



SmokeBot

Mobile Robots with Novel Environmental Sensors
for Inspection of Disaster Sites with Low Visibility

Project start: January 1, 2015

Duration: 3.5 years

Deliverable 3.2

Software Toolkit – Indoor Surface Temperature
Mapping

Due date: month 18 (June 2016)

Lead beneficiary: LUH

Dissemination Level: PUBLIC

Main Authors:

Björn Zeise (LUH)

Version History:

0.1: Initial version, BZ, June 2016

0.2: Minor changes, BZ, June 2016

Contents

Abstract	4
1 Introduction	5
2 Temperature Mapping	5
2.1 Data Acquisition and Verification	6
2.2 Intrinsic and Extrinsic Calibration	6
2.3 Temperature Projection	7
2.4 Disturbance Removal	8
3 Hardware	8
3.1 Sensor Stack	8
3.2 Calibration Targets	9
4 Software Toolkit	10
4.1 Intrinsic Camera Calibration	12
4.2 Extrinsic Calibration between Laser Range Finder and Camera	13
4.2.1 Data Acquisition	13
4.2.2 Data Verification	14
4.2.3 Calibration	14
4.3 Temperature Mapping	15
4.4 Reflection Removal	16
5 Discussion	17
5.1 Temperature Mapping	17
5.2 Reflection Removal	17
6 Conclusion and Future Work	18
References	19
Appendix A: Structure of Calibration Files	20
Appendix B: Temperature Mapping and Reflection Removal (Paper)	22

Abstract

This document is part of the software toolkit published as Deliverable 3.2 within the scope of the Horizon 2020 project SmokeBot. SmokeBot's objective is to improve the application of mobile robots in disaster scenarios with low visibility conditions. A key element of the project is to develop perception algorithms that support the cooperation of humans and machines in search and rescue missions. This is done with the aid of a novel, multimodal sensor unit, which combines radar, gas, thermal and traditional vision sensors.

This deliverable describes the methods used in Task 3.2 ("Mapping of Surface Temperature onto 3D Structures") and explains the usage of the software toolkit. Task 3.2 is part of Work Package 3 ("Thermal Imaging"), which aims at combining temperature data and 3D structures in order to create partial 3D surface temperature maps. The objective of Task 3.2 in particular is to perform a geometrical calibration between a 3D laser range finder and a thermal imaging camera in order to project temperature values onto 3D structures of the environment. Additionally, methods for dealing with thermal disturbances (see also Deliverable 3.1 and [1]) shall be implemented to enable a reliable environmental representation.

1 Introduction

In today's search and rescue applications, thermal imaging cameras are widely used for tasks such as human detection [2] [3], exploration of forest/smoldering fires [4] [5] and detection of potential hazards [6]. In SmokeBot, thermal imaging cameras will especially be used for the latter task.

Connecting measurements of 3D range sensors with 2D images has been accomplished in several works before. In the field of temperature mapping, there are approaches of [7], [8] and [9]. In these works, the most common method of temperature mapping is applied, namely the calculation of intersections between rays pointing from the range sensor to the detected 3D point and the camera image plane. While our work on Task 3.2 is mainly based on this principle, we also introduce new ideas regarding the calibration target and disturbance handling (see also Deliverable 3.1).

The considerations in this document apply to experiments performed with a thermal imaging camera (FLIR A655sc) working in the long-wave infrared spectrum ($7.5\ \mu\text{m}$ to $14\ \mu\text{m}$). Note that maybe not all presented results are applicable to other kinds of thermal imaging systems. The perception of 3D structures in SmokeBot relies on several range sensors with different sensing capabilities under low-visibility conditions. Those sensors are a 3D laser range finder (Velodyne Puck VLP-16), a radar camera and a thermal stereo vision system. The temperature mapping approach presented in this deliverable is meant to be working with all of these sensors. Due to the current stage of the project with neither radar camera nor thermal stereo vision system available, we used laser range finder data to implement and evaluate our approach.

The document is organized as follows:

- First, we explain the general calibration/mapping method used in Task 3.2.
- Afterwards, we describe which hardware was used during our experiments.
- Then, we present the software toolkit consisting of data acquisition, calibration, temperature mapping and reflection removal algorithms.
- In the end, we conclude the deliverable with a short discussion.

2 Temperature Mapping

Connecting 3D information from range sensors with 2D information from cameras demands the knowledge of the spatial setup between the individual sensors. For this purpose, the sensors have to be calibrated using a calibration target that provides common features, which means that it must be possible to detect those features in both the laser scans and the camera images. In our research, we investigated two different calibration targets – a heated calibration trihedron and a heated single plane target.

We divided the temperature mapping toolkit into four main components. First, the calibration data has to be captured which means that matched data pairs (laser scans and thermal images) are written to files. The second component is the actual calibration procedure that finds the geometrical transformation (consisting of a rotation matrix and a translation vector) between

the sensors. We do this by minimizing the error between detected features in the scans and images. The third component uses the estimated transformation and projects the thermal images onto the point cloud. The last component handles thermal disturbances.

2.1 Data Acquisition and Verification

Data acquisition is performed by taking "snapshots" of current data pairs. For this purpose, incoming messages containing laser scans and images are first synchronized using a maximum allowed time difference between the messages' arrival time. After that, features from both the laser scans and thermal images are extracted.

For an individual laser scan, such a "feature" consists of one or more plane equations. Those can be determined using a plane segmentation algorithm. With the help of calibration patterns on the target, it is possible to estimate those plane equations from the thermal images as well. This way, a spatial relation between the sensors' coordinate frames can be established. Only if features in both the laser scan and the corresponding thermal image are found, the data pair is marked as *valid*.

After multiple valid data pairs were recorded, they are saved to the hard disk. Although the acquisition algorithms are designed to create as much valid data pairs as possible, the generation of faulty data pairs cannot be completely excluded. That is why the user should verify the created data pairs manually before proceeding to the calibration step. If no verification is performed, calibration results will be less accurate.

2.2 Intrinsic and Extrinsic Calibration

The intrinsic calibration of the laser range finder is not part of the toolkit. Instead, we use the calibration parameters as provided by the manufacturer. To determine the intrinsic camera parameters, one of the ROS nodes provides an intrinsic calibration mode (see Section 4.1 for more details).

The core of the whole toolkit is the extrinsic calibration between the laser range finder and the thermal imaging camera. Regarding the calibration target, we investigated two different versions: a heated calibration trihedron and a heated single plane. The trihedron calibration procedure is mainly based on the approach presented in [10], where both sensors are rigidly mounted pointing at the calibration target. Then, we first find the individual transformations between the sensors' frames and the calibration target's frame (see Figure 1). After that, the transformation between both sensors is estimated. The resulting transformation parameters are used as initial values for the minimization procedure

$$\underset{R_{lc}, t_{lc}}{\operatorname{argmin}} \sum_{i=1}^I \sum_{j=1}^3 \sum_{k=1}^K \left\| \hat{n}_c^{(i,j)} \left(R_{lc} p_{lc}^{(i,j,k)} + t_{lc} \right) - d_c^{(i,j)} \right\|^2, \quad (1)$$




Figure 1: Model of the (trihedron) calibration setup illustrating the individual coordinate frames. On each plane of the trihedron, there is a circle grid calibration pattern.

where we try to find the rotation matrix \mathbf{R}_{lc} and the translation vector \mathbf{t}_{lc} that satisfy the Hesse normal form of each plane (control variable j). In Equation (1), we match each 3D point $\mathbf{p}_{lc}^{(i,j,k)}$ of a plane detected with the laser range finder into the corresponding plane detected by the thermal imaging camera (defined by $\hat{\mathbf{n}}_c^{(i,j)}$ and $d_c^{(i,j)}$). In total, there are K points per plane which are first transformed from the laser frame to the camera frame. Afterwards, the distance between the transformed point and the plane detected by the camera is minimized in a non-linear least square way using a total sum of I data pairs.

The procedure for the second calibration target (a single calibration plane) is the same as described above except for the number of planes which reduces to one. This simplifies Equation (1) to

$$\underset{\mathbf{R}_{lc}, \mathbf{t}_{lc}}{\operatorname{argmin}} \sum_{i=1}^I \sum_{k=1}^K \left\| \hat{\mathbf{n}}_c^{(i)} (\mathbf{R}_{lc} \mathbf{p}_{lc}^{(i,k)} + \mathbf{t}_{lc}) - d_c^{(i)} \right\|^2. \quad (2)$$

A more detailed description of the calibration procedure is given in [11].

2.3 Temperature Projection

The assignment of temperature values to 3D points of the point cloud is performed using the pinhole camera model and the intrinsic/extrinsic calibration parameters estimated as explained in the previous section. The projection equation is the following:

$$s \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = \mathbf{K} (\mathbf{R}_{lc} \mathbf{p}_l + \mathbf{t}_{lc}). \quad (3)$$

In Equation (3), s is a scaling factor for the 2D image coordinates u and v , \mathbf{K} is the intrinsic camera matrix, \mathbf{R}_{lc} and \mathbf{t}_{lc} define the transformation from the laser frame to the camera frame,




Figure 2: Temperature mapping result using a false-colored point cloud.

and \mathbf{p}_l is a 3D point measured by the laser range finder. An exemplary mapping result can be seen in Figure 2. For better visibility, the points were false-colored, which means that their grayscale intensities were artificially colored using a color map.

2.4 Disturbance Removal

A detailed description of possible misinterpretations in thermal images and how to handle these disturbances is given in Deliverable 3.1. At the present time, the temperature mapping toolkit offers the possibility to remove thermal reflections when looking at a surface whose pose relative to the camera position can be tracked over multiple viewing angles. Therefore, the regarded surface has to be prepared with special markers. The necessity of this will be omitted as soon as a 3D model of the environment (i.e. a spatial 3D map of the environment consisting of multiple registered 3D scans) will be available in a later stage of the project.

3 Hardware

In this section, we describe the general hardware setup used during our investigations. We especially focused on the creation of different calibration targets. The target's features must be detectable in both the laser scans and the thermal images. While it is usually no problem to detect planes in the point clouds, the patterns seen by the thermal imaging camera must show a large intensity contrast in the thermal images. This is achieved by exploiting the different emissivity characteristics of metal (low-emissive) and dielectric (high-emissive) materials.

3.1 Sensor Stack

The sensors are rigidly mounted in a stack during calibration and operation (see Figure 3). If the spatial configuration between them is changed, a re-calibration is needed. In later stages of the project, the sensor stack is planned to be extended with the radar sensor. A first integration approach of this can be seen in Figure 4. The picture was taken at the first review meeting in



Figure 3: Rigidly mounted sensor stack consisting of thermal imaging camera (top) and 3D laser range finder (bottom).




Figure 4: Sensor stack including radar on Taurob Tracker platform during integration week in April 2016.

April 2016 in Vienna, where we presented results on environment perception under heavy smoke conditions.

3.2 Calibration Targets

In order to estimate reliable calibration parameters from only a couple of observations, we decided to use a trihedron as calibration target. This way, we are able to detect three different planes in only one observation. We began our investigations using a heated trihedron as depicted in Figure 5. The trihedron's frame is constructed of aluminum profiles. The planes are made of high-emissive PVC plates. In our current algorithms, the plates' different colors are not used. In an early stage of development, we tried to identify the individual planes by investigating the laser points' remission values. Unfortunately, this turned out to be inapplicable. The circle grid pattern on the individual plates are made of low-emissive




Figure 5: First iteration of the heated calibration trihedron with large circles.




Figure 6: Second iteration of the heated calibration trihedron with smaller circles.




Figure 7: Front side of the single plane calibration target.



Figure 8: Back side of the single plane calibration target.

aluminum paint, which was applied by using a pattern/mask and aluminum-containing spray. The trihedron is heated from behind using a heat fan.

Using the first iteration of the trihedron target, it was very difficult to detect all the circles of the grid pattern. We traced this back to the large circle diameters (huge effect of perspective distortion) in combination with only small gaps between the rows and columns. In the second version of the calibration trihedron, we applied another circle grid pattern to the plates (see Figure 6). This version works well with the algorithms presented in the next sections.

After successfully calibrating the system with the help of a trihedron, we decided to implement a more portable calibration target. Therefore, we applied heat pads on a single aluminum plate (Figure 8). On the other side of the plate, we glued small squares made of polylactide (printed using a 3D printer) forming a checkerboard pattern (Figure 7). In order to generate a large intensity contrast in the thermal images, half of the squares are coated with aluminum paint.

4 Software Toolkit

The software toolkit consists of the following components and functionalities:

- `trihedron_data_acquisition` (ROS node): Records and saves valid data pairs of laser scans and corresponding thermal images using the heated trihedral calibration target.
- `single_plane_data_acquisition` (ROS node): Records and saves valid data pairs of laser scans and corresponding thermal images using the heated plate with checkerboard pattern. Additionally, an intrinsic camera calibration can be performed with this node.
- `trihedron_calibration`: Uses valid trihedron data pairs in order to extrinsically calibrate the sensors according to Section 2.2.
- `single_plane_calibration`: Uses valid single plane data pairs in order to extrinsically calibrate the sensors according to Section 2.2.
- `thermal_mapping` (ROS node): Projects thermal images onto 3D points using the extrinsic calibration parameters determined by one of the calibration methods.

- `reflection_removal` (ROS node): Eliminates moving thermal reflections based on background subtraction.

The components marked as *ROS nodes* can only be run within the ROS framework¹. Besides the *roscore* running, there have to be publishers providing image and laser scan messages. Several parameters used in the ROS-related toolkit components can be adjusted (see Figure 9). The preset default values demonstrated functionality with evaluation datasets.

Several tools read and write files (point clouds, images, calibrations) to the hard disk. Those files are written to folders created relatively to the program's execution path. Care must be taken, especially when switching from one tool to the next one. An overview on the structure of the calibration files (ending with *.yaml*) is given in Appendix A.

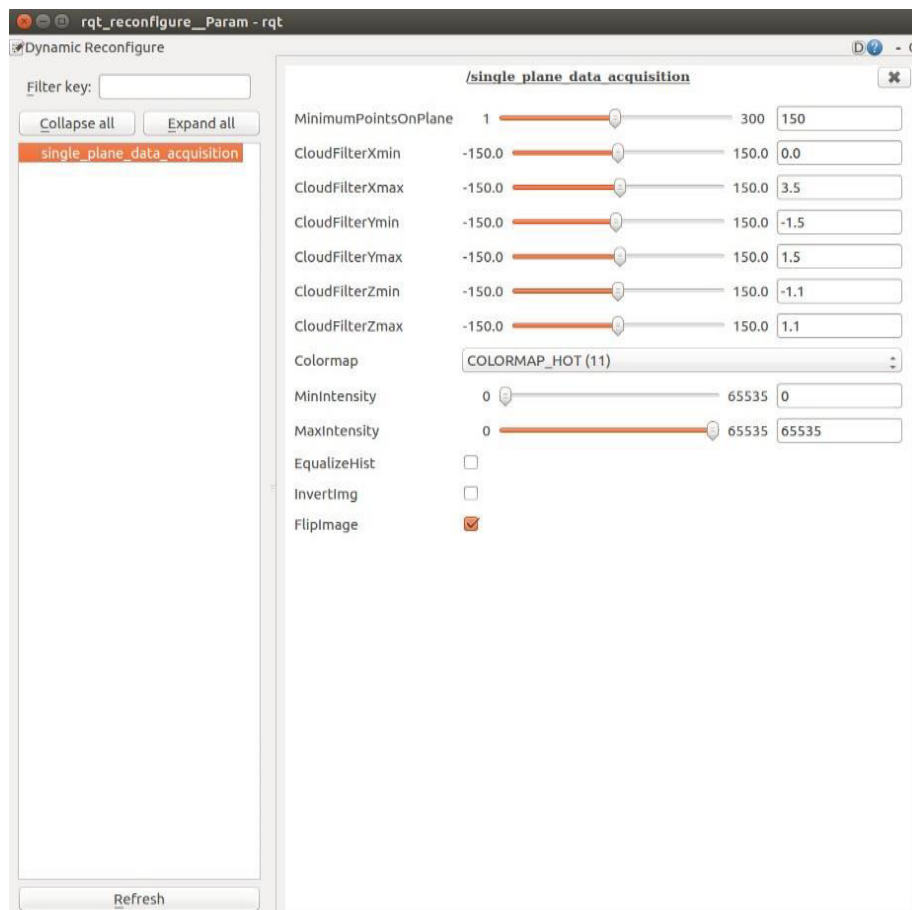


Figure 9: Parameter adjustment using the ROS node `rqt_reconfigure`.

¹ <http://www.ros.org/>

4.1 Intrinsic Camera Calibration

The intrinsic camera calibration, i.e. the estimation of camera matrix K and array of distortion coefficients d , is performed using standard computer vision algorithms from the OpenCV² library. The calibration is done using a checkerboard pattern placed in front of the camera. The algorithm tries to detect the pattern in each thermal image. If it is successfully detected, the image is marked as *valid* and it can be used for calibration. The user recognizes this when the colored pattern is shown in the calibration window (see Figure 10).

Assuming that the camera image and the 3D point cloud are available as ROS topics, the calibration can be carried out as follows:

- 1) Make sure that ROS messages of types *sensor_msgs/Image* and *sensor_msgs/PointCloud2* are published on the topics */camnode/image_raw* and */velodyne_points*, respectively. Remapping of topic names could prevent problems here.
- 2) With the help of a file called “initial_intrinsics.yml” in the node’s execution directory, it is possible to set an initial guess for the intrinsic camera matrix.
- 3) Start the intrinsic calibration node:

```
roslaunch temperature_mapping single_plane_data_acquisition intrinsic
```

- 4) By pressing the button “1”, the user is able to take a thermal image (only if the pattern was detected). Those images are stored in a temporary buffer.
- 5) When the minimum number of images is reached, the user is able to save the images as files on the hard disk by pressing the button “2”. A timestamped folder will be created in the execution directory, containing the subfolders *pattern_points* and *raw_images*.



Figure 10: Intrinsic camera calibration tool.

² <http://opencv.org/>

- 6) With at least the minimum number of images stored in the temporary buffer, the camera calibration can be performed by pressing the button “3”. A file called “thermal_intrinsics.yml” will be created.

4.2 Extrinsic Calibration between Laser Range Finder and Camera

Assuming the intrinsic calibration parameters to be known, the extrinsic calibration can be performed following the steps of data acquisition, data verification and calibration.

4.2.1 Data Acquisition

During data acquisition, planes are segmented from the 3D point cloud provided by the laser range finder. In the case of the trihedral target, the algorithm looks for three planes oriented orthogonally to each other (Figure 12). Assuming the single plane target, the size of the detected plane is taken into account (Figure 14). The implementation was done using the Point Cloud Library³. The plane equations according to the image data (Figure 11 and Figure 13) could be determined using the same method as described in Section 4.1.

Depending on the applied calibration target (trihedron or single plane), data acquisition is done by:

- 1) Make sure that ROS messages of types *sensor_msgs/Image* and *sensor_msgs/PointCloud2* are published on the topics */camnode/image_raw* and */velodyne_points*, respectively. Remapping of topic names could prevent problems here.
- 2) Make sure that the file called “thermal_intrinsics.yml”, which contains the intrinsic camera parameters, is in the node’s execution directory.




Figure 11: Detection of circle grid patterns in the thermal image using the trihedral calibration target.




Figure 12: Trihedron’s planes segmented from the point cloud.

³ <http://pointclouds.org/>

- 3) Start the data acquisition node, depending on the calibration target:

```
roslaunch temperature_mapping trihedron_data_acquisition
```

or

```
roslaunch temperature_mapping single_plane_data_acquisition
```

- 4) By pressing the button “1”, the user is able to record a data pair (only if both the plane(s) and the pattern(s) were detected). Those data pairs are stored in a temporary buffer.
- 5) When the minimum number of data pairs is reached, the user is able to save them as files on the hard disk by pressing the button “2”. A timestamped folder will be created in the execution directory, containing the subfolders *pointclouds*, *pattern_points* and *raw_images*.

4.2.2 Data Verification

Although the acquisition algorithms are designed to find only valid data pairs, it is possible that faulty data pairs are created. In particular, this can be the case for the segmented point clouds. The user has to verify each point cloud visually before continuing with the calibration step. A common way to quickly verify the point clouds is to open the *.pcd* files created inside the *pointclouds* folder during data acquisition with the PCL tool *pcd_viewer*.

4.2.3 Calibration

The calibration programs must be executed from inside the timestamped folder created during data acquisition. It is necessary to copy the file “thermal_intrinsic.yml” created during intrinsic camera calibration into this folder.

The calibration algorithms read the point cloud and pattern point files and match each 3D point into the corresponding plane defined by a plane equation calculated based on image data (pattern points). This matching is expressed as a non-linear least square minimization problem




Figure 13: Detection of checkerboard pattern in the thermal image using one individual plate as calibration target.




Figure 14: Target plane segmented from point cloud during single plane calibration.

that can be solved using the Levenberg-Marquardt algorithm. For our implementation, we used the library *lmfit*⁴ that provides a generic routine for minimizing a sum of squares.

Calibration is performed as follows:

- 1) Make sure that the file called “thermal_intrinsic.yml”, which contains the intrinsic camera parameters, is in the node’s execution directory.
- 2) Start the extrinsic calibration node (“PATH” depending on location of executable):

```
./PATH/trihedron_calibration
```

or

```
./PATH/single_plane_calibration
```

- 3) The algorithm will automatically perform the calibration. A file called “calibration_file.yml” will be written to the execution directory.

4.3 Temperature Mapping

The temperature mapping algorithm takes point clouds and thermal images as input. In combination with the file “calibration_file.yml” created during the previous steps, both a colored point cloud (Figure 15) and an intensity point cloud are published.

The temperature mapping algorithm can be started as follows:

- 1) Make sure that ROS messages of types *sensor_msgs/Image* and *sensor_msgs/PointCloud2* are published on the topics */camnode/image_raw* and */velodyne_points*, respectively. Remapping of topic names could prevent problems here.




Figure 15: Exemplary false-colored point cloud published on topic */pointcloud_rgb* after temperature mapping.

⁴ <http://apps.jcns.fz-juelich.de/doku/sc/lmfit>

- 2) Make sure that the file called “calibration_file.yml”, which contains the intrinsic camera parameters as well as the extrinsic transformation matrix, is in the node’s execution directory.

- 3) Start the thermal mapping node:

```
roslaunch temperature_mapping thermal_mapping
```

- 4) Observe the published topics `/pointcloud_rgb` and `/pointcloud_intensity`.

4.4 Reflection Removal

As explained in Section 2.4, the current version of reflection removal relies on a surface of known structure that was prepared with markers (see Figure 16). The reflection removal algorithm is based on background subtraction using data pairs taken from several different viewing angles. A detailed description of the procedure can be found in [11].

Reflection removal can be performed as follows:

- 1) Make sure that ROS messages of types `sensor_msgs/Image` and `sensor_msgs/PointCloud2` are published on the topics `/camnode/image_raw` and `/velodyne_points`, respectively. Remapping of topic names could prevent problems here.
- 2) Start the reflection removal node:

```
roslaunch temperature_mapping reflection_removal
```




Figure 16: Highly reflective aluminum plate prepared with markers (red) at the corners and high-emissive (non-reflective) shapes of duct tape (white).

5 Discussion

At this point, we would like to discuss some important considerations when making use of the temperature mapping toolkit.

5.1 Temperature Mapping

In order to generate reliable calibration parameters, the camera's focus and/or zoom settings are not allowed to be changed, because otherwise the intrinsic camera matrix and the distortion coefficients would change as well. This implies that the whole sensor setup is calibrated for data acquisition at a specific distance. In our experiments, we calibrated the system at a distance of about two meters. At larger distances, the patterns/plates may be too small to be recognized and therefore have to be enlarged.

Regarding the patterns used for camera calibration, the use of checkerboard squares showed more reliable results than circle grid patterns. For a future version of the trihedral calibration target, this should be taken into account.

We evaluated the resulting point clouds visually and found that sometimes, wrong temperature projections are performed. On the one hand, we trace this back to the parallax effect, which means that individual points can be seen in the point cloud, but cannot be seen in the thermal image and vice versa. On the other hand, the 3D laser scanner used in our experiments provides only 16 vertical scan lines, which makes the resulting point cloud vertically sparse. During our investigations, we intentionally did not use a dense point cloud (e.g. generated by a time-of-flight camera), because neither the radar camera nor the thermal stereo vision will provide such dense point clouds in later stages of the project.

5.2 Reflection Removal

The proper functioning of the reflection removal algorithm strongly depends on the parameters chosen for calculating the average image during background subtraction. Especially the weighting parameters and the reflection threshold influence the removal results. Choosing improper parameters can lead to a blurring effect in the reflective areas. Depending on how the




Figure 17: Original thermal image before reflection removal (false-colored).




Figure 18: Thermal image without reflection (false-colored).

weighting parameters are selected, it may be possible that it takes some calculation runs (callbacks) until a satisfying result is visible in the corrected thermal image.

The current implementation is only able to eliminate moving reflections that seem to change their position on the reflecting surface due to a varying viewing angle of the camera. In a later stage of the project, it may be possible to keep track of the reflections in order to remove them even if the camera does not change its point of view. Exemplary images (before/after reflection removal) can be seen in Figure 17 and Figure 18. In this example, the thermal reflection was realized using a can filled with hot water which was placed in direction of reflection.

6 Conclusion and Future Work

In this document, we described the hardware and software components developed within the scope of Task 3.2 (Mapping of Surface Temperature onto 3D Structures). Additionally, the basics of the calibration method were explained. For each of the software toolkit's components, we gave an overview on how to use it.

During our investigations, we first designed a setup for extrinsically calibrating a sensor stack consisting of a thermal imaging camera and a 3D laser range finder. Therefore, we created two different calibration targets: a heated trihedron as well as a heated single plane target.

Based on the hardware setup, we developed a software toolkit that is able to

- acquire calibration datasets,
- calibrate the sensor stack (intrinsic camera parameters and extrinsic transformation between camera and laser scanner),
- project temperature values onto 3D structures, and
- eliminate moving thermal reflections (currently in specifically prepared datasets only).

Using our calibration approach, we successfully calibrated the system in order to map surface temperatures onto 3D structures. In a later stage of the project, additional 3D range data from the radar sensor and the thermal stereo vision system will be available. Since the laser range finder will not work under low-visibility conditions, it then has to be evaluated whether a direct calibration (radar/thermal stereo → thermal camera) or an indirect calibration (radar/thermal stereo → laser range finder → thermal camera) is more useful for temperature mapping.

Future work regarding WP3 will involve the development of a thermal stereo vision method in order to generate additional 3D information in smoky environments. This data, together with the 3D information provided by the radar camera, will be used to create a 3D environment model, which is necessary to handle the thermal disturbances in a satisfying way. Additional focus will be on hazard detection and prediction (WP4), where temperature mapping will be used to recognize hot spots.

References

- [1] B. Zeise, S. P. Kleinschmidt, and B. Wagner, “Improving the interpretation of thermal images with the aid of emissivity’s angular dependency,” in *Proceedings of the IEEE International Symposium on Safety, Security, and Rescue Robotics*, West Lafayette (IN, USA), 2015, pp. 1–8.
- [2] H. Mano, K. Kon, N. Sato, M. Ito, H. Mizumoto, K. Goto, R. Chatterjee, and F. Matsuno, “Treaded control system for rescue robots in indoor environment,” in *Proceedings of the IEEE International Conference on Robotics and Biomimetics*, Guilin (China), 2008, pp. 1836–1843.
- [3] M. Z. Aziz and B. Mertsching, “Survivor search with autonomous UGVs using multimodal overt attention,” in *Proceedings of IEEE Safety Security and Rescue Robotics*, Bremen (Germany), 2010, pp. 1–6.
- [4] L. Merino, F. Caballero, J. R. Martínez-de Dios, J. Ferruz, and A. Ollero, “A cooperative perception system for multiple UAVs: Application to automatic detection of forest fires,” *J. Field Robot.*, vol. 23, no. 3–4, pp. 165–184, 2006.
- [5] D. W. Casbeer, R. W. Beard, T. W. McLain, S.-M. Li, and R. K. Mehra, “Forest fire monitoring with multiple small UAVs,” in *Proceedings of the American Control Conference*, Portland (OR, USA), 2005, pp. 3530–3535.
- [6] F. Amon, A. Hamins, and J. Rowe, “First responder thermal imaging cameras: establishment of representative performance testing conditions,” in *Proceedings of SPIE - The International Society for Optical Engineering*, Orlando (FL, USA), 2006, p. 62050Y–62050Y–12.
- [7] M. I. Alba, L. Barazzetti, M. Scaioni, E. Rosina, and M. Previtali, “Mapping infrared data on terrestrial laser scanning 3D models of buildings,” *Remote Sens.*, vol. 3, no. 9, pp. 1847–1870, 2011.
- [8] D. Borrmann, J. Elseberg, and A. Nüchter, “Thermal 3D mapping of building façades,” in *Intelligent Autonomous Systems 12*, S. Lee, H. Cho, K.-J. Yoon, and J. Lee, Eds. Berlin et al.: Springer, 2013, pp. 173–182.
- [9] S. Vidas, P. Moghadam, and M. Bosse, “3D thermal mapping of building interiors using an RGB-D and thermal camera,” in *Proceedings of the IEEE International Conference on Robotics and Automation*, Karlsruhe (Germany), 2013, pp. 2311–2318.
- [10] X. Gong, Y. Lin, and J. Liu, “3D LIDAR-camera extrinsic calibration using an arbitrary trihedron,” *Sensors*, vol. 13, no. 2, pp. 1902–1918, 2013.
- [11] B. Zeise and B. Wagner, “Temperature Correction and Reflection Removal in Thermal Images using 3D Temperature Mapping,” in *Proceedings of the 13th International Conference on Informatics in Control, Automation and Robotics*, Lisbon (Portugal), in press.

Appendix A: Structure of Calibration Files

Exemplary *initial_intrinsics.yml* file (input of intrinsic calibration):

```
%YAML:1.0
calibrationDate: "Thu Jun  9 13:24:43 2016\n"
cameraMatrix: !!opencv-matrix
  rows: 3
  cols: 3
  dt: d
  data: [ 770.0,  0.,  320.0,
          0., 770.0,  240.0,
          0.,   0.,   1. ]
```

Exemplary *thermal_intrinsics.yml* file (output of intrinsic calibration; input of extrinsic calibration):

```
%YAML:1.0
calibrationDate: "Mon Jun 13 15:12:33 2016\n"
cameraMatrix: !!opencv-matrix
  rows: 3
  cols: 3
  dt: d
  data: [ 7.9779318641666737e+02, 0., 3.9059562691094516e+02,
          0., 7.9779318641666737e+02, 2.8346012164769610e+02,
          0., 0., 1. ]
distCoeffs: !!opencv-matrix
  rows: 5
  cols: 1
  dt: d
  data: [ -1.1367139173063034e-01, -6.8754577114807203e-01,
          9.9828733379039186e-04, 6.7004617733327334e-03,
          1.8637294421462314e+00 ]
```

Exemplary *calibration_file.yml* file (output of extrinsic calibration; input of thermal mapping):

```
%YAML:1.0
calibrationDate: "Mon Jun 13 17:40:09 2016\n"
cameraMatrix: !!opencv-matrix
  rows: 3
  cols: 3
  dt: d
  data: [ 7.9779318641666737e+02, 0., 3.9059562691094516e+02,
          0., 7.9779318641666737e+02, 2.8346012164769610e+02,
          0., 0., 1. ]
distCoeffs: !!opencv-matrix
  rows: 5
  cols: 1
  dt: d
  data: [ -1.1367139173063034e-01, -6.8754577114807203e-01,
          9.9828733379039186e-04, 6.7004617733327334e-03,
          1.8637294421462314e+00 ]
extrinsicRotTrans: !!opencv-matrix
  rows: 3
  cols: 4
  dt: f
  data: [ -2.02291012e-02, -1.30996799e+00, -3.50462832e-03,
          8.42574081e+01, -4.25581560e-02, 5.14101516e-03,
          -9.01831210e-01, -6.02842426e+00, 9.46532905e-01,
          -3.29444781e-02, -4.46327366e-02, 1.28000809e+02 ]
```

Appendix B: Temperature Mapping and Reflection Removal (Paper)

Temperature Correction and Reflection Removal in Thermal Images using 3D Temperature Mapping

Björn Zeise and Bernardo Wagner

*Institute of Systems Engineering, Real Time Systems Group, Leibniz Universität Hannover,
Appelstr. 9A, D-30167 Hannover, Germany
{zeise, wagner}@rts.uni-hannover.de*

Keywords: Infrared Thermography, 3D Temperature Mapping, Emissivity Estimation, Temperature Correction, Thermal Reflection Removal, Mobile Service Robotics.

Abstract: Many mobile robots nowadays use thermal imaging cameras (TICs) in order to enhance the environment model that is created during exploration tasks. In conventional thermography, thermal images always have to be carefully revised by human operators, which is not practicable in autonomous applications. Unknown surface emissivities are the main source of misinterpretations in thermal images. In this work, we present two methods dealing with these misinterpretations by exploiting the TIC's changing point of view. While the first approach classifies the regarded material in order to estimate improved surface temperature values, the second one is capable of detecting and removing thermal reflections. The spatial relationship between the thermal images and the regarded surface is made by using a rigidly mounted sensor stack consisting of a TIC and a 3D laser range finder, whose extrinsic calibration is described. During evaluation, we demonstrate the functionality of both approaches.

1 INTRODUCTION

In the domain of mobile service robotics, there are plenty of possible use cases for thermal imaging cameras (TICs) – not only in the context of search and rescue (Aziz and Mertsching, 2010), but also for inspection tasks (Vidas et al., 2013) or traffic surveillance (Iwasaki et al., 2013). Connecting 2D thermal images with 3D structural information brings benefits to robotic applications, e.g. when robot operators have to make quick decisions in demanding situations or in the context of self preservation regarding an autonomously acting robot. The projection of thermal images onto 3D structures is called temperature mapping and depicts one of the topics covered in this work.

While the general procedure of temperature mapping is similar to RGB mapping, there are possible sources of misinterpretations in thermal images influencing the temperature mapping results. This work focuses on two of them, namely the temperature misinterpretations arising from unknown emissivity values as well as from thermal reflections (see Figure 1). While a human operator would probably have no problems figuring out that the thermal image shows reflections or incorrectly interpreted temperature val-

ues on metal surfaces, this is a rather hard task for a robot. Misinterpreted environment information can lead to false assessment of the current situation which in turn can endanger the accomplishment of the whole mission.

This work brings the following contributions to the domain of thermography in mobile robotics: First, we describe how to calibrate a TIC and a 3D laser range finder (LRF) using a heated calibration trihedron. We estimate the extrinsic calibration parameters in order to provide 3D points measured by the LRF with temperature information. The second contribution is an extension of our previous work (Zeise et al., 2015) aiming at improving temperature measurements of dielectric and metal surfaces. Using the robot's capability of changing its point of view, we exploit the emissivity's viewing angle dependency in order to classify the surface material and correct the measured temperatures accordingly. In contrast to our previous work, this is performed not only for individual surface points and lines, but also for 2D images of a mixed-material surface. The third contribution is an algorithm that identifies and eliminates moving thermal reflections.

The remainder of the paper is organized as follows: In Section 2, we give an overview on related




Figure 1: Exemplary thermal reflection: (a) False-colored thermal image of a person containing a highly reflective metal surface which reflects the person's thermal radiation and (b) a RGB image of the same scene.

work. In Section 3, we describe how to find the extrinsic calibration parameters needed to perform temperature mapping. Section 4 explains our approaches to correcting temperatures and removing reflections from thermal images. We close with an evaluation of the proposed methods in Section 5, concluding the presented work in Section 6.

2 RELATED WORK

Finding the extrinsic calibration parameters of a LRF and a camera has been investigated in several works. In (Zhang and Pless, 2004), the extrinsic calibration of a camera and a 2D LRF is described. After finding initial guesses for intrinsic and extrinsic parameters with the help of a planar checkerboard pattern, the calibration result is further refined using non-linear minimization. Regarding the calibration between a 3D LRF and a camera, a similar procedure was used in works such as (Pandey et al., 2010) and (Gong et al., 2013). In these approaches, planes are detected in both the laser and camera observations to determine the transformation between the sensor frames.

Temperature mapping is a well-known problem not only in the robotics domain. The most common method is to use a ray tracing algorithm, that calculates the intersections of the laser rays and the camera image plane. This principle has been applied by e.g. (Alba et al., 2011), (Borrmann et al., 2013) and (Vidas et al., 2013) with different kinds of range sensors.

The effect of a TIC's varying viewing angle has been investigated in (Litwa, 2010) and (Muniz et al., 2014) for dielectrics, and in (Iuchi and Furukawa, 2004) for metals. In (Zeise et al., 2015), we recently showed an approach to reducing misinterpretations in thermal images resulting from unknown emissivity values. For this purpose, we exploited the different emissivity characteristics of dielectrics and met-

als. Using the TIC's viewing angle we were able to improve the interpretation of temperature values of low-emissivity surface points and lines.

The removal of thermal reflections can be achieved either hardware or software-based. A hardware-based solution suppressing thermal reflections with the help of an infrared polarizing filter was presented in (Vollmer et al., 2004). This approach showed partially good result, but also many limitations (expensive infrared polarizing filter; compliance to strict spatial measurement setup). The general principle of most software-based methods is to extract background and foreground layers from the image. This relies on the assumption that the input image is a linear superposition of an object layer and one or more reflection layers. In (Planas-Cuchi et al., 2003), a user-assisted, single-image approach is presented. Several authors make use of multiple images of the same scene taken with different camera configurations. Varying the polarizer setting (Farid and Adelson, 1999), changing the focus (Schechner et al., 2000) or applying flashlight (Agrawal et al., 2005) allows to separate the layers. In order to use these methods, two still images of the same scene have to be taken, which is not always possible on a mobile robot. The reflection handling method most relevant to our work is the use of the camera's changing point of view. Approaches to this have been presented in (Criminisi et al., 2005), (Li and Brown, 2013) and (Szeliski et al., 2000).

3 EXTRINSIC LASER-CAMERA CALIBRATION

The main challenge of temperature mapping lies in proper geometric calibration of the sensors. The LRF¹ and the TIC are mounted in a rigid setup, that can be seen in Figure 2(a). Pointing the sensor stack at a calibration target of known dimensions, we first find the transformations between the individual sensors and the calibration target. After that, the transformations between the sensor coordinate frames (i.e. the TIC's and the LRF's coordinate frames) can be calculated. Since we focus on finding the extrinsic calibration parameters of the sensors, we assume the intrinsic calibration parameters for both sensors to be known.

Since the calibration procedure is mainly based on the approach of (Gong et al., 2013), we only give

¹By using the term LRF, we mean a 3D LRF unless otherwise indicated.




Figure 2: Sensors used during calibration: (a) sensor setup consisting of LRF (bottom), TIC (center) and Kinect v2 (top, not used in this work), and (b) calibration setup between the sensor stack and the calibration target.

a short explanation at this place. The general calibration setup is depicted in Figure 2(b). Let us define different coordinate frames (X, Y, Z) by using the indices c, l and t in order to express quantities with respect to the camera, laser and trihedron coordinate systems (see also Figure 3). The extrinsic calibration between the LRF and the TIC is defined by a (3×3) rotation matrix \mathbf{R}_{lc} and a (3×1) translation vector \mathbf{t}_{lc} . This transformation can also be expressed as a transformation from the laser coordinate frame to the trihedron coordinate frame $(\mathbf{R}_{tl}, \mathbf{t}_{tl}) = (\mathbf{R}_{lc}^{-1}, \mathbf{t}_{lc}^{-1})$, followed by a transformation to the camera coordinate frame $(\mathbf{R}_{tc}, \mathbf{t}_{tc})$. A 3D point \mathbf{p}_l with respect to the laser coordinate frame can be projected onto the TIC's image plane using the pinhole camera model:

$$s \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = \mathbf{K}(\mathbf{R}_{lc}\mathbf{p}_l + \mathbf{t}_{lc}), \quad (1)$$

where $[u \ v]^T$ is a 2D point in the image plane scaled by the factor s and \mathbf{K} is the camera matrix containing intrinsic calibration parameters.

In order to find the individual geometric transformations, a common calibration target is needed. We use a heated trihedron whose planes made of PVC are orthogonally oriented. On each of these planes, a pattern of small circles was created using aluminum-containing spray. Observing these three planes with the LRF allows to find the transformation from the trihedron coordinate frame to the laser coordinate frame $(\mathbf{R}_{tl}, \mathbf{t}_{tl})$. In the thermal image, the circles of the patterns let us find the corresponding transformation from the trihedron coordinate frame to the camera coordinate frame $(\mathbf{R}_{tc}, \mathbf{t}_{tc})$.

In the following subsections, we show how to find the initial transformations for the optimization procedure and how to refine the calibration parameters.




Figure 3: Model of the calibration setup: Coordinate frames of the TIC (X_c, Y_c, Z_c) , of the LRF (X_l, Y_l, Z_l) as well as of the calibration trihedron (X_t, Y_t, Z_t) . On each plane of the trihedron, there is a pattern consisting of small circles.

3.1 Finding Initial Transformations

Finding the rotation matrix \mathbf{R}_{tl} and translation vector \mathbf{t}_{tl} that transform 3D points with respect to the trihedron coordinate frame to 3D points with respect to the laser coordinate frame is accomplished using the plane equations for each of the trihedron's planes. In the following explanations, we utilize a superscript notation using i, j and k (or variations) to refer to individual laser-camera data pairs ($i \in 1, 2, \dots, I$), individual planes ($j \in 1, 2, 3$) and individual 3D points ($k \in 1, 2, \dots, K$) lying in a plane. The plane equations are acquired with the help of a RANSAC-based plane extraction method. Each of the planes detected by the LRF is defined by a unit normal vector $\hat{\mathbf{n}}_l^{(j)}$ and a distance $d_l^{(j)}$ from the laser coordinate frame's origin to the plane. The corresponding plane equation is $\mathbf{p}_l^{(j)} \hat{\mathbf{n}}_l^{(j)} - d_l^{(j)} = 0$, where $\mathbf{p}_l^{(j)}$ is an arbitrary point lying in the j th plane.

The individual columns of \mathbf{R}_{tl} depict the planes' unit normal vectors $\hat{\mathbf{n}}_l^{(j)}$ with the result that:

$$\mathbf{R}_{tl} = \begin{bmatrix} \hat{n}_{l,x}^{(1)} & \hat{n}_{l,x}^{(2)} & \hat{n}_{l,x}^{(3)} \\ \hat{n}_{l,y}^{(1)} & \hat{n}_{l,y}^{(2)} & \hat{n}_{l,y}^{(3)} \\ \hat{n}_{l,z}^{(1)} & \hat{n}_{l,z}^{(2)} & \hat{n}_{l,z}^{(3)} \end{bmatrix}. \quad (2)$$

The translation vector \mathbf{t}_{tl} can be calculated multiplying \mathbf{R}_{tl} with the vector of plane distances:

$$\mathbf{t}_{tl} = \mathbf{R}_{tl} \begin{bmatrix} d_l^{(1)} \\ d_l^{(2)} \\ d_l^{(3)} \end{bmatrix}. \quad (3)$$

This vector points from the laser coordinate frame's origin to the trihedron's corner, which is also the origin of the trihedron coordinate frame.

The determination of rotation matrix \mathbf{R}_{tc} and translation vector \mathbf{t}_{tc} that transform 3D points with respect to the trihedron coordinate frame to 3D points with respect to the camera coordinate frame is quite similar to the procedure described above. Using the circle grid patterns, we find the rotation matrices and translation vectors for transforming the individual pattern coordinate frames to the camera coordinate frame. These matrices and vectors can then be related to the trihedron coordinate frame. The columns of \mathbf{R}_{tc} , just like for the laser transformation, are the planes' unit normal vectors, i.e. the third column $\mathbf{r}_3^{(j)}$ of the rotation matrix found for each pattern/plane $j \in \{1, 2, 3\}$:

$$\mathbf{R}_{tc} = \begin{bmatrix} \mathbf{r}_3^{(1)} & \mathbf{r}_3^{(2)} & \mathbf{r}_3^{(3)} \end{bmatrix}. \quad (4)$$

The translation vector \mathbf{t}_{tc} that points from the camera coordinate frame's origin to the trihedron's origin can be calculated with the help of one of the pattern-camera transformations by shifting the pattern's origin to the trihedron's origin.

3.2 Refining Calibration Parameters

A refinement of the calibration parameters can be accomplished by matching the laser points lying in the trihedron's planes into the plane equations estimated based on the thermal image data. Hence, we try to find \mathbf{R}_{tc} and \mathbf{t}_{tc} that satisfy the Hesse normal form of each trihedron's plane. This can be expressed as an optimization problem of the following form:

$$\underset{\mathbf{R}_{tc}, \mathbf{t}_{tc}}{\operatorname{argmin}} \sum_{i=1}^I \sum_{j=1}^3 \sum_{k=1}^K \left\| \hat{\mathbf{n}}_c^{(i,j)} (\mathbf{R}_{tc} \mathbf{p}_l^{(i,j,k)} + \mathbf{t}_{tc}) - d_c^{(i,j)} \right\|^2, \quad (5)$$

where K is the total number of points detected by the LRF in one specific plane j and data pair i , and I is the total number of laser-camera data pairs taken. A solution for this non-linear least squares problem can be found using the Levenberg-Marquardt algorithm.

4 HANDLING OF MISINTERPRETATIONS

Most misinterpretations in thermal images originate from an unknown emissivity ϵ , which is a surface-specific property dependent especially on the type of surface material and the TIC's viewing angle. The viewing angle is the angle between the camera's optical axis and the surface normal of the regarded point.

According to (Martiny et al., 1996), the signal measured by the TIC can be expressed as:

$$S_{sum} = \epsilon S_{obj} + (1 - \epsilon) S_{amb} \quad (6)$$

with

$$S_{obj} = \frac{R}{\exp\left(\frac{B}{T_{obj}}\right) - F} \quad (7)$$

and

$$S_{amb} = \frac{R}{\exp\left(\frac{B}{T_{amb}}\right) - F}, \quad (8)$$

if the the atmospheric amount of radiation power is neglected. In Equation 6, S_{sum} , S_{obj} and S_{amb} represent output signals measured per image pixel. According to the subscripts, they depict the total output signal, the output signal corresponding to the regarded object's radiation power and the output signal corresponding to the reflected ambience's radiation power. In Equations 7 and 8, T_{obj} and T_{amb} are the true object/ambient temperatures, while R , B and F are camera-specific parameters provided by the TIC's manufacturer. From Equation 6, it can be seen that the higher the emissivity value, the smaller the influence of ambient reflections on the thermal image.

Since in mobile robot exploration a preparation of the environment before measurement (e.g. using high emissivity coatings) is not applicable, we developed two methods for improving the interpretation of thermal images regarding unknown surface emissivities and thermal reflections.

4.1 Estimation of Unknown Emissivities

This section refers to our previous work in (Zeise et al., 2015), where we showed the feasibility of temperature correction for surfaces regarded from a known viewing angle. In our approach, we estimated improved surface temperature values of metal and dielectric surface points exploiting the different emissivity characteristics of metals and dielectrics at a varying viewing angle.

In general, the approach divides into two subroutines. The first step is to determine the regarded surface point's material class, i.e. metal or dielectric. This can be done by observing one individual point's output signal at different viewing angles. As we derived from Equation 6, the total output signal depends linearly on the emissivity. Due to the fact that the emissivity in general shows a qualitative behavior as depicted in Figure 4, it is sufficient to investigate the individual pixel's output signal in order to assign one of the properties *metal* or *dielectric* to the point.




Figure 4: Qualitative comparison of metal (blue, dashed line) and dielectric (red, solid line) emissivity characteristics depending on the viewing angle.

The second step is to estimate improved temperature values for each individual point on the surface. This can be done using non-linear least squares optimization minimizing the error between the measured output signals at different viewing angles and theoretically calculated output signals that depend mostly on emissivity. Equations describing emissivity as a function of the viewing angle can be found in the literature, e.g. in (Howell et al., 2011). The minimization is performed as:

$$\operatorname{argmin}_{\mathbf{p}} \sum_{z=1}^Z (S_z^m - S_z^c(\mathbf{p}))^2, \quad (9)$$

where $\mathbf{p} = \{n, k, T_{obj}, T_{amb}\}$ is the set of parameters to be found. In this set of parameters, the refractive index n and the extinction coefficient k , together with the corresponding viewing angle, describe the graph of emissivity ϵ . The signals S_z^m and S_z^c represent measured/calculated values at Z different viewing angles.

In order to use our approach, some assumptions have to be made. On the one hand, both the surface temperature T_{obj} and the ambient temperature T_{amb} are assumed to be unknown but constant. On the other hand, T_{obj} must be higher than T_{amb} . Since the assumption that $T_{amb} = \text{const.}$ is in conflict with the occurrence of thermal reflections, it is necessary to identify and remove them before correcting the surface temperatures.

4.2 Thermal Reflection Removal

Handling of thermal reflections can be accomplished by using background subtraction. The relation of temperature values and 3D information can be used to investigate temperature changes of specific surface points. For our investigations, we assume a static environment, which implies that the temperatures of the regarded surfaces do not change over time. Hence, if an individual surface point's measured temperature does not change when the camera moves, it means

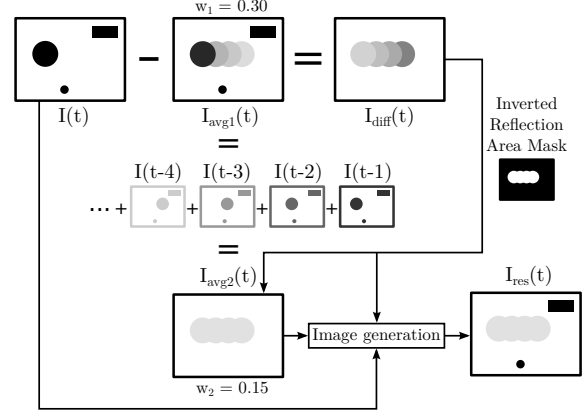


Figure 5: Illustration of the reflection removal procedure: Weighted moving averages of thermal images are used to handle reflections (see the text for a detailed description).

that the temperature is mostly the temperature of the surface itself. In contrast, a changing temperature of one specific surface point implies a superposition of the actual surface temperature and a thermal reflection.

In Figure 5, our approach to thermal reflection removal is illustrated. Let $I(t), I(t-1), \dots, I(t-n)$ be a set of registered thermal images of a static scene taken from different points of view. The weighted moving average $I_{avg}(t)$ at time t can be calculated using the following equation:

$$I_{avg}(t) = wI(t) + (1-w)I_{avg}(t-1), \quad (10)$$

where w is a weighting factor influencing the importance of the latest image $I(t)$. A high value of w means more influence of the current image on the moving average. We use this factor to create two individual moving averages: In the first, $I_{avg1}(t)$, the latest thermal image has a major influence on the average image. Subtracting this average image from the current thermal image gives us $I_{diff}(t)$ that can be regarded as a reflection mask. Every pixel in $I_{diff}(t)$ that is not zero/white represents a potential reflection. The second moving average image, $I_{avg2}(t)$, is created using a smaller weighting factor in combination with the reflection mask. Masking out the non-reflective pixels allows to generate an image that contains only reflective pixels. In contrast to the mask itself and the first average image, the reflective areas in $I_{avg2}(t)$ are very smooth. This smooth image is then used to fill the reflective areas during image generation. Depending on whether the regarded pixel was marked as reflective or not in the reflection mask, the image generation either uses data from $I_{avg2}(t)$ or from $I(t)$ to generate the resulting image $I_{res}(t)$.




Figure 6: Mapping results: False-colored point cloud showing a person standing in front of a cupboard.

5 EVALUATION

The evaluation was performed using the *Velodyne Puck* LRF and the *Flir A655sc* TIC. The LRF provides dense, horizontal scans with a field of view of 360° , while the vertical field of view is only 30° consisting of 16 scan lines. The TIC has a spatial resolution of 640×480 pixels, working in a spectral range between $7.5 \mu\text{m}$ and $14 \mu\text{m}$. The camera's field of view is $45^\circ \times 34^\circ$.


5.1 Laser-Camera Calibration and Temperature Mapping

As stated in Section 3, the intrinsic calibration parameters for the LRF were taken as provided by the manufacturer. The intrinsics of the TIC were estimated using a heated aluminum plate covered with squares made of aluminum and PVC. The parameter estimation was performed using standard computer vision algorithms (Zhang, 2000).

The laser-camera calibration can be assessed looking at the temperature mapping results (Figure 6). Since the vertically sparse point cloud complicates an objective temperature mapping evaluation, we do not provide information on the reprojection error.

5.2 Temperature Correction

The temperature correction algorithm was evaluated using the heated, low-emissive aluminum plate ($56 \text{ cm} \times 70 \text{ cm}$) depicted in Figure 7(a). On the plate, there are several stripes of high-emissive duct tape forming shapes of a triangle, a square and a circle. The aim of this evaluation is to first identify the regarded material class for every individual pixel in the thermal image. After that, we use this information to estimate improved values of the surface points' tem-



(a) Reflective plate

(b) Material classification result

Figure 7: Reflective plate: (a) RGB image of the low-emissive aluminum plate and (b) the result of the material classification procedure (white areas = *metal*, black areas = *dielectric*, gray areas = not specified).

peratures.

We took more than 50 laser-camera data pairs of the heated plate while continuously increasing the sensors' viewing angle. Using the four elliptic tags near the corners of the plate, we were able to register consecutive images. In addition, we used the LRF's 3D point cloud data to calculate every individual pixel's viewing angle.

The results of the material classification are depicted in Figure 7(b). The algorithm distinguishes between metal and dielectric points by adding up the differences between consecutive measurements taken from increasing viewing angles. If the resulting value is greater than zero, the algorithm tags the surface point as *metal*, otherwise as *dielectric*.

The temperature correction is performed for every individual surface point. Depending on the material class, parameters n , k , T_{obj} and T_{amb} are estimated using non-linear optimization. To solve the least squares minimization problem, we make use of the Levenberg-Marquardt algorithm.

The plate's true surface temperature was determined using an additional surface thermometer. On the metal areas, we measured a temperature of about 317 K . The dielectric shapes had a slightly lower temperature of about 316 K . The direction of reflection was covered with a wall having a temperature of about 295 K , ensuring a constant ambient temperature T_{amb} .

The optimization results are depicted in Figure 8. The aluminum temperatures in the original thermal image are obviously misinterpreted. While in the original image (i.e. without taking emissivity into account) the metal surface had an average temperature of about 297 K , our method was able to correct this value to about 305 K for most of the metal surface points. Due to a high emissivity of the dielectric




Figure 8: Temperature correction results (false-colored): (a) Original thermal image with metal areas interpreted as too cold (dark) and (b) the corrected image with improved temperature interpretations.

points, there are almost no misinterpretations in these areas. Hence, the corrected dielectric temperatures are nearly the same as the ones originally measured.

While in the border areas between dielectric and metal points we expected wrong estimations, the results regarding the metal areas are in need of improvement. This can mostly be traced back to the wavelength dependency of the optical constants n and k (as introduced in Section 4.1). Since all measurements taken by the TIC take into account the whole wavelength spectrum between $7.5 \mu\text{m}$ and $14 \mu\text{m}$, n and k , which can vary at different wavelengths, cannot be unambiguously determined. This is also the reason for erroneously estimated high temperature values represented by the bright areas in Figure 8(b).

5.3 Reflection Detection and Removal

To evaluate the capabilities of our reflection removal approach, we used the same aluminum plate as for the temperature correction experiment. As stated before, the bare aluminum surface is highly reflective to thermal radiation. We placed a can filled with hot water (approx. 350 K) in the direction of reflection. We created a dataset, taking thermal images while simultaneously changing the camera's point of view. Applying our algorithm to this dataset leads to the results depicted in Figure 9.

The resulting image shows that the thermal reflection could be most widely removed. During our experiments, we noticed a thin, black borderline in the corrected image at the edge of the aluminum plate. We trace this back to the fact that our algorithm cuts out the plate in order to ensure an accurate image registration. Assuming a complete 3D environment model in our future research, these borderlines will not be present anymore.




Figure 9: Reflection removal result for an exemplary dataset (false-colored): (a) shows the original image containing a reflection of a can (marked by the ellipse), while (b) depicts the processed image without reflection.

A possible shortcoming of our approach is that only moving reflections can be detected at the moment. For later stages of our work, one solution for this could be to keep track of reflections as soon as they were detected.

6 CONCLUSION AND FUTURE WORK

In this work, we faced the problem of handling misinterpretations in thermal images making use of spatial knowledge of the regarded scene. We acquired information of the camera's viewing angle by using a rigidly mounted sensor setup consisting of a TIC and a LRF, which we extrinsically calibrated with the help of a heated calibration trihedron.

We first presented an approach that – using the emissivity's viewing angle dependency – is able to improve temperature measurements of surfaces with unknown emissivity values. Our method showed good results in determining the material class (dielectric/metal) of regarded surface points. While the general functionality of the temperature correction was demonstrated, the results are yet limited due to several unsolved dependencies (e.g. wavelength-dependent optical constants).

The second method presented deals with thermal reflections. Exploiting the TIC's varying viewing angle, our algorithm based on background subtraction was able to detect and remove reflections from the images. Experiments showed that moving thermal reflections could successfully be removed.

Future work will focus on integrating the presented approaches into a 3D simultaneous localization and mapping (SLAM) algorithm. By doing this, our algorithms will benefit from the spatial knowledge of the scene on the one hand, while enhancing the generated environment model with improved tem-

perature mapping on the other hand. Additional effort can be spent on tuning the performance of the temperature correction algorithm, e.g. by using a GPU implementation.

ACKNOWLEDGEMENTS

This work has partly been supported within H2020-ICT by the European Commission under grant agreement number 645101 (SmokeBot).

REFERENCES

- Agrawal, A., Raskar, R., Nayar, S. K., and Li, Y. (2005). Removing photography artifacts using gradient projection and flash-exposure sampling. *ACM Trans. Graph.*, 24(3):828–835.
- 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.
- Aziz, M. Z. and Mertsching, B. (2010). Survivor search with autonomous UGVs using multimodal overt attention. In *IEEE Safety Security and Rescue Robotics*, pages 1–6, Bremen, Germany.
- Bormann, D., Elseberg, J., and Nüchter, A. (2013). Thermal 3D mapping of building façades. In Lee, S., Cho, H., Yoon, K.-J., and Lee, J., editors, *Intelligent Autonomous Systems 12*, number 193 in *Advances in Intelligent Systems and Computing*, pages 173–182. Springer Berlin Heidelberg, Berlin, Heidelberg.
- Criminisi, A., Kang, S. B., Swaminathan, R., Szeliski, R., and Anandan, P. (2005). Extracting layers and analyzing their specular properties using epipolar-plane-image analysis. *Computer Vision and Image Understanding*, 97(1):51–85.
- Farid, H. and Adelson, E. H. (1999). Separating reflections and lighting using independent components analysis. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 1, pages 262–267, Fort Collins, CO, USA.
- Gong, X., Lin, Y., and Liu, J. (2013). 3D LIDAR-camera extrinsic calibration using an arbitrary trihedron. *Sensors*, 13(2):1902–1918.
- Howell, J. R., Siegel, R., and Mengüç, M. P. (2011). *Thermal radiation heat transfer*. CRC Press, Boca Raton, FL, USA, 5th edition.
- Iuchi, T. and Furukawa, T. (2004). Some considerations for a method that simultaneously measures the temperature and emissivity of a metal in a high temperature furnace. *Review of Scientific Instruments*, 75(12):5326–5332.
- Iwasaki, Y., Kawata, S., and Nakamiya, T. (2013). Vehicle detection even in poor visibility conditions using infrared thermal images and its application to road traffic flow monitoring. In Sobh, T. and Elleithy, K., editors, *Emerging Trends in Computing, Informatics, Systems Sciences, and Engineering*, number 151 in *Lecture Notes in Electrical Engineering*, pages 997–1009. Springer New York, New York, NY, USA.
- Li, Y. and Brown, M. S. (2013). Exploiting reflection change for automatic reflection removal. In *IEEE International Conference on Computer Vision*, pages 2432–2439, Sydney, Australia.
- Litwa, M. (2010). Influence of angle of view on temperature measurements using thermovision camera. *IEEE Sensors Journal*, 10(10):1552–1554.
- Martiny, M., Schiele, R., Gritsch, M., Schulz, A., and Wittig, S. (1996). In situ calibration for quantitative infrared thermography. In *International Conference on Quantitative InfraRed Thermography*, pages 3–8, Stuttgart, Germany.
- Muniz, P. R., Cani, S. P. N., and Magalhães, R. d. S. (2014). Influence of field of view of thermal imagers and angle of view on temperature measurements by infrared thermovision. *IEEE Sensors Journal*, 14(3):729–733.
- Pandey, G., McBride, J., Savarese, S., and Eustice, R. (2010). Extrinsic calibration of a 3D laser scanner and an omnidirectional camera. *IFAC Proceedings Volumes*, 43(16):336 – 341.
- Planas-Cuchi, E., Chatris, J. M., López, C., and Arnaldos, J. (2003). Determination of flame emissivity in hydrocarbon pool fires using infrared thermography. *Fire Technology*, 39(3):261–273.
- Schechner, Y. Y., Kiryati, N., and Basri, R. (2000). Separation of transparent layers using focus. *International Journal of Computer Vision*, 39(1):25–39.
- Szeliski, R., Avidan, S., and Anandan, P. (2000). Layer extraction from multiple images containing reflections and transparency. In *IEEE Conference on Computer Vision and Pattern Recognition*, volume 1, pages 246–253, Hilton Head Island, SC, USA.
- Vidas, S., Moghadam, P., and Bosse, M. (2013). 3D thermal mapping of building interiors using an RGB-D and thermal camera. In *IEEE International Conference on Robotics and Automation*, pages 2311–2318, Karlsruhe, Germany.
- Vollmer, M., Henke, S., Karstädt, D., Möllmann, K. P., and Pinno, F. (2004). Identification and suppression of thermal reflections in infrared thermal imaging. In *InfraMation Proceedings*, volume 5, pages 287–298, Las Vegas, NV, USA.
- Zeise, B., Kleinschmidt, S. P., and Wagner, B. (2015). Improving the interpretation of thermal images with the aid of emissivity’s angular dependency. In *IEEE International Symposium on Safety, Security, and Rescue Robotics*, pages 1–8, West Lafayette, IN, USA.
- Zhang, Q. and Pless, R. (2004). Extrinsic calibration of a camera and laser range finder (improves camera calibration). In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 3, pages 2301–2306, Sendai, Japan.
- Zhang, Z. (2000). A flexible new technique for camera calibration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(11):1330–1334.