

Generation of Large-Scale 3D Thermograms in Real-Time Using Depth and Infrared Cameras

by S. Schramm*, P. Osterhold*, R. Schmoll* and A. Kroll*

* University of Kassel, Department of Measurement and Control, Mönchebergstraße 7, Kassel, Germany, sebastian.schramm@mrt.uni-kassel.de

Abstract

In recent years, various methods for the generation of 3D thermograms have been developed. A well-established approach is the fusion of data from depth and long-wave infrared (LWIR) cameras. However, these models generated in real-time have the limitation that the model size is limited due to inefficient data storage. Newer algorithms from Computer Vision promise to overcome this limitation. Within this work, one of these 3D modelling algorithms is extended by the overlay of thermal data, which allows the creation of large-scale 3D thermograms. The results will show the advantages over current existing systems.

1. Introduction

3D thermal imaging allows simultaneous analysis of geometry and temperature data. This reduces projection errors and simplifies the interpretation compared to 2D thermograms. The user is supported, for example, in the measurement of objects with many undercuts, where a representation as an overall integrated result is received (instead of a report with many 2D images). With major advances in the data quality of low-cost 3D depth cameras and improvements of the modeling algorithms, linking thermographic data to these models is becoming increasingly popular. Most existing systems that create 3D thermograms from combinations of these sensor data in real-time are limited in the model size. Newer algorithmic approaches try to overcome this limitation (see chapter 2). With the developed 3D thermal imaging system (see chapter 3) three different algorithms are compared experimentally (see chapter 4) and afterwards the most appropriate 3D modelling method is extended by temperature superimposition to enable large-scale 3D thermograms (see chapter 5). The use of the system is demonstrated at example of an industrial molding machine (see chapter 6).

2. Related Work

In recent years, various systems have been developed for the creation of 3D thermograms. Especially due to the different methods of generating the 3D information, the properties and purposes of the systems differ considerably. Beside special applications, where e.g. CAD data is directly fused with thermal images [1] or the combination with color cameras [2], two 3D sensors have been found to be particularly suitable:

Laser scanners (LIDAR) are mainly used for large 3D thermograms in outdoor applications [3, 4]. Due to the time-of-flight working principle, the uncertainty of the measurement does not increase at long distances. However, the resulting models are very sparse because of the limited amount of combined point of view. Due to the duration of a 3D measurement and the usually static position of the measuring system, only low requirements on the model generation speed are necessary.

For the second major class of 3D thermal imaging systems, depth cameras are used for geometry measurements [5–7]. These are low-cost sensors but can usually measure distances up to 3 meters with sufficient accuracy, as the measurement error increases quadratic with the depth. Since these systems are usually hand-held by the operator and moved around the object to be measured, modelling must take place in real-time to visualize the progress of the measurement. These calculations can currently only be performed on a graphics card and the memory of the graphics card is the bottleneck for the size of the measurement objects with the amount of data generated [5]. Commonly used voxel techniques, which divide a measurement volume into a large number of small volume elements, are often memory inefficient, as all volume elements are usually stored in the memory of a graphic card, although only a fraction of the elements contain information of surfaces.

With KinectFusion the first real time geometric modelling system from depth sensor data was invented in 2011 [8]. Based on this work, many algorithmic extensions have been developed in recent years. InfiniTAM [9] and Kintinuous [10] allocate only memory for voxels that are located close to a measured object surface. In addition, parts of the data are transferred from the memory of the graphics card into the RAM. In Kintinuous [10] the measurement volume is shifted based on the movement of the cameras and the model parts, which are no longer in the current measurement volume, are dynamically swapped. In contrast in InfiniTAM [9] the data is grouped into submaps, which are swapped between the memories. Another storage technique is the use surface elements (surfels) instead of volume elements, like in ElasticFusion [11]. These can be stored more easily without surrounding information, so that only surface data needs to be stored. In order to ensure sufficient accuracy for large point clouds, all algorithms also use Loop-Closure (LC) methods [9–11]. Since a small offset in the superposition of individual point clouds occurs at each time step, these errors sum up to

visible displacements in the model over the duration of the measurement. The LC recognizes object features (e.g. after a rotation around the object) and then tries to realign the shifted features.

All three algorithms also use registration algorithms that include color information in addition to the geometry information. This means that features are recognized in the color image and their displacement between the consecutive camera poses is used to overlay the point clouds. This should enable the modelling of areas with few geometric but even color features (e.g. painted flat surfaces).

3. 3D Thermal Measurement System

In order to test the algorithms, a prototypical 3D thermal measurement system was built. For this purpose, a latest generation depth camera (Intel RealSense D415) is rigidly connected to a lightweight long-wave infrared camera (optris PI 450) via a self-designed acrylic base plate. The technical data of the sensors used can be found in Table 1. The mounted 3D thermal imaging system is shown in Fig. 1. A laptop with a high-end graphics card (GeForce GTX 980M, 4 GB VRAM) and Ubuntu 18.04 as operating system is used for the generation of the 3D thermograms.

Table 1: Selected technical specifications of the used sensors.

Item	Intel RealSense D415		optris PI 450
	RGB	Depth (VIS-NIR)	
Image resolution	1920 px × 1080 px	1280 px × 720 px	382 px × 288 px
Field-of-View	69° × 42°	50° × 40°	53° × 38°
Frame rate	60 Hz		80 Hz
Spectral Range	0.4 μm to 0.65 μm	0.4 μm to 0.865 μm	7.5 μm to 13 μm
Typical Uncertainty	-	Depth in z-axis: ≤ 2 %	Temp.: ± 2 K or ± 2 %



Fig. 1. Assembled 3D thermal imaging system, consisting of thermal imaging (top) and depth camera (bottom).

4. Experimental Comparison of 3D Modelling Method

The three algorithms [9–11] are compared for their suitability for the generation of 3D thermograms by means of experimental measurements and their quantitative and qualitative evaluation. Therefore, no integration of thermal data was needed.

4.1. Evaluation criteria

For the evaluation of the three algorithms two main criteria were used:

- Most important is the accuracy of the generated 3D models as it determines the quality of the later 3D thermograms. The accuracy is determined by the quality of the registration procedures and the resolution of the models. In the case of particularly large target objects that are completely circled, the loop closure mentioned in chapter 2 has an additional influence. The accuracy is evaluated quantitatively by comparison with geometric reference measurements and qualitatively by the visual differences of the 3D models.
- The second criterion is the speed of the calculation. This is important because the real-time feedback to the operator allows the operator to intervene on the measurement. Systems as in [5] for example can then indicate areas where a new data acquisition is necessary, e.g. by color-coding. Although the computing times differ considerably between different computers, it can be assumed that the trend of the comparison of the three algorithms is comparable when using a sufficiently fast graphics card.

Since all algorithms were selected on the basis that they promise to capture large 3D models and this has already been shown in the works [9–11], this will not be considered separately.

4.2. Experiments

The algorithms are evaluated by measuring three objects of different geometric properties. As scenes, an infrared calibrator, a laboratory furnace and an office scene (two desks with some objects on it) are considered. In the case of the laboratory furnace, the difficulty is due to the side wall of the furnace, as it contains very few geometrical and visual features. If the algorithms support the export of textured models, the fusion with the visual camera data is shown in the following.

Multiple measurements were taken for each measurement object (three measurements for the IR-calibrator, three different movements around the laboratory furnace with two measurements each and two measurements around the office desk scene). The raw data of the sensors (depth and color image) were stored and the data was played back for each individual algorithm, in such a way that all images of a measurement are processed consecutively. This ensures that the results are comparable despite the system is hand-held.

For the assessment of the modeling accuracy, some sample dimensions of the measurement objects are defined. For the calibrator, the diameter of the calibration surface determined from the point clouds is compared with the specification in the data sheet (measurement error e_{D1} in mm). For the laboratory furnace, the deformation of the model is primarily considered by determining the modeled angle of orthogonal edges to each other (e_{W2} in °). In the office desk scene, five distances (all four desk edges and a distance between two objects on the desk measured with a laser rangefinder) are compared with data from the model (e_{D3} in cm).

4.3. Results

The quantitative results of the accuracy measurement for all three scenes are shown in Table 2. There, both the mean value of the measurement error μ and the variance in the form of the empirical standard deviation σ are given. It is obvious that ElasticFusion provides for the highest accuracy in the experiments.

Table 2: Quantitative Results of the experimental accuracies.

Algorithm	IR-Calibrator		Laboratory Furnace		Office Desk	
	e_{D1}		e_{W2}		e_{D3}	
	μ in mm	σ in mm	μ in °	σ in °	μ in cm	σ in cm
Kintinuous	2.00	3.46	2.39	1.38	4.04	23.37
ElasticFusion	0.67	0.58	1.52	0.93	1.53	16.01
InfiniTAM	1.33	1.53	2.20	0.57	1.84	19.49

This is additionally supported by the qualitative comparison of the results (see Figure 2). In InfiniTAM, a large amount of noise related defects is modelled, since all measured data is inserted into the model without outlier removal. In the case of Kintinuous, wrong registration occurs on the left side of the calibrator and the model is looking partly deformed. Due to the surfel approach, ElasticFusion gives no dense model, but the geometry and color of the calibrator is well reproduced.



Fig. 2. Visual comparison of the modelling quality of the three investigated algorithms with the same input data. From left to right: InfiniTAM (no export of RGB textured models), Kintinuous, ElasticFusion and visual image.

Figure 3 shows a representative curve of the registration time for one measurement of the office scene. It can be seen that InfiniTAM is the fastest algorithm. ElasticFusion and Kintinuous both have about the same ground speed, but with Kintinuous, there are significantly more outliers upwards. The few outliers of ElasticFusion are due to loop closures, which have been integrated into the model at those points in time. With InfiniTAM, a frame rate of 30 Hz can be achieved with the used computing system, with ElasticFusion a rate of 20 Hz can still be observed.

Since ElasticFusion is by far the quantitatively and qualitatively most accurate algorithm and still provides time sufficient performance in terms of calculation speed, it is used as the basis for the 3D thermal measurement system.

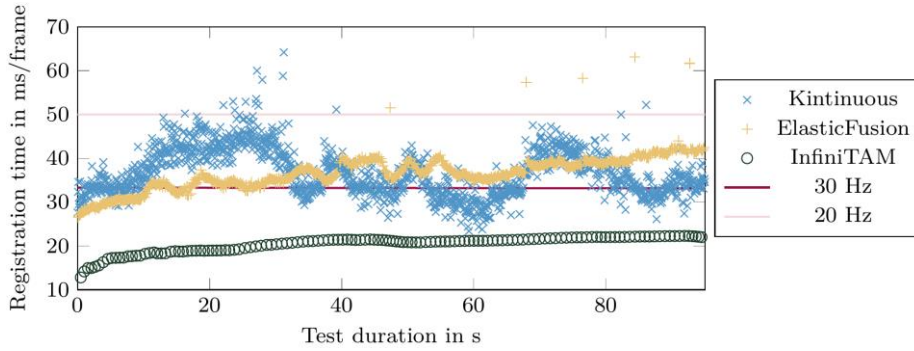


Fig. 3. Comparison of the registration time needed to process one frame. The horizontal lines show the virtual limits for obtaining a frame rate of 20 respectively 30 Hz.

5. Approach for Extended 3D Thermography

ElasticFusion is extended by integrating the temperature data into the 3D model. Therefore, the data processing pipeline of the algorithm is augmented by the input of false color values (see section 5.2.1) of the temperature measurement (with the colormap and the temperature measuring range it is possible to recalculate the temperature values). At first, both cameras must be geometrically calibrated (intrinsically and extrinsically) so that the temperature values can be transformed into the coordinate system of the depth camera.

5.1. Geometrical Calibration

For the data fusion of cameras, they must be geometrically calibrated both intrinsically and extrinsically. In intrinsic calibration, the parameters for mapping object points $[X_c Y_c Z_c]$ (in m) to the image plane $[u v]$ (in px) of a camera are determined:

$$\begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = \frac{1}{z_c} \underbrace{\begin{bmatrix} f_u & 0 & c_u \\ 0 & f_v & c_v \\ 0 & 0 & 1 \end{bmatrix}}_{\mathbf{K}} \begin{bmatrix} X_c \\ Y_c \\ Z_c \end{bmatrix} \quad (1)$$

With the camera matrix \mathbf{K} , consisting of the focal length scaling factors f_u and f_v (in px) and the image centers c_u and c_v (in px). In addition to the parameters derived from the pinhole camera model, distortions caused by the optics must be corrected. For this purpose, the Brown distortion model from [12] is used, whose five parameters $[k_1 k_2 p_1 p_2 k_3]$ are also estimated during the intrinsic calibration.

Within the extrinsic calibration, the transformations (rotation matrix \mathbf{R} and translation vector \mathbf{t}) between two camera coordinate systems c_1 and c_2 are examined:

$$\begin{bmatrix} X_{c2} \\ Y_{c2} \\ Z_{c2} \\ 1 \end{bmatrix} = [\mathbf{R} \quad \mathbf{t}] \begin{bmatrix} X_{c1} \\ Y_{c1} \\ Z_{c1} \\ 1 \end{bmatrix} = \begin{bmatrix} r_{11} & r_{12} & r_{13} & t_1 \\ r_{21} & r_{22} & r_{23} & t_2 \\ r_{31} & r_{32} & r_{33} & t_3 \end{bmatrix} \begin{bmatrix} X_{c1} \\ Y_{c1} \\ Z_{c1} \\ 1 \end{bmatrix} \quad (2)$$

In this case, these are the transformations between the coordinate systems of the depth and the thermal camera. As both are rigidly connected to each other as described in chapter 3, the calibration only needs to be carried out once after mounting the cameras onto the frame.

A target is required for the calibration that has detectable features with a defined distance from each other. This target is placed in different poses (changed translation and rotation) in front of the multi-camera system in a manner that it is always in the field-of-view and depth-of-field of both cameras. Using image processing algorithms, these features are detected in the images and their positions are determined. By means of the optimization algorithm according to [13] the calibration parameters are optimized for the smallest deviations from the known feature distance of a virtual reprojection. For the extrinsic calibration, it is necessary that features are visible in the spectral range of both the long-wave infrared camera and the NIR cameras of the depth sensor. For an overview of the use of different targets for 3D thermal imaging systems, please refer to [14].

Since the depth camera is already geometrically calibrated and rectified by the manufacturer, only the infrared camera is calibrated intrinsically and is then extrinsically calibrated with the left near infrared camera of the depth sensor (since the camera coordinate system of the depth image is placed there). 35 images of a backside heated, chessboard pattern calibration target were taken. The calculation of the calibration parameters is then carried out with the self-developed (based on the Zhang algorithm as implemented in OpenCV) and freely available *MRT Calibration Toolbox*¹. To

¹ <https://github.com/MT-MRT/MRT-Camera-Calibration-Toolbox>

optimize the parameters, 1000 calibration runs are performed with subsets of 10 images each. Then the five images with the highest reprojection error are removed (a high reprojection error usually indicates incorrectly detected features) and then the calibration is performed again.

The results of the calibration are listed in Table 3. The translation between the cameras corresponds to the values resulting from the placement of the cameras on the acrylic frame. Since the remaining reprojection error should not be used as the sole quality measure for calibration, Figure 4 also shows a 3D thermogram of the IR calibrator. There it can be seen that the thermal layer fits exactly at the transition between the calibration surface (15,24 cm diameter) and the front side.

Table 3: Results of the intrinsic calibration of the thermal camera (left) and the extrinsic calibration of the thermal and the depth camera (right).

Thermal camera int.	
Parameter	Value
f_u in px	373.5
f_v in px	373.3
c_u in px	203.6
c_v in px	151.4
k_1 in -	$-4.712 \cdot 10^{-1}$
k_2 in -	$4.125 \cdot 10^{-1}$
p_1 in -	$-1.110 \cdot 10^{-3}$
p_2 in -	$4.692 \cdot 10^{-4}$
k_3 in -	$-1.265 \cdot 10^{-1}$
RMSE in px	0.558

Thermal – Depth ext.	
Parameter	Value
r_{11} in -	0.99998
r_{12} in -	$-1.92 \cdot 10^{-3}$
r_{13} in -	$5.99 \cdot 10^{-3}$
r_{21} in -	$1.73 \cdot 10^{-3}$
r_{22} in -	0.9995
r_{23} in -	$3.19 \cdot 10^{-2}$
r_{31} in -	$6.05 \cdot 10^{-3}$
r_{32} in -	$-3.19 \cdot 10^{-2}$
r_{33} in -	0.9995
t_x in mm	-34.64
t_y in mm	-34.42
t_z in mm	-0.263
RMSE in px	0.418




Fig. 4. Front view (left) and detailed top view (right) of the thermal layer positioning at a 3D thermogram of an IR-calibrator (Fluke 4180).

5.2. Algorithmic Extension

To enable the integration of thermographic data into ElasticFusion, the data structure and program must be extended appropriately. The complete flow of the thermographic data thread running parallel to ElasticFusion is shown in Figure 5 left. After the initialization of the program, the calibration data of the infrared camera is loaded. Afterwards the radiometric object and environment properties (T_{amb} , ϵ , τ) required for the temperature calculation are loaded. The rectification of the camera data is done with a rectification map (see section 5.2.2). After initialization of the camera, the thermal images are received and checked for a non-uniformity correction (NUC) (see section 5.2.3). After the rectification, the data is pushed into a buffer, which ElasticFusion is allowed to access. In ElasticFusion the most recent image of the buffer is loaded, the measured temperature values are assigned to the individual surfels according to the geometric calibration from section 5.1 and integrated into the data structure (see section 5.2.1). Since there is a maximum time offset of approx. 8.3 ms due to the high frame rates of the cameras, no temporal calibration was applied (both cameras have no trigger).

5.2.1. Data structure

After assigning a temperature value to a surfel, it must be stored within the surfel. The data structure given by ElasticFusion offers 4 bytes of free memory, which is used for the temperature data (see Figure 5 right). However, to enable real-time false-color display, the temperature values must be saved as false colors, since it is not possible to convert all points, which are visible in the viewer, in a single time step. Since the temperature values are only saved as false colors,

a color palette with many single color steps should be chosen. The manufacturer's radiometric camera calibration was used to calculate the temperature values. The remaining 8 bit of free memory could be used to save the measuring conditions in the future to enable data selection criteria like in [5].

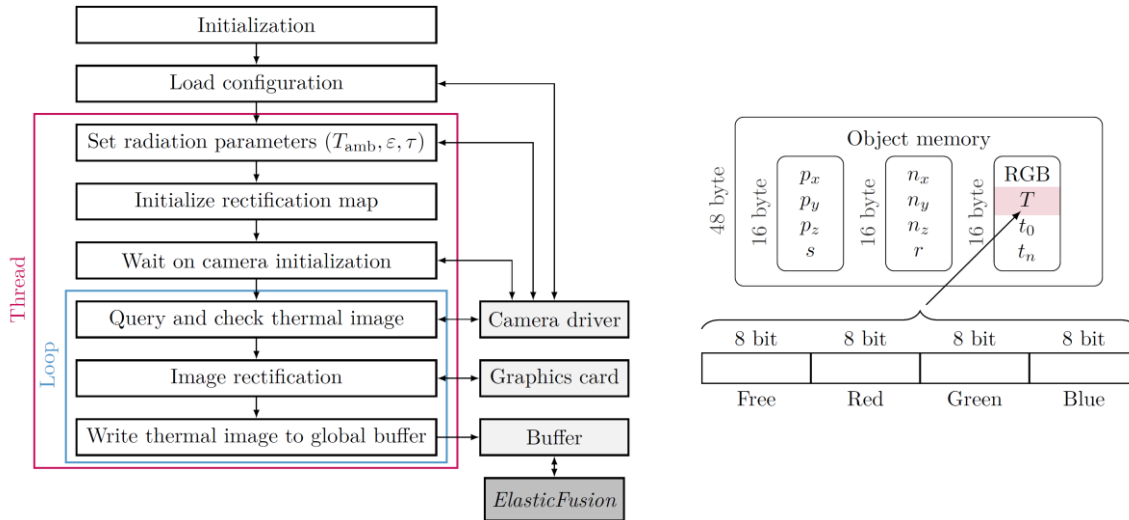


Fig. 5. Flowchart of the thermal image data processing (left) and the new data structure of a surfel (right).

5.2.2. Distortion Calculation

The distortion calculation is not solved for each frame according to Brown's model equations, as this would be too computationally demanding for real-time operation. Instead, the equation is solved once at the beginning for the calibration parameters and the assignment rule is stored in the rectification map. With this look-up table the individual pixels of the incoming thermal images are now assigned to their location in the rectified image. By outsourcing this assignment task to the graphics card of the system additional computing time can be saved. In a test, this improvement resulted in savings of 32.7% computing time compared to an execution on the main processor.

5.2.3. Integration of a NUC

Since previous systems like in [5] could only measure a maximum object size of approx. 1 m^3 (depending on the voxel-size and the graphics memory), the duration of a measurement was usually not more than 20 to 30 seconds. As the object size is now significantly increased, the measurement duration increases accordingly. With previous systems, a NUC was carried out before starting the measurement and the Auto-NUC was switched off during the measurement and modelling process. Due to the use of an uncooled microbolometer camera this will now lead to large errors, especially in scenes with strongly varying temperature radiation. Now the Auto-NUC of the camera can be switched on via input parameters and the frequency can be adapted to the requirements of the scene (signal is sent via the manufacturer's driver). A test has been integrated into the evaluation loop, which checks if a NUC is currently performed and then sends a pause signal to the ElasticFusion modeling process. During the NUC the camera retransmits the last image taken before the NUC with its frame rate, so that the integrated test is based on the calculation of differential images.

6. Industrial Test Object

As a demonstration object, a thermoplastic injection molding machine was examined with the new 3D thermal imaging system. The resulting 3D color model and the 3D thermal image are shown in Fig. 7. The measured object surface is 46 m^2 . Almost a third of the graphics card's memory was filled with surfel information, so that a total area of about 140 m^3 can now be recorded in one measurement. With a system according to [5], the recording of approximately 100 single 3D thermograms and a time-consuming semi-manual merging post-process would have been necessary to build a comparable result. The recording duration was 14 minutes, such that the set Auto-NUC (every 15 seconds) enabled a temperature drift-free recording. The exported point cloud, which contains both thermal and visual information, has a size of 213 MB.

The thermal characteristic of the casting mold is clearly visible, as it was heated to $200 \text{ }^\circ\text{C}$. In this case, a logarithmic temperature scale was chosen to also show the changes in the lower temperature range. Since no analysis of the emissivity was carried out, the measured temperatures are Blackbody temperatures (emissivity was set to 1).

Problems were caused mainly by three various reasons during the measurement. Areas with visual and near-infrared semi-transparent materials, such as the glass pane of the door, are poorly modelled, as the depth camera and the

visual camera can only measure inaccurately there. The registration of the point clouds with the overall model is problematic in those areas, where only few geometric or visual features are available to calculate the poses. The absence of the features leads to the fact that the optimization of the cost functions no longer leads to a unique solution at registration. This can be seen, for example, in the control cabinet at the bottom of the front. Additionally, during the execution of the NUC the geometrical modelling is stopped. The registration of the point clouds becomes more difficult if the operator makes a large movement with the 3D thermography system during NUC, since the underlying iterative closest point (ICP) algorithm is only working for small displacements. In the worst case, this can lead to incorrect registration and thus to the termination of the measurement. Therefore, for correct measurement, the user should not move the system too fast during the NUC (which is audible).




Fig. 6. Visual 3D model (top) and 3D thermogram (bottom) of a thermoplastic injection molding machine with heated mold cavity. In the color model areas of poor modeling quality are shown: a) Semi-transparent materials, b) Parts with few features, c) Small parts that are difficult to resolve, such as cables.

7. Conclusion

As shown in the previous case study, objects with a significantly larger measurement volume (more than two orders of magnitudes) than before can now be thermally measured in three dimensions with a single measurement run using a combination of depth and infrared cameras. This leads to a drastic time saving, which makes potential service and maintenance tasks with such a system economically viable. By testing and selecting a suitable algorithm, geometrically accurate models will be generated. The case study reveals minor methodological constraints (semi-transparent materials, presence of features) that must be considered during the measurement.

8. Summary and Outlook

This work shows the generation of large-scale 3D thermograms by the fusion of depth and infrared camera data. The approach is based on ElasticFusion [11], which performs the task of 3D modelling from the depth camera data and outperforms two other investigated methods [9, 10]. It is demonstrated that large objects can be measured with a high resolution in real-time. In the future, methods for the evaluation of the temperature data measuring conditions from [5] will be used to extend the presented system in order to integrate only appropriate data into the model. Furthermore, work is currently in progress to automatically determine the heat loss from object surface areas and temperatures. For this purpose, it must be possible to select individual model parts and specify their emissivity. The integration of data from an inertial measurement unit could additionally improve the data availability during the execution of a NUC.

REFERENCES

- [1] Sels S., Verspeek S., Ribbens B., Bogaerts B., Vanlanduit S., Penne R. and Steenackers G., "A CAD matching method for 3D thermography of complex objects," *Infrared Physics & Technology*, vol. 99, pp. 152–157, 2019.
- [2] Ham Y. and Golparvar-Fard M., "An automated vision-based method for rapid 3D energy performance modeling of existing buildings using thermal and digital imagery," *Advanced Engineering Informatics*, vol. 27, no. 3, pp. 395-409, 2013.
- [3] Lagüela S., Martínez J., Armesto J. and Arias P., "Energy efficiency studies through 3D laser scanning and thermographic technologies," *Energy and buildings*, vol. 43, no. 6, pp. 1216–1221, 2011.
- [4] Costanzo A., Minasi M., Casula G., Musacchio M. and Buongiorno M. F., "Combined use of terrestrial laser scanning and IR thermography applied to a historical building," *Sensors*, vol. 15, no. 1, pp. 194–213, 2015.
- [5] Ordonez Müller A. O. and Kroll A., "Generating high fidelity 3-D thermograms with a handheld real-time thermal imaging system." *IEEE Sensors Journal*, vol. 17, no. 3, pp. 774–783, 2016.
- [6] Vidas S. and Moghadam P., "HeatWave: A handheld 3D thermography system for energy auditing," *Energy and Buildings*, vol. 66, pp. 445-460, 2013.
- [7] Hellstein P. and Szvedo M., "3D thermography in non-destructive testing of composite structures," *Measurement Science and Technology*, vol. 27, no. 12, ref.no. 124006, 2016.
- [8] Newcombe R. A., Izadi S., Hilliges O., Molyneaux D., Kim D., Davison A. J., Kohli P., Shotton J., Hodges S. and Fitzgibbon A. W., "Kinectfusion: Real-time dense surface mapping and tracking," *ISMAR*, pp. 127–136, 2011.
- [9] Prisacariu V. A., Kähler O., Golodetz S., Sapienza M., Cavallari T., Torr P. H. S. and Murray D. W., "InfiniTAM v3: A Framework for Large-Scale 3D Reconstruction with Loop Closure", arXiv preprint arXiv:1708.00783, 2017.
- [10] Whelan T., Kaess M., Johannsson H., Fallon M., Leonard J. J. and McDonald J., "Real-time large-scale dense RGB-D SLAM with volumetric fusion," *The Int. J. of Robotics Research*, vol. 34, no. 4–5, pp. 598–626, 2015.
- [11] Whelan T., Salas-Moreno R. F., Glocker B., Davison A. J. and Leutenegger S., "ElasticFusion: Real-time dense SLAM and light source estimation," *The Int. J. of Robotics Research*, vol. 35, no. 14, pp. 1697–1716, 2016.
- [12] Kaehler A. and Bradski G., "Learning OpenCV: Computer vision in C++ with the OpenCV library," O'Reilly, 2016.
- [13] Zhang Z., "A flexible new technique for camera calibration," *IEEE Transactions on pattern analysis and machine intelligence*, vol. 22, pp. 1330–1334, 2000.
- [14] Rangel J., Soldan S. and Kroll A., "3D thermal imaging: Fusion of thermography and depth cameras," QIRT, Bordeaux, France, 2014.