West | | | | | | | | |
| 10 | 0:06 | 0:03 | 0:23 | 0:22 | 6:05 | 1:17 | 0:12 | 8:28 |
| 25 | 0:08 | 0:03 | 0:23 | 0:53 | 6:04 | 1:11 | 0:11 | 8:53 |
| 50 | 0:07 | 0:03 | 0:21 | 1:45 | 6:05 | 1:15 | 0:09 | 9:45 |
| Main Station Graz | | | | | | | | |
| 10 | 0:12 | 0:03 | 0:14 | 0:15 | 3:46 | 1:05 | 0:24 | 5:59 |
| 25 | 0:12 | 0:03 | 0:13 | 0:38 | 3:46 | 1:09 | 0:23 | 6:24 |
| 50 | 0:12 | 0:03 | 0:14 | 1:16 | 3:48 | 1:17 | 0:25 | 7:15 |

Table 1: Calculation time comparison (minutes:seconds)

Several calculation runs for both railway stations were performed and are listed in Table 1 where the time values are shown in minutes. The time was measured for the calculation of the main mesh, for the creation of the planar triangle classes, for the calculation of the K-d intersection tree, the shooting of the rays and setting the visibility map values, the combination of the several parts of the model to a complete geometry to be optimized for the rendering process, the packing of the textures to texture atlases and the last calculation of the K-d tree for the complete model by combining the K-d trees for each model part.

The Vienna West model contains 210 sub objects and consists of 65928 triangles. The person paths used in this model have 123307 entries. The model has more small parts and this leads to the longer combination time of several sub-meshes and the longer lasting texture packing as seen in Table 1.

The second model, the main railway station of Graz, was made of 140 objects and 170540 triangles. There were a total of 24857

points in the person path trajectories of this model. The longer lasting creation of the main mesh is caused by the larger number of triangles.

In both parts it is obvious that only the actual intersection calculation is influenced by the number of rays. The time for this operation is growing linear. The highest amount of time is needed for the optimization of the model for rendering. This is caused by splitting the model for the calculation of the visibility maps and the subsequently necessary recombination of the whole model.




Figure 4: Vienna West station

Figure 4 shows the old version of the railway station Vienna West. The red sections of the visibility points are the regions receiving the most attention. The simulated people look into the direction they walk and therefore they look up when they go upstairs. That is the reason for the high attention regions on the top of the back wall seen here.

A user is able to see in this example that the advertisement in the center is really well placed. The attention of people going upstairs is on this and therefore it could be argued to sell the ad space for a high amount.




Figure 5: Main station Graz

In Figure 5 it is possible to see the model of the main station in Graz. The regions of high interest were positioned near the

elevator and behind the section where it is possible to go down to reach the trains. In this case an ad space might be installed on the center column on the right side of the picture because the attention of many people lies on it. The shop in the rear left part of the image is positioned rather well because the focus of many persons lie on it and therefore it seems that nearly everyone on his way through the railway station is able to see it.

It is possible to re-simulate the person path analysis after setting some obstacles into the station and afterwards recalculate the lines of sight of the simulated persons in this model.




Figure 6: Person path after recalculation

Figure 6 shows the new person path trajectory after a ticket machine object was set as an obstacle into its center. The recalculation of visibility maps was also tested with different numbers of rays. For 10 rays it needed about 1 minute and 45 seconds, for 25 rays it needed 4 minutes and 10 seconds and for 50 rays 8 minutes and 11 seconds. It describes a linear increase but the calculation needs much more time in the complete model than in the prepared one.

This simple example already shows the consequences of the newly set ticket machine. People would have to go around it and are no longer able to go directly on their way to the elevator or stairs behind it. It would be wrong to place this obstacle because it would constrain the smooth person flow.

7 Conclusion and Future Work

We have shown an application to analyze human behavior in an interior environment on the example of railway stations. Besides the person path analysis the possibility to show lines of sight improves the applicability of the information shown. When the model and person paths are created the user can start the calculation of the visibility maps and the visualization with a few mouse clicks to select the model and the person path file.

The basic data for the 3D models were easy to access and therefore it is possible to create more in a pretty short time, but it is important to get some simulation data for a newly created environment. The CAD plan adds more structure to the reconstructed environment created with photos which produces more actual content on the other side.

The main goal of our application was not to create a fully realistic model with many graphic effects to represent the real world, but to display the simulation results in a recognizable environment.

In addition to the primary exploration of the model the user gets more visual information by the possibility to change the environment to find out more about the new situation in person path simulation and visibility. In some examples in the conclusion we showed simple interpretations of the given data to show the worth of ad space, positions of sales rooms and apparent errors in setting obstacles in a railway station.

There are still many possibilities to improve the whole application. The first point is the reconstruction algorithm, which is focused on city modeling and usually not reliable for building interiors. Due to the large number of glass panes in such interiors, a significant amount of manual modeling had to be performed to arrive at adequate models. To overcome this problem, significant research work would have to be done, so that problematic, reflecting and refracting surfaces can be identified and correctly reconstructed.

The visualization of person path can also be improved by several techniques. If there is a large number of paths within a small area it is difficult to identify the direction and density of the main streams. Bundling of them in small regions, and color coding them according to density would be one of the many possible new ways to visualize them. Another way would be to find the major paths and follow them with the camera to see how they go through the building to see more about the behavior of the simulated persons by watching their way.

The basic visibility maps have all the same size. It would be better to find a way that the basic size of the used maps depend on the size of the plane it covers. A solution has to be created to get a good balance between accuracy and performance to create these faces.

The visualization of the regions of interest in the model might also be improved. On several backgrounds the contrast of visibility maps might not be enough to identify them. A method to always find a good balance between them and the background textures might improve the interpretation of the whole application.

Performance can also be possibly improved if the whole process is analyzed to optimize its parts. Also the use of parallel computation might be improved in order to use more cores of current computer systems simultaneously. Especially the recombination of several paths needs a rather high percentage of the whole preprocessing time, as shown in table 1. The recalculation of the visibility maps can also be improved because a large number of recalculated rays would lead to a rather slow recalculation process.

Another research topic will be the automatic creation of textures for the reconstructed point clouds. Since this is the most time consuming part of the model creation, any improvement in this part of the process will significantly improve the turn-around time for creating useable models for our application.

8 Acknowledgement

This work was part of the 3D-Inside research project, which is supported by the Austrian Research Promotion Agency FFG as part of the KIRAS support program for security research under contract nr. xxx.

References

- AKENINE-MÖLLER, T., HAINES, E. AND HOFFMAN, N., 2008, *Real-Time Rendering*, Third Edition, A. K. Peters.
- DEBEVEC, P.E., TAYLER, C.J., AND MALIK, J., 1996, Modeling and rendering architecture from photographs: a hybrid geometry- and image-based approach. In *Proceedings of the 23rd Annual Conference on Computer Graphics and interactive Techniques SIGGRAPH '96*. ACM, NY, 11-20.
- GORADIA, R., KANAKANTI, A., CHANDRAN, S., DATTA, A., 2007, Visibility Map for Global Illumination in Point Clouds. In *Proceedings of GRAPHITE 2007*, Perth,
- HARTLEY, R., 1994, Euclidean reconstruction from uncalibrated views. In J.L.Mundy, A. Zisserman, D. Forsyth, Eds., *Applications of invariance in Computer Vision*, Lecture Notes in Computer Science 825, pp. 237-256/ Springer, 1994.
- HAVRAN, V., BITTNER, J. On improving KD-trees for ray shooting, In *Proc. of WSCG 2002 Conference*
- ISHARA, A., ZACH, C. AND BISCHOF, H., 2007, Towards Wiki-based Dense City Modeling. In *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, 1-8
- KATABIRA, K., SUZUKI, T., ZHAO, H., NAKAGAWA, Y., AND SHIBASAKI, R., 2008, An Analysis of Crowds Flow Characteristics by using Laser Range Scanners. In *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, Vol. 37, Part B5, pp. 953-956. Beijing, 2008.
- NAKAMURA, K., ZHAO, H., SHIBASAKI, R., SAKAMOTO, K., AND SUZUKAWA, N., 2005, Visualizing Passenger Flow in Railway Station using Laser Scanners. In *Proceedings of the 9th Computers in Urban Planning and Urban Management Conference (CUPUM)*, 2005.
- POLLEFEYS, M., VAN GOOL, L., PROESMANS, M., 1996, Euclidean 3D reconstruction from image sequences with variable focal lengths. In B.F. Buxton and R. Cipolla, Ed., *Computer Vision-ECCV'96*, Lecture Notes in Computer Science, pp. 31-42. Springer, 1996.
- SINHA, S.N., STEEDLY, D., SZELISKI, R. AND AGRAWALA, M., 2008, Interactive 3D architectural modeling from unordered photo collections. In *ACM SIGGRAPH Asia 2008 Papers*, J.C. Hart, Ed. ACM, NY, 1-10.
- SUN, Y., PAIK, J.K., KOSCHAN, A. AND ABIDI, M.A., 2002, 3D reconstruction of indoor and outdoor scenes using a mobile range scanner. In *Proceedings of the 16th International Conference on Pattern Recognition 2002*,
- TSAI, R.Y., 1986. An Efficient and Accurate Camera Calibration Technique for 3D Machine Vision. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, Miami Beach, 1986.
- WOLFART, E., SEQUEIRA, V., NG, K., BUTTERFIELD, S., GONCALVES, J.G.M., HOGG, D., 1999. Hybrid approach to the construction of triangulated 3D models of building interiors. In *Proceedings of the ICVS'99 International Conference on Computer Vision Systems*, Springer 1999.