

Mutual Calibration for 3D Thermal Mapping^{*}

Dorit Borrmann^{*} Hassan Afzal^{*} Jan Elseberg^{*}
Andreas Nüchter^{*}

** Jacobs University Bremen gGmbH, Automation Group, Campus Ring
1, D-28759 Bremen, Germany
d.borrmann|h.afzal|j.elseberg|a.nuechter@jacobs-university.de*

Abstract: Three-dimensional digital heat distribution maps are needed to assess the energy efficiency of real estates. The availability of such maps are of great importance for reducing the ecological footprint of houses, buildings, and factories. Designing estates has reached the point, where so-called Passivhaus buildings make extensive use of the intrinsic heat from internal sources such as waste heat from lighting, white goods, and other electrical devices, but without using dedicated heaters. In our approach for creating high-precise heat distribution maps a robot is equipped with a 3D laser scanner, a thermal camera, and a color camera. Data from all the sensors are combined to model the environment precisely. This paper describes the setup of the sensors and the processing of the acquired data, including the automatic co-calibration needed to fulfill this task.

Keywords: 3D thermal mapping, autonomous mobile robots, calibration, thermal imaging, laser scanning

1. INTRODUCTION

Modern design of real estates in the context of green or sustainable buildings makes use of all available heat sources to meet the Passivhaus, the Zero-energy building, or even the Energy-plus building standard. These heat sources include for example electrical equipment and even the body heat from the people or animals inside the building. Analysing energy related issues is a challenging task. Thermal imaging is typically used to record temperatures. However, as the acquired images document the precise temperatures without any spatial dimensions thermography alone does not suffice for quantitative measurements. Precise digital 3D models have been created for years with terrestrial laser scanning technology. Registration algorithms from the geodesy and robotics community combine laser scan data acquired at different positions into complete models of the environment. In this paper, we combine two technologies, thermography and 3D laser scanning. We present a novel calibration board and a set of new algorithms for performing effectively the mutual camera-laser calibration. In addition, we discuss projection, occlusion, and resolution errors and problems occurring from low resolution cameras. We present new methods to cope with these issues.

The acquired data can be used by architects and construction engineers during the construction of new and the modification of existing buildings. It helps to reach energy savings, to ensure reliability of data processing centers, and in many other applications such as digital preservation in the context of cultural heritage. Wardlaw et al. (2010)

^{*} This work was supported by SEE-ERA.NET project ThermalMapper under the project number ERA 14/01. (Project Webpage, 2011)




Fig. 1. Laser scan with reflectance, color and thermal information (best viewed in color).

evaluate the impact of 3D models onto the user's ability to interpret thermal information with the conclusion that a good representation of a 3D model can help the user, e.g., to locate heat sources or to detect thermal bridges. One problem that often remains in the 3D representation is the identification of heat sources. To solve this problem an optical camera is added to our setup.

This paper presents our approach for the autonomous generation of 3D models enhanced with temperature and color information (cf. Fig. 1). The following section presents the state of the art in related research topics. Section 3 presents our approach for the creation of 3D thermal models and section 4 evaluates the procedure. Section 5 concludes the paper.

2. BACKGROUND AND STATE OF THE ART

To assess the energy efficiency of houses thermal cameras are commonly used. These cameras measure temperatures precisely, but return only 2D images of the environment and therefore the loss of energy can only be roughly quantified. Images are projections to 2D. From a sequence of images it is in principle possible to perform a 3D reconstruction. These approaches are called bundle adjustment or structure from motion (SFM), if the focus lies on solving simultaneous localization and mapping (SLAM), i.e., on recovering the 3D structure of the environment and the sensors poses (position and orientation). Since reliable solutions to image based 3D reconstruction for thermal images have not been presented yet, we use the emerging technology of terrestrial laser scanning. Laser scanning methods are well established in the surveying community and in robotics. Terrestrial 3D laser scanning systems yield precise 3D point clouds. Scanning from different poses enables one to digitize a complete indoor environment and to resolve occlusions. Registration algorithms from the geodesy and robotics community are available to automatically align scan views from different poses (Brenner et al., 2008; Surmann et al., 2003)

Related work in inspection robotics includes work in human detection with thermal cameras using temperature signatures (Markov and Birk, 2007). Högner and Stilla (2007) present a modified van as surveying vehicle for acquiring thermal images in urban environments. However, in the focus are outdoor environments and image-based techniques like SFM. Prakash et al. (2006) present stereo imaging using thermal cameras, but focus on small scale applications. Iwaszczuk et al. (2011) suggest an approach to map terrestrial and airborne infrared images onto existing building models. The model is textured by extracting polygonal parts from the image and mapping those onto the model using standardized masked correlation.

Only a little work has been done for combining 3D scanners and thermal cameras. Cabrelles et al. (2009) present a methodology to exhaustively record data related to a World Heritage Monument using terrestrial laser scanning, close range photogrammetry and thermal imagery. They use four different sensors for data acquisition: a reflectorless total station, a terrestrial laser range scanning sensor, a digital photo camera and a thermal camera and use a total of eight natural control points with the help of the total station to relate the geometry between different sensors. Pelagottia et al. (2009) present a first automatic approach for multispectral texture mapping. Their method is based on the extraction of a depth map in the form of an image from the model geometry whose pixels maintain an exact correspondence with the vertices of the 3D model. Afterwards the registration with the chosen texture is done which is based on maximization of mutual information. Alba et al. (2011) combine the data from a terrestrial laser scan with images acquired by a bi-camera system, i.e., a system consisting of an optical camera and a thermal camera. The fusion of the two data sets is achieved by using some control points that are measured with both optical camera and laser scanner. 3D environment mapping using 3D scanners on mobile robots have been subject of past research (Surmann et al., 2003; Nüchter, 2009). Up to




Fig. 2. The robot Irma3D, with a 3D laser scanner, a thermal camera and a webcam. In the background the calibration patterns for the optical camera (left) and the thermal camera (right) are displayed.

our knowledge modeling using 3D scanning and thermal imaging has not been done yet.

3. ADVANCED MUTUAL CALIBRATION BETWEEN THE 3D SENSOR AND THE THERMAL CAMERA

3.1 Experimental Setup and Data Acquisition

The setup for simultaneous acquisition of 3D laser scan data, thermal, and optical images is the robot Irma3D (cf. Fig. 2). Irma3D is built of a Volksbot RT-3 chassis. Its main sensor is a Riegl VZ-400 laser scanner from terrestrial laser scanning. Two cameras are mounted on top of the scanner. The Logitech QuickCam Pro 9000 webcam has a video resolution of 1600×1200 . The optris PI160 thermal camera has an image resolution of 160×120 pixels and a thermal resolution of 0.1°C . It acquires images at a frame rate of 120 Hz and with an accuracy of 2°C . The laser scanner acquires data with a field of view of $360^\circ \times 100^\circ$. To achieve the full horizontal field of view the scanner head rotates around the vertical scanner axis when acquiring the data. We take advantage of this feature when acquiring image data. Since the cameras are mounted on top of the scanner, they are also rotated. We acquire 9 images with each camera during one scanning process to cover the full 360° .

3.2 Data Processing Procedure

After acquiring the 3D data it has to be merged with the color information. This processing consists of four steps that will be explained in this section.

Intrinsic Calibration of Thermal and Optical Cameras. Each sensor perceives the world in its own local coordinate system. To join the perceived information we need the specific parameters of these coordinate systems. Each camera has unique parameters that define how a point (X, Y, Z) in world coordinates is projected onto the image plane. These parameters are calculated through a process known as geometric camera calibration. Given the focal length (f_x, f_y) of the camera and the camera center (c_x, c_y) image coordinates (x, y) are calculated as:

$$\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} f_x & 0 & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} X/Z \\ Y/Z \\ 1 \end{bmatrix}. \quad (1)$$

Given the radial distortion coefficients k_1, k_2, k_3 and the tangential distortion coefficients p_1, p_2 and $r = \sqrt{x^2 + y^2}$ the corrected image points (x_c, y_c) are calculated as

$$\begin{pmatrix} x_c \\ y_c \end{pmatrix} = \begin{pmatrix} x(1 + k_1 r^2 + k_2 r^4 + k_3 r^6) + 2p_1 y + p_2(r^2 + 2x^2) \\ y(1 + k_1 r^2 + k_2 r^4 + k_3 r^6) + p_1(r^2 + 2y^2) + 2p_2 x \end{pmatrix} \quad (2)$$

To determine the parameters of optical cameras chessboard patterns are commonly used because the corners are reliably detectable in the images. A number of images showing a chessboard pattern with known number and size of squares are recorded. In each image the internal corners of the pattern are detected and the known distance between those in world coordinates allows to formulate equations (1) and (2) as a non-linear least squares problem and solve for the calibration parameters (Bradski and Kaehler, 2008). Starting from a guess computed by minimizing an algebraic distance in closed form, we solve the non-linear least square problem by applying the Levenberg-Marquardt algorithm (Zhang, 1999). For solving the calibration problem using the method of Zhang (1999) the calibration pattern has to be planar.

Luhmann et al. (2010) have explored the calibration procedure using different types of thermal cameras. Generally an object with a unique pattern having distinct targets is used which eases labeling and increases accuracy of the calibration process. The points are actively or passively heated. In case of passive heating different material causes the pattern to show up. Luhmann et al. (2010) developed a pattern consisting of targets of self-adhesive foil on an aluminum plate. While the targets emit radiation related to their own temperature the reflective metal surface reflects the cold temperature of space thus leading to an enormous contrast in temperature. Unfortunately this concept is not applicable for the co-calibration of the thermal camera and a laser scanner as it is very difficult to position the board in a way that the sky is reflected without occlusions and the board is completely visible in the laser scan. For low resolution thermal cameras a chessboard pattern is error-prone even after heating it with an infrared lamp. For pixels that cover the edge of the squares the temperature is averaged over the black and white parts thus blurring the edges. Instead a pattern with clearly defined heat sources such as small light bulbs is suggested as it shows up nicely in thermal images. Fig. 2 shows our pattern in the background. It is composed of 30 tiny 12 Volt lamps, each with a glass-bulb diameter of 4 mm. The overall size of the board is 500 mm (width) \times 570 mm (height). Identifying the heat sources in the image enables us to perform intrinsic calibration in the same way as for optical cameras. To detect the light bulbs in the thermal image a thresholding procedure is applied to create a binary image showing regions of high temperature. A further thresholding step discards effectively all regions that are too big or too small. If the remaining number of regions is equal to the number of light bulbs in the pattern the regions are sorted according to the pattern to allow for easy determination of correspondences. To calculate the exact center of the features, the mean is calculated by weighing all the pixels in the region by its temperature value.

Extrinsic Calibration – Cameras and Laser Scanner. After calculating the internal parameters of the cameras

we need to align the camera images with the scanner coordinate system, i.e., extrinsic calibration. The three rotation and three translation parameters are known as the extrinsic camera parameters and define the geometric relation between camera and laser scanner. Once all the points are in the camera coordinate system the projection to the image can be defined up to a factor s using equation (3) (Bradski and Kaehler, 2008):

$$s \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} f_x & 0 & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} r_{11} & r_{12} & r_{13} & t_1 \\ r_{21} & r_{22} & r_{23} & t_2 \\ r_{31} & r_{32} & r_{33} & t_3 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} \quad (3)$$

Suppose there are n images of the calibration pattern and m planar points on the pattern considering the distortions as independent and identically distributed noise then the maximum likelihood estimate of the transformation between the scanner and camera coordinate system is obtained by minimizing

$$\sum_{i=1}^n \sum_{j=1}^m \|\mathbf{p}_{ij} - \hat{\mathbf{p}}(\mathbf{A}, \mathbf{D}, \mathbf{R}_i, \mathbf{t}_i, \mathbf{P}_j)\|^2 \quad (4)$$

where \mathbf{A} is the intrinsic matrix, \mathbf{R}_i the rotation matrix, \mathbf{t}_i the translation vector, and \mathbf{D} the distortion parameters. $\hat{\mathbf{p}}(\mathbf{A}, \mathbf{D}, \mathbf{R}_i, \mathbf{t}_i, \mathbf{P}_j)$ defines the projection of point \mathbf{P}_j in image i , according to equation (3) and (2). This approach assumes that we have a number of points that are identifiable in both the laser scan and the image. For this purpose we attach the calibration pattern onto a board. For the optical camera this is a printed chessboard pattern and for the thermal camera light bulbs arranged in a regular grid pattern. The calibration patterns are depicted in the background of Fig. 2 The position of the points of these patterns are known. Algorithm 1 detects the points in a laser scan.

Algorithm 1 Calibration pattern detection in a laser scan.

Require: point cloud, specification of calibration pattern

- 1: discard points outside the area of the expected board
 - 2: find the most prominent plane using RANSAC (Random SAmple Consensus) (Fischler and Bolles, 1981)
 - 3: project a generated plane model into the center of the detected plane
 - 4: use ICP (Iterative Closest Point) algorithm (Besl and McKay, 1992) to fit the plane model to the data points
 - 5: **if** each point from the plane model has a corresponding point in the point cloud **then**
 - 6: **return** position of the light bulbs according to ICP result
 - 7: **end if**
-

3D to 2D Projection and Color Mapping. During the data acquisition phase laser scans and images are acquired simultaneously. After determining the relations between scanner and cameras in the calibration step this relation is used directly to color the point cloud according to the images.

Scan Registration. Laser scans acquired at different positions are registered into one common coordinate system using 6D SLAM from *3DTK – The 3D Toolkit* (Andreas

Nüchter et al., 2011). The complete model of the environment can be inspected in the viewer from 3DTK enhanced with either reflectance values, thermal data or color from photos (cf. Fig. 1). Switching between the different views enables the user to detect sources of wasted energy and to locate them clearly in the more realistic optical view.

4. EVALUATION OF THE CALIBRATION PROCEDURE

4.1 Extrinsic calibration

The extrinsic calibration is achieved by first finding the best transformation for each pair of image and laser scan given the detected features in image and laser scan and the intrinsic calibration of the camera. The transformation is defined as the translation x, y, z and the rotation $\theta_x, \theta_y, \theta_z$. Especially for low resolution images these n transformations are scattered, as shown exemplarily in Fig. 3(a) for a calibration data set consisting of 31 data pairs. The translation differs by several cm in each coordinate. For the best calibration results the transformation has to be chosen that minimizes the reprojection error for all image pairs. The reprojection error r_k of the transformation calculated from data pair k is defined as

$$r_k = \sum_{i=0}^n \sum_{j=0}^m \frac{|p_{i,j,k} - c_{i,j}|}{m \cdot n}, \quad (5)$$

where n is the number of data pairs, m is the number of feature points on the calibration pattern, $c_{i,j}$ is the j th point on the calibration pattern in the i th thermal image and $p_{i,j,k}$ is the j th point on the calibration pattern in the i th scan when projected onto the image using the k th transformation.

Different methods have been applied to find the best transformation.

mean The average over all translation and rotation vectors is calculated.

emed The transformation is determined as the element-wise median of all transformations, i.e., the median of all values x , and of all values $y, z, \theta_x, \theta_y, \theta_z$, respectively.

med The transformation is chosen that contains the median of the translation vectors. The median translation vector is the vector with the minimum summed distances to all other vectors.

cmean The transformation vectors are filtered based on the distance of the translation vectors to the median. The transformation is then chosen as the mean of the remaining vectors.

min From all transformations the one is chosen, that has the minimum reprojection error. Each translation vector is used to project all feature points onto their corresponding images. The reprojection error for a translation vector is the average distance between feature points in the image and the corresponding projected scan features.

Fig. 3(b) shows the results for the different methods for two calibration data sets. The reprojection error for each of the n image and scan pairs is plotted as crosses. Here the deviation of each hypothesis becomes obvious that is due to the low resolution of the thermal image.




Fig. 4. Comparison of calibration using the second data set. Left: projection using transformation with highest reprojection error. Right: using transformation with lowest reprojection error. (best viewed in color)

The lines represent the different methods. From only two examples it becomes clear that no single method is always advisable to use. The median method, that is the winning method for the first calibration process performs significantly worse than the other methods in the second experiment. However, the best transformation from one data pair outperforms the joint methods in both tries. To yield the best calibration results in general cases, it is suggested to compare the single transformations to the combined methods and chose the best one. We decided against the use of any advanced regression methods for finding the optimal transformation for several reasons. First, the best results for the used methods lie close together. Therefore it is very likely that they actually represent a very good transformation. Second, the results depend heavily on the quality of the input data. As the input data will always be noisy it is expected that the benefits in terms of quality are not worth the extra effort of a reliable optimization method. Last, many optimization methods are prone to be trapped in local optima. This risk is elevated by the highly scattered data available in our calibration data sets.

Fig. 4 shows an example of the projection with different calibration results from the second calibration data set. The left image shows the result using the transformation with the highest reprojection error, the right image the projection using the transformation with the lowest reprojection error. It is obvious that the lower reprojection error also produces visibly better results. The image shows a computer monitor. In the right image the wall behind the right edge of the monitor is assigned warmer temperatures while in the left image the hot spot aligns nicely with the monitor edge in the point cloud.

Projection/Occlusion/Resolution Errors. Due to the different points of view the sensors see different parts of the world. An area that is visible for one sensor might be occluded for the other sensor. When mapping the color information to the point cloud this causes wrong correspondences and therefore faulty colored points. This impact is increased by the low resolution of the thermal camera. With only 120 by 160 pixels per image each pixel corresponds to many 3D points seen by the laser scanner leading to errors at jump edges. Consequently small calibration inaccuracies have a large impact on the results. To solve this problem we take advantage of the fact that if a point belongs to an object there will be more points on that object. We take all points that are projected onto one

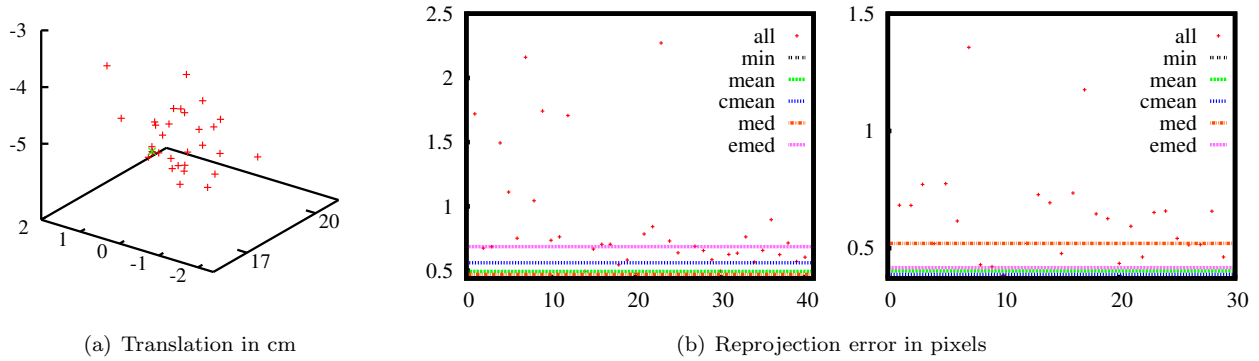


Fig. 3. (a) Distribution of the translation calculated from the sensor data pairs. (b) Reprojection error for extrinsic calibration. The x -axis represents the indices of the data pairs.




Fig. 5. The error correction algorithm. The original image (top) and the corrected point cloud (bottom).

pixel and its neighboring pixels. The points are clustered depending on their distance to the scanner. A heuristic based on distance to the 3D scanner and size of the cluster determines effectively which points are considered and enhanced with color information. This removes also some correct color information but the improvement prevails. Fig. 5 demonstrates this procedure. The top image shows the original colored point cloud. The red boxes mark areas with faulty colored points. These are eliminated with the procedure as seen in the bottom image.

Low Resolution Cameras. Approaches commonly used when combining laser scan data and optical images use expensive high resolution cameras. We used a simple webcam in our experiments. However, the resolution of the thermal camera was only 120×160 pixels. This suggests

that the approach is also applicable for low resolution optical cameras. To verify this we created an image pyramid of the optical images to compare the results achieved with different resolutions, i.e., 1200×1600 , 900×1200 , 600×800 , 300×400 , and 120×160 . Experiments showed that internal and external calibration could successfully be performed for all resolutions given careful positioning of the calibration pattern during calibration. At distances between 50 and 100 cm the distance was large enough to capture the entire board reliably with the scanner and the pattern was still detectable in the camera images. Looking at results achieved with different resolution images shows that the quality decreases and error-proneness increases with decreasing resolution. This effect is diminished by smart application of the error correction algorithm depending on the scenario. Even at resolutions as low as 120×160 an outdoor scene benefits from the added color information. Fig 6 presents results of this experiment. The raw point cloud is presented and the point cloud colored with different resolution images. The bottom left corner of each image shows the scaled input images. The robustness of the approach can be improved by using more data pairs for the calibration process. For the lowest resolution, i.e., 120×160 , only 21 images could be used for the intrinsic calibration, and 15 data pairs for the extrinsic calibration. Given the higher noise level in the feature detection due to the low resolution it is obvious that more data pairs are necessary to minimize the error. Applying the correction algorithm with varying thresholds for decreasing resolution further improves the projection results.

5. CONCLUSIONS

This paper presents a procedure to automatically join sensor information from a laser scanner, a thermal, and a photo camera into one complete 3D thermal model with additional color information. The problems arising in the calibration process are effectively solved by novel methods. To further improve the model in future work we will focus on two parts. First, remaining errors in the projection step caused by occlusions and the low resolution of the thermal camera can be eliminated by taking into account that changes in the depth of the laser scan image and changes in temperature or color are likely to correlate to each other. Second, the data from different views can be used to achieve smoother transitions and to remove errors from the projection step.




Fig. 6. From top left to bottom right: (1) 3D Point cloud without color information. (2) Point cloud coloring with using the full 1600×1200 color image. (3)-(6) Image pyramid with the corresponding point cloud coloring. (best viewed in color)

A video explaining our approach and showing results is given under the following link:

<http://youtu.be/K-1GaGKdNcG>.

REFERENCES

- Alba, M.I., Barazzetti, L., Scaioni, M., Rosina, E., and Previtali, M. (2011). Mapping infrared data on terrestrial laser scanning 3d models of buildings. *Remote Sensing*, 3(9), 1847–1870. doi: 10.3390/rs3091847. URL <http://www.mdpi.com/2072-4292/3/9/1847/>.
- Andreas Nüchter et al. (2011). 3DTK – The 3D Toolkit. <http://slam6d.sourceforge.net/>.
- Besl, P. and McKay, N. (1992). A Method for Registration of 3-D Shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2), 239 – 256.
- Bradski, G. and Kaehler, A. (2008). *Learning OpenCV, Computer Vision with OpenCV library*. O’Reilly Media, first edition.
- Brenner, C., Dold, C., and Ripperda, N. (2008). Coarse orientation of terrestrial laser scans in urban environments. *ISPRS Journal of Photogrammetry and Remote Sensing, Theme Issue Terrestrial Laser Scanning*, 63(1), 4–18.
- Cabrelles, M., Galcera, S., Navarro, S., Lerma, J.L., Akasheh, T., and Haddad, N. (2009). Integration of 3d laser scanning, photogrammetry and thermography to record architectural monuments. In *Proceedings of the 22nd CIPA Symposium*. Kyoto, Japan.
- Fischler, M.A. and Bolles, R.C. (1981). Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography. *Communications of the ACM*, 24, 381 – 395.
- Högner, L. and Stilla, U. (2007). Texture extraction for building models from ir sequences of urban areas. In *Proceedings of the 2007 Urban Remote Sensing Joint event: URBAN / URS*.
- Iwaszczuk, D., Hoegner, L., and Stilla, U. (2011). Matching of 3D building models with IR images for texture extraction. In *Joint Urban Remote Sensing Event*. Munich, Germany.
- Luhmann, T., Piechel, J., Ohm, J., and Roelfs, T. (2010). Geometric calibration of thermographic cameras. In *International Archives of Photogrammetry, Remote Sensing and Spatial Information*, volume 38 part 5. Newcastle upon Tyne, UK. Commission V Symposium.
- Markov, S. and Birk, A. (2007). Autonomous victim detection in rescue operations by modelbased reproduction of 3d scenes from 2d thermal images. In *KI 2007: Advances in Artificial Intelligence, Proceedings of 30th Annual German Conference on AI (KI ’07)*, number 4667 in LNAI. Osnabrück, Germany.
- Nüchter, A. (2009). *3D Robotic Mapping: The Simultaneous Localization and Mapping Problem with Six Degrees of Freedom*. Number 52 in Springer Tracts in Advanced Robotics. Springer.
- Pelagottia, A., Mastio, A.D., Ucheddu, F., and Remondino, F. (2009). Automated multispectral texture mapping of 3d models. In *17th European Signal Processing Conference (EUSIPCO 2009)*. Glasgow, Scotland. URL <http://www.eurasip.org/proceedings/eusipco/eusipco2009/contents/papers/1569192658.pdf>.
- Prakash, S., Pei, L.Y., and Caelli, T. (2006). 3d mapping of surface temperature using thermal stereo. In *Proceedings of International Conference on Control, Automation, Robotics and Vision (ICARCV ’06)*.
- Project Webpage (2011). Project ThermalMapper. <http://www.faculty.jacobs-university.de/anuechter/thermalmapper.html>.
- Surmann, H., Nüchter, A., and Hertzberg, J. (2003). An autonomous mobile robot with a 3D laser range finder for 3D exploration and digitalization of indoor environments. *Journal Robotics and Autonomous Systems (JRAS)*, 45(3–4), 181–198.
- Wardlaw, J., Gryka, M., Wanner, F., Brostow, G., and Kautz, J. (2010). A new approach to thermal imaging visualisation. EngD Group Project, University College London.
- Zhang, Z. (1999). Flexible camera calibration by viewing a plane from unknown orientations. In *in ICCV*, 666–673.