

Close range 3D thermography: real-time reconstruction of high fidelity 3D thermograms

Antonio Rafael Ordóñez Müller

Schriftenreihe Mess- und Regelungstechnik der Universität Kassel

Band **07**



kassel
university



Schriftenreihe Mess- und Regelungstechnik der Universität Kassel

Band 7 / Vol. 7

Herausgegeben von / Edited by

Univ.-Prof. Dr.-Ing. Andreas Kroll

Close range 3D thermography: real-time reconstruction of high fidelity 3D thermograms

Antonio Rafael Ordóñez Müller



kassel
university 
press

This work has been accepted by the Faculty of Mechanical Engineering of the University of Kassel as a thesis for acquiring the academic degree of Doktor der Ingenieurwissenschaften (Dr.-Ing.).

Supervisor: Univ.-Prof. Dr.-Ing. Andreas Kroll, Universität Kassel

Co-Supervisor: Univ.-Prof. Dr.-Ing. Anton Maas, Universität Kassel

Defense day: 19th April 2018

Bibliographic information published by Deutsche Nationalbibliothek
The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie;
detailed bibliographic data is available in the Internet at <http://dnb.dnb.de>.

Zugl.: Kassel, Univ., Diss. 2018
ISBN 978-3-7376-0624-0 (print)
ISBN 978-3-7376-0625-7 (e-book)
DOI: <http://dx.medra.org/10.19211/KUP9783737606257>
URN: <https://nbn-resolving.org/urn:nbn:de:0002-406258>

© 2018, kassel university press GmbH, Kassel
www.upress.uni-kassel.de

Printed in Germany

Abstract

Infrared thermography is a powerful technique that enables the non-contact measurement of an object's surface temperature and presents the results in form of thermal images. The analysis of these images provides valuable information about an object's thermal behavior and allows in many cases the diagnostic of abnormal conditions. However, the fidelity with which thermal images depict the temperature distribution of an object's surface is limited because the position of the camera with respect to the surface affects the accuracy of the measurements. 3D thermography offers the possibility to overcome this and other limitations that affect conventional 2D thermography but most 3D thermographic systems developed so far generate 3D thermograms from a single perspective or from few noncontiguous points of view and do not operate in real time. As a result, the 3D thermograms they generate do not offer much advantage over conventional thermal images. But reductions in cost, size and weight of thermographic cameras, the recent introduction to the market of depth cameras that are also small, low-cost and light and the significant increase of computational power of commercial computers over the last years have unlocked the possibility to implement affordable handheld 3D thermal imaging systems that can be easily maneuvered around an object and that can generate high-fidelity 3D thermograms in real time.

This thesis exploits these technological advancements and explores the multiple aspects involved in the real-time generation of high-fidelity 3D thermograms at close range using a handheld 3D thermal imaging system. It provides a structure for the knowledge accumulated in the field of 3D thermography, offers new insight into the radiometric, geometric and temporal calibration of the sensors and makes various contributions to improve the design of 3D thermographic systems, the fidelity of the 3D thermograms and the interaction with the user. In regards to the radiometric calibration, it shows through a series of original experiments how the thermal background radiation affects the accuracy of the temperature measurements and provides a plausible explanation for this poorly documented behavior. On the subject of geometric calibration, it presents and analyzes the results obtained with an improved calibration target and demonstrates that the orientation of the thermographic camera cannot be estimated with enough accuracy

due of the poor spatial resolution of the thermal images and that manual fine-tuning may be necessary. Concerning the temporal calibration, this work provides a detailed analysis of the problem of assigning timestamps to the depth and thermal data. It also proposes a distinction between the concepts of data fusion and data validation where the former addresses the problem of determining which temperature measurements corresponds to which vertex on the 3D model and the latter addresses the problem of validating the measurement conditions before a temperature value is assigned to the corresponding vertex. In regards to sensor data fusion, it discusses the most effective approach to relate depth and thermal data and presents an effective method to reduce registration errors due to the parallax effect. In regards to data validation, it presents a new approach to determine when temperature values in the thermal texture of the 3D thermogram should be updated. Finally, this work shows that the use of shading and special color palettes helps improve the discernibility of features in renderings of 3D thermograms and proposes the use of additional textures to inform the user during the scan which parts of the model have been scanned under good measurement conditions and which need to be rescanned.

Contents

Abstract	iii
Contents	v
List of Figures	ix
List of Tables	xi
Acronyms	xiii
List of Symbols	xv
1 Introduction	1
1.1 Motivation	1
1.2 State of the art	3
1.3 Objectives and outline	5
2 Theoretical background	7
2.1 Radiometric quantities	7
2.2 Electromagnetic radiation	9
2.2.1 Black body	10
2.2.2 Real body	12
2.2.3 Other theoretical bodies	13
2.2.4 Propagation in a medium	14
2.2.5 Radiative properties at an interface	15
2.2.6 Radiative properties of opaque bodies	20
2.3 Optical systems	21
2.3.1 Thin lens	21
2.3.2 Depth of field	22
2.3.3 Field of view	23
2.3.4 Optical aberrations	24
2.3.5 Pinhole camera model	24
2.4 Summary	26
3 2D Thermal imaging	29
3.1 Radiometric chain	29
3.2 Thermographic camera	31
3.2.1 Basic components	31

3.2.2	Characterization	34
3.2.3	Switch-on behavior	34
3.3	Thermal image	36
3.3.1	Accuracy	36
3.3.2	Discernibility	41
3.4	Summary	42
4	3D thermal imaging: overview	43
4.1	System	43
4.2	Summary	46
5	3D thermal imaging: calibration	47
5.1	Radiometric calibration	47
5.1.1	Overview	47
5.1.2	Related work	48
5.1.3	Experiments	49
5.1.4	Analysis	52
5.1.5	Discussion	55
5.2	Geometric calibration	56
5.2.1	Overview	56
5.2.2	Related work	58
5.2.3	Calibration procedure	59
5.2.4	Experimental results and analysis	62
5.2.5	Discussion	67
5.3	Temporal calibration	68
5.3.1	Overview	68
5.3.2	Related work	69
5.3.3	Proposed approach	69
5.3.4	Experimental results	70
5.3.5	Discussion	72
5.4	Summary	73
6	3D thermal imaging: generation	75
6.1	3D model generation	75
6.1.1	Overview	75
6.1.2	Related work	75
6.1.3	Original approach and modifications	77
6.1.4	Experimental results	82
6.1.5	Discussion	83
6.2	Data fusion	84
6.2.1	Overview	84
6.2.2	Related work	85
6.2.3	Proposed approach	88
6.2.4	Analytical and experimental results	91
6.2.5	Discussion	92
6.3	Data validation	93
6.3.1	Overview	93

6.3.2	Related work	93
6.3.3	Proposed approach	94
6.3.4	Experimental results	96
6.3.5	Discussion	98
6.4	Data representation and visualization	100
6.4.1	Overview	100
6.4.2	Related work	101
6.4.3	Proposed approach	102
6.4.4	Experimental results	105
6.4.5	Discussion	108
6.5	Summary	109
7	3D Thermal imaging: case study	111
7.1	Scanning an industrial furnace	111
7.2	Discussion	116
7.3	Summary	117
8	Conclusions and outlook	119
8.1	Conclusions	119
8.2	Outlook	122
A	Radiative properties	127
B	Reflection and refraction at an interface	129
B.1	Snell's law	129
B.2	Fresnel equations	130
C	Device specifications	131
D	Radiometric calibration - supplemental material	133
E	Geometric calibration - error margins	135
F	Data fusion errors	139
F.1	Registration error due to timestamp errors	139
F.2	Registration error due to depth camera pose errors	141
	Publications	143
	Bibliography	145

List of Figures

1.1	Thermal images captured with a FLIR i7 thermal imaging camera	2
2.1	Solid angle	8
2.2	Electromagnetic spectrum and five bands infrared subdivision scheme . .	10
2.3	Spectral radiance of a black body	11
2.4	Radiance invariance	15
2.5	Directional spectral emissivities	21
2.6	Process of image formation using a symmetric thin lens	21
2.7	Depth of field	22
2.8	Field of view	23
2.9	Pinhole camera model	24
3.1	Radiometric chain	30
3.2	Atmospheric spectral transmittance based on the HITRAN database . . .	30
3.3	Main building blocks of a thermographic camera	32
3.4	Temperature measurements during the camera's warm up phase	35
3.5	Dependence of temperature measurement on angle of observation	37
3.6	Reflections on plates of different material	38
3.7	Effect of focusing on temperature measurement	38
3.8	Effect of distance to object on temperature measurement	39
3.9	Effect of camera movement on temperature measurement	39
3.10	Effect of thermal background radiation on temperature measurement . . .	40
3.11	Discernibility of objects in thermal images	41
4.1	Thermal camera and depth sensor rigidly attached to acrylic frame	44
4.2	RViz screenshot	46
5.1	Experimental setup	49
5.2	Color images of test patterns with holes of different radii r	50
5.3	Thermal images of the test patterns	50
5.4	Thermal images of test patterns at different distances	51
5.5	Thermal images of the calibration disc when it covers the entire FOV of the thermographic camera. $T_{\text{calibrator}} = 120^{\circ}\text{C}$	51
5.6	Temperature measured at the center of the calibration disc for patterns with different hole sizes and for different temperatures of the calibrator .	52
5.7	Difference between the maximum and minimum temperature measure- ments for each temperature setting of the calibrator	52
5.8	Temperature measured at the center of the calibrator for patterns with different hole sizes and at different distances d from the calibrator	53

5.9	Temperature profile from the center of the calibration disc to the upper right corner of the thermal image	54
5.10	Calibration target	60
5.11	Calibration images	61
5.12	Histograms of intrinsic parameters of thermographic camera	64
5.13	Error contribution caused by each individual distortion coefficient	65
5.14	Accuracy of the registration before and after fine-tuning the intrinsic and extrinsic parameters	68
5.15	Variability of image timestamps before filtering	71
5.16	Variability of image timestamps after filtering	71
5.17	Motion of 3D thermal imaging system to calibrate $\Delta t_{\text{delay_avg}}$ (top view)	72
5.18	Effects of inaccurate calibration of $\Delta t_{\text{delay_avg}}$	72
6.1	TSDF based surface description	78
6.2	Volume, FOV and coordinate systems at the beginning of a scan	79
6.3	Effect of voxel size on quality of 3D model	82
6.4	Effect of increased depth measurement errors on quality of 3D model	83
6.5	Interpolation of thermal camera pose from depth camera poses	88
6.6	Parallax effect	89
6.7	Linear interpolation error in a curved trajectory	92
6.8	Parallax effect in 3D thermogram	92
6.9	Data structure within a voxel of the thermal volume	96
6.10	Effect of validating the distance to the surface	97
6.11	Effect of validating the emission angle	99
6.12	Effect of shading for small and large temperature ranges	106
6.13	Visualizing the measurement distance for each temperature value	107
6.14	Visualizing the emission angle for each temperature value	107
7.1	3D thermogram of the access door of an operating industrial furnace	112
7.2	Registration errors due to depth measurement errors	113
7.3	Registration errors due to parallax when an object in the foreground leaves the FOV of the depth camera	114
7.4	Registration errors due to parallax when an object has not yet entered the operating range of the depth camera	115
7.5	Registration errors due to reflections	115
A.1	Radiative properties	128
B.1	Refraction of light at an interface	129
B.2	Reflection and refraction of light at an interface	130
D.1	Thermal images of test patterns taken at different distances d from the calibrator	133
D.2	Temperature measured at the center of the calibrator for patterns with different hole sizes and at different distances d from the calibrator	134
F.1	Variability of image timestamps after filtering	140
F.2	Rotational uncertainty	142

List of Tables

2.1	Radiometric quantities	9
5.1	Calibration results	62
5.2	Percentage change in the individual parameters from their estimated value that causes an increase of 0.1 pixels in the MRE	63
5.3	Intrinsic calibration using the same images but different subsets of features	66
C.1	Fluke 4180 precision infrared calibrator	131
C.2	PrimeSense Carmine 1.09 depth sensor	131
C.3	Optris PI 450 and Variocam hr head 600 thermographic cameras	132

Acronyms

CPU	central processing unit. 42 , 69 , 72 , 122
DOF	depth of field. 22 , 31 , 37 , 50 , 52 , 59 , 94 , 96 , 100 , 108
FIR	far-infrared. 10
FOV	field of view. vii , 23 , 32 , 33 , 42–44 , 49 , 50 , 55 , 59 , 64 , 77 , 78 , 82 , 84 , 92 , 99 , 100 , 103 , 113 , 117 , 120 , 122
FPA	focal plane array. 31–34 , 36 , 38 , 54 , 72 , 88 , 118 , 121
GPU	graphics processing unit. 42 , 77 , 115
HFOV	horizontal field of view. 23 , 64
ICP	iterative closest point. 43 , 77 , 79 , 82 , 108 , 115 , 119 , 122 , 138
IFOV	instantaneous field of view. 23 , 38 , 40 , 71 , 108 , 119
IPS	indoor positioning system. 85
ITS-90	International Temperature Scale of 1990. 46
LTE	local thermodynamic equilibrium. 13 , 14 , 17
LWIR	long-wavelength infrared. 10 , 31
MAE	mean absolute error. 56
MDTD	minimum detectable temperature difference. 33
MIR	mid-infrared. 10
MRE	mean reprojection error. 56 , 59 , 62 , 64–66 , 73 , 118
MRTD	minimum resolvable temperature difference. 33
MVS	multi-view stereo. 75

MWIR	mid-wavelength infrared. 10
NIR	near-infrared. 10
NUC	nonuniformity correction. 32–34 , 41 , 44 , 48 , 121
PCA	principal component analysis. 85
PCB	printed circuit board. 57
QWIP	quantum well infrared photodetectors. 31
RMSE	root mean squared error. 56
RTE	radiative transfer function. 14
SfM	structure from motion. 75 , 86
SIF	scene influence factor. 47 , 48
SIFT	scale-invariant feature transform. 85
SLERP	spherical linear interpolation. 87
SNR	signal to noise ratio. 31
SWIR	short-wavelength infrared. 10
TLS	terrestrial laser scanner. 76
ToF	time of flight. 86
TSDF	truncated signed distance function. 77–79 , 82 , 88 , 89 , 96 , 103 , 105 , 108 , 115
UAV	unmanned aerial vehicle. 3
UGV	unmanned ground vehicle. 3
VFOV	vertical field of view. 23 , 64

List of Symbols

Roman symbols

A	area
c	diameter of circle of confusion
c_u, c_v	pixel coordinates where the optical axis intersects the image
\mathbf{c}	voxel's centroid position vector
\mathbf{P}	point cloud array
d	aperture diameter; distance
D	focus distance
D_N	nearest focus distance
D_F	farthest focus distance
\mathbf{D}	minimum distance map
E	irradiance
E_λ	spectral irradiance
f	focal length; scaling factor
f_x, f_y, f_u, f_v	scaling factors
h	height
i, j, k	voxel indices in volumetric array; discrete time
I	radiant intensity; image
I_λ	spectral intensity
J	radiosity
J_λ	spectral radiosity
\mathbf{K}	camera matrix
k_u, k_v	inverse of detector size
k_1, k_2, k_3	radial distortion coefficients
k'_1, k'_2, k'_3	radial distortion coefficients (using normalized coordinates)
L	radiance
L_λ	spectral radiance
M	radiant exitance; number of images
M_λ	spectral exitance

M_v	vertex map
M_n	normal map
N	f-number or f-stop; number of features per image; median filter length
\mathbf{n}	normal vector
\bar{n}	complex refractive index
\bar{n}_λ	spectral complex refractive index
n	refractive index
p_1, p_2	tangential distortion coefficients
p'_1, p'_2	tangential distortion coefficients (using normalized coordinates)
\mathbf{p}	position vector
\mathbf{P}	color palette
\mathbf{q}	quaternion
r	radius
\mathbf{r}	position vector
\mathbf{R}	rotation matrix
t	time
\mathbf{t}	translation vector
T	temperature
\mathbf{T}	geometric transformation (rotation and translation)
u, v	pixel coordinates
\mathbf{v}	vertex position vector
\mathbf{V}	volumetric array
w	width, weight
x, y, z	Cartesian coordinates

Greek symbols

α	absorptivity; shearing factor; interpolation weight
α_λ	spectral absorptivity
δ_x, δ_y	offset due to geometric distortion
ε	emissivity
ε_λ	spectral emissivity
θ	polar angle
κ	extinction coefficient
λ	wavelength
$\rho_{\text{bi-d}}$	bidirectional total reflectivity
$\rho_{\lambda, \text{bi-d}}$	bidirectional spectral reflectivity
$\rho_{\text{d-h}}$	directional-hemispherical total reflectivities
$\rho_{\lambda, \text{d-h}}$	directional-hemispherical spectral reflectivities

$\rho_{\text{h-d}}$	hemispherical-directional total reflectivities
$\rho_{\lambda,\text{h-d}}$	hemispherical-directional spectral reflectivities
ρ_{s}	specular reflectivity
$\rho_{\lambda,\text{s}}$	spectral specular reflectivity
$\sigma_{\lambda,\text{a}}$	spectral absorption coefficient
$\sigma_{\lambda,\text{s}}$	spectral scattering coefficient
τ	transmissivity
Φ	radiant flux
Φ_{λ}	spectral radiant flux
ϕ	azimuth angle
ϕ_{s}	timing jitter
Ω	solid angle

Chapter 1

Introduction

1.1 Motivation

Infrared thermography is a technique that enables the non-contact measurement of an object's surface temperature and presents the results in form of thermal images. The analysis of these images, also known as thermograms, provides valuable information about an object's thermal behavior and allows in many cases the diagnostic of abnormal conditions. Infrared thermography is extensively used in the industry sector to inspect heat generating facilities, such as furnaces, boilers, engines, heat exchangers and radiators and to monitor the quality of components produced in thermal processes such as press hardening. Infrared thermography is also used in the building sector to carry out energy audits and detect abnormal losses and gains of heat. Thermograms help auditors determine whether a piece of equipment is overheating and where it is happening, whether fabricated components meet the expected temperature profile and whether a building needs insulation and where it should go. As such, infrared thermography is a valuable tool in the assessment of energy efficiency.

The fidelity with which thermograms generated by a thermographic camera depict the real surface temperature of an object depends, however, on a number of factors. Among these factors are the position of the camera with respect to the object's surface during the inspection. On the one hand, surfaces that are captured outside the depth of field of the thermographic camera look blurred in the thermal image and their estimated temperatures are often incorrect, as their values result from averaging the temperature values measured by the neighboring pixels. On the other hand, the accuracy of the measurements depends on the orientation of the surface with respect to the camera's optical axis. In other words, there is an inherent limitation in trying to accurately portray the 3D temperature profile of a three dimensional object by means of 2D images, when

depth information is not available. Moreover, it is often difficult to recognize objects and their constituent parts in thermal images because, despite the color mapping, these provide less information than regular color images (one channel vs three channels of information) and because the human brain uses other type of information to infer the geometry of an object and identify it, for example, color shading. Figure 1.1 shows two objects that would be easily recognized in a conventional color image but are difficult to identify in a thermal image. For this reason, the documentation of thermographic inspections is often accompanied with color images. Finally, it is difficult to carry out accurate heat loss/gain estimations by means of 2D images, especially for objects with complex geometries, because in general the true size of the inspected surface area cannot be estimated from the thermal image.




Figure 1.1: Thermal images captured with a FLIR i7 thermal imaging camera

3D thermography has the potential to overcome many of the limitations that affect conventional 2D thermography. On the one hand, the geometry of the object is embedded in the three dimensional representation of the results. This enables a more accurate estimation of the heat loss/gain experienced by the object's surfaces. It also facilitates the identification of the inspected objects and their features. On the other hand, knowledge about the position of the thermographic camera with respect to the object during the inspection can be exploited to improve the accuracy of the 3D thermogram. This can be achieved, for example, by making sure that only data captured from perspectives that provide appropriate measurement conditions are used to generate the 3D thermogram. Furthermore, geometric data and other information that can be derived from them can be used to assist users with little expertise in order to acquire higher quality data during the inspection. Alternatively, this information could be used as feedback to completely automatize the inspection process. Finally, 3D thermography can help improve the visualization and analysis of the temperature distribution of large objects with complex geometries such as the interior of large machines or buildings. Instead of using several overlapping thermal images to document the inspection, the entire scene

can be depicted by a single 3D thermogram which can then be viewed and analyzed with a suitable software package.

The use of 3D thermography has been proposed for applications such as interior and exterior building energy auditing [19, 28, 41], medical diagnostic [7, 25], biometrics [58], industrial inspection [2], search and rescue [59], lab testing [14, 65] and the recording of architectural monuments [13]. However, most of the 3D thermographic systems that have been developed so far generate 3D thermograms with data acquired from a single perspective or from few noncontiguous points of view, as they have been implemented with cumbersome equipment that is difficult to move. As a consequence, few strategies to process and integrate thermal data taken from multiple perspective have been proposed. Furthermore, the 3D thermograms are not generated in real time so the person scanning the object does not have feedback on the quality of the data being captured which varies with the pose of the camera relative to the object. Consequently, these systems do not exploit the full potential of 3D thermography. But reductions in cost, size and weight of thermographic cameras and the increase of computational power over the last years as well as the recent introduction to the market of new low-cost and lightweight depth sensors have unlocked the possibility to implement affordable handheld 3D thermal imaging systems that can be easily maneuvered around the object to be inspected. Such a system could be mounted, for example, on a remote controlled or autonomous *unmanned ground vehicle (UGV)* or *unmanned aerial vehicle (UAV)* to carry out inspections of large machines or buildings from different perspectives. This opens up the opportunity for novel research on how to further improve the design of 3D thermographic systems, the fidelity of the resulting 3D thermograms and the interaction with the user.

1.2 State of the art

The generation of high-fidelity 3D thermograms depends on both the design of an accurate and reliable 3D thermographic system and the acquisition of thermal and geometric data under the best possible measurement conditions.

Many different types of 3D thermographic systems have been proposed in the literature. These systems employ distinct combinations of sensors and methods to generate the 3D thermograms. The works usually include procedures to estimate the parameters of the mathematical models that describe the sensors and to determine the position and orientation of the sensors relative to each other. The methods employed derive from the fields of computer vision and photogrammetry and have been adapted to fit the particular characteristics and requirements of thermographic cameras. A comparative analysis of different methods that have been used to calibrate the sensors, reconstruct

the 3D model and fuse geometric and thermal data is presented in greater depth in chapters 5 and 6.

In some works, 3D thermograms are generated from a single perspective. Once the parameters of the system are known, the fusion of geometric and thermal information is straightforward: the temperature value associated with each vertex of the 3D model is obtained by projecting the vertex onto the thermal image. 3D thermograms generated this way, however, are affected by the same problem that affects conventional 2D thermal images: the accuracy of the measurements depends on the orientation of the surface with respect to the camera's optical axis. The works of [77] and [3] recognize this problem and attempt to solve it using the geometric information embedded in the 3D thermogram. They propose to adjust the measured temperature value with a correction factor based on the position and orientation of the camera with respect to the object's surface. This idea presents two inconveniences. On the one hand, the thermal radiation emitted by many surfaces drops very fast for emission angles greater than 70° and equals zero at 90° . This means that as the angle approaches 90° , the correction factor would approach infinity. Moreover, since it is difficult to get accurate estimates of the emission angle with current depth sensing technologies, especially at large angles, this approach could potentially introduce more errors than it would correct for. On the other hand, as the emissivity of a surface decreases, its reflectivity increases. This means, errors due to background reflections could potentially be amplified by the correction factor. Another problem with generating 3D thermograms from a single perspective is that there might not be a single perspective from which to cover the entire area of interest. Some regions might be occluded by objects in the foreground or by the geometry of the scanned object itself. In [23], for example, a 3D thermogram generated from a single perspective is used to determine the heat loss of the main facade of a concrete building but, from the sensor's point of view, parts of the facade are blocked from sight by trees. In this particular case, the occluded regions are small and probably do not have a strong influence on the accuracy of the heat loss estimation. Nonetheless, it is clear that a 3D thermographic system capable of generating high-fidelity 3D thermograms should do so using data from multiple perspectives.

In other works, 3D thermograms are generated with data captured from different perspectives either by rotating the object in front of the 3D thermographic system or by placing the 3D thermographic system at discrete locations with respect to the object or scene. Despite being generated with a more comprehensive data set, the fidelity of the resulting 3D thermograms falls short of the full potential of 3D thermography because little attention has been paid to determine an appropriate way to merge data captured under different measurement conditions. The thermographic system used in [19], for example, assigns temperature values to the point cloud generated by a 3D laser scanner

without taking into account the position of the thermal camera with respect to the objects surface. As a consequence, the thermal texture of the 3D thermogram consists of a mix of temperature measurements with varying degrees of accuracy and spatial resolution, but there is no way to tell which measurements are more accurate than others. In [60], on the other hand, thermal images are treated as regular color images and the color assigned to each vertex of the 3D model results from averaging the color associated with it in thermal images taken from different points of view. The angle of emission is used as weight in the calculation of the average color. Averaging colors does not improve the accuracy of the temperature values represented by the thermal texture. In fact, the averaged color might not even form part of the original color palette.

The system proposed in [84] also generates 3D thermograms with data captured from different perspectives but, unlike the aforementioned systems, data are captured while the sensors are in motion. Because the geometric and thermal data are updated at least 30 times per second, the resulting 3D thermograms are able to describe the temperature distribution at the object's surface with great detail. The temperature value assigned to each vertex of the 3D model results from averaging temperature measurements corresponding to the same point on the object but taken from different perspectives. Weights based on various measurement conditions are used in the calculation. This approach produces smooth thermal textures which may even hide specular reflections on shiny surfaces but, as will be explained in section 6.3, it can reduce the accuracy of the temperature measurements. The system was modified later on to generate and visualize the 3D thermograms in real time [56]. This improvement allows the user to check the integrity of the 3D model during the scan but the user still gets no feedback on the quality of the measurements that are being made so he or she cannot corrective actions during the scan. In fact, a review of the literature revealed that no work exploits the geometric information gathered by the 3D thermographic system to provide the user with feedback on the quality of the thermal data.

Finally, in other works the thermal images are mapped onto computer generated 3D models [14, 31, 34]. The main disadvantage of these approaches is that there is not always a 3D model of the object to be inspected and when there is one, it might not accurately describe the real object.

1.3 Objectives and outline

A review of the state of the art related to field of close range 3D thermography reveals that various aspects concerning the generation of high-fidelity 3D thermograms have not been addressed, are not well understood or still have potential for improvement. With

this in mind, this thesis attempts to accomplish three objectives. The first objective is to identify the different aspects involved in the generation of high-fidelity 3D thermograms based on the work that has already been contributed to the field of 3D thermography. The second objective is to provide a deeper understanding of the problems that affect each of these aspects either by means of experiments or through mathematical analysis and, where possible, find an alternative approach that helps improve the current state of the art. The third objective is to develop new approaches to effectively exploit the geometric information available to 3D thermographic systems in real time to improve the fidelity and visualization of the generated 3D thermograms.

The remainder of this thesis is organized as follows. Chapter 2 introduces the main theoretical concepts needed to understand the fundamentals of infrared thermography as well as the related terminology. Chapter 3 uses these to explain the principle of operation of a typical 2D thermal imaging system and to describe the limitations and main issues that affect the fidelity of conventional thermal images. Chapter 4 introduces the 3D thermal imaging system used in this work. Chapter 5 elaborates on the different aspects involved in the calibration of the 3D thermal imaging system. Similarly, chapter 6 elaborates on the different aspects involved in the generation, representation and visualization of high-fidelity 3D thermograms. Each aspect is worked out in a separate section which includes a discussion of related previous works and the contributions made in this work to the corresponding subject. Finally, Chapter 8 provides a summary of the contents and findings of this thesis and an outlook of future work. The appendices include additional graphs that illustrate the radiative properties explained in chapter 2, formulas used to generate some plots, a quick reference to the technical specifications of the sensors, additional experimental results and a more thorough analysis of certain topics.

Chapter 2

Theoretical background

This chapter briefly presents the theoretical background underlying 3D thermal imaging systems. It is meant to provide the reader with basic knowledge to understand the subsequent chapters as well as an explanation of terms and concepts used in this work.

2.1 Radiometric quantities

Radiometric quantities are used to describe the propagation of radiative energy in a medium as well as its behavior at the interface between two different media.

Solid angle

The solid angle is an extension of the notion of the planar angle from 2D to 3D space. A planar angle represents a section of a circle and its magnitude is proportional to the length of the subtended arc. In analogy, a solid angle represents a section of a sphere and its magnitude is proportional to the area of the subtended spherical cap. Solid angles are expressed in a dimensionless unit called steradian (sr). By definition, the solid angle Ω is the ratio of the area A of a spherical cap to the square of the sphere radius r (2.1a). Often, equations that describe how the radiative energy propagates through a medium from one surface to another are expressed in spherical coordinates and as a function of the differential solid angle $d\Omega$ (2.1b), where dA is the effective area perpendicular to the direction (θ, ϕ) . These two concepts are illustrated in Figure 2.1.

$$\Omega = \frac{A}{r^2} \tag{2.1a}$$

$$d\Omega = \frac{dA}{r^2} = \frac{r \sin\theta d\phi \cdot r d\theta}{r^2} = \sin\theta d\theta d\phi \tag{2.1b}$$




Figure 2.1: Solid angle

When working with radiometric quantities, the product $\cos\theta d\Omega$ comes up very often. In analogy to the concept of projected (or effective) area element $dA_p = \cos\theta dA$, the projected (or effective) solid angle element is defined by (2.2).

$$d\Omega_p = \cos\theta d\Omega = \cos\theta \sin\theta d\theta d\phi \quad (2.2)$$

Radiant Flux and Derivatives

The base quantity that describes the amount of radiative power that leaves or reaches a surface is the radiant flux. It is represented by the symbol Φ and its unit is the watt. Other radiometric quantities like irradiance, radiosity, radiant exitance, intensity and radiance as well as their corresponding spectral counterparts describe density functions of the radiant flux in terms of area, solid angle and wavelength for different flux directions. These quantities are summarized in Table 2.1.

From the definitions in Table 2.1 it follows that the spectral exitance, the spectral intensity and the spectral radiance are related to the radiant and spectral flux by (2.3a), (2.3b) and (2.3c), respectively.

$$M_\lambda = \frac{d^2\Phi}{dA d\lambda} = \frac{d\Phi_\lambda}{dA} \quad (2.3a)$$

$$I_\lambda = \frac{d^2\Phi}{d\Omega d\lambda} = \frac{d\Phi_\lambda}{d\Omega} \quad (2.3b)$$

$$L_\lambda = \frac{d^3\Phi}{dA \cos\theta d\Omega d\lambda} = \frac{d^2\Phi_\lambda}{dA \cos\theta d\Omega} \quad (2.3c)$$

Here, dA denotes a differential area element located at the center of the sphere and the variable θ represents the angle between the normal to the area element dA and the direction of the ray specified by $d\Omega$.

Table 2.1: Radiometric quantities

Quantity	Symbol	Description	Unit
Base quantity			
Radiant flux	Φ	Radiant electromagnetic power received by or leaving ^a a surface	W
Single density quantities			
Irradiance	E	Radiant flux received by a surface per unit area	W m^{-2}
Radiosity	J	Radiant flux leaving ^a a surface per unit area	W m^{-2}
Radiant exitance	M	Radiant flux emitted by a surface per unit area	W m^{-2}
Radiant intensity	I	Radiant flux received by or leaving ^a a surface per unit solid angle	W sr^{-1}
Spectral flux	Φ_λ	Radiant flux per unit wavelength	W m^{-1}
Joint density quantities of two variables			
Spectral irradiance	E_λ	Radiant flux received by a surface per unit area per unit wavelength	W m^{-3}
Spectral radiosity	J_λ	Radiant flux leaving ^a a surface per unit area per unit wavelength	W m^{-3}
Spectral exitance	M_λ	Radiant flux emitted by a surface per unit area per unit wavelength	W m^{-3}
Spectral intensity	I_λ	Radiant flux received by or leaving ^a a surface per unit solid angle per unit wavelength	$\text{W sr}^{-1} \text{m}^{-1}$
Radiance	L	Radiant flux received by or leaving ^a a surface per unit solid angle per unit projected area	$\text{W sr}^{-1} \text{m}^{-2}$
Joint density quantity of three variables			
Spectral radiance	L_λ	Radiant flux received by or leaving ^a a surface per unit solid angle per unit projected area per unit wavelength	$\text{W sr}^{-1} \text{m}^{-3}$

^a ... emitted, reflected and transmitted by ...

For more information regarding radiometric quantities the reader is referred to [11].

2.2 Electromagnetic radiation

In thermodynamics, temperature is a measure of the total internal energy of an object that results from translational, rotational and vibrational motion of its constituent particles (molecules, atoms and subatomic particles). The higher the object's temperature, the higher its internal energy. Subatomic particles are constantly absorbing and emitting energy in form of electromagnetic radiation as they interact with each other and move back and forth between different energy states. The higher the object's temperature, the

higher the frequency with which the subatomic particles switch energy states and also the greater the energy that they emit. At the object's boundaries, this electromagnetic energy is radiated outwards to the surroundings. Thus, all objects at a temperature above 0 K emit some amount of electromagnetic radiation. The total electromagnetic power that an object emits strongly depends on its surface temperature. Thermal imaging systems exploit this fact to estimate an object's surface temperature by measuring the electromagnetic power it radiates. Most of the electromagnetic power radiated by objects commonly found in measurement and inspection tasks is emitted in the infrared range of the electromagnetic spectrum. For this reason, thermal imaging is also usually referred to as infrared imaging.

Several subdivision schemes are used to describe the infrared spectrum. These typically consist of three bands —*near-infrared (NIR)*, *mid-infrared (MIR)* and *far-infrared (FIR)*— or five bands —*near-infrared (NIR)*, *short-wavelength infrared (SWIR)*, *mid-wavelength infrared (MWIR)*, *long-wavelength infrared (LWIR)* and *far-infrared (FIR)*. The boundaries of each band are defined somewhat arbitrarily and vary from standard to standard. Criteria commonly used for defining these boundaries are the responsivity of infrared detector or the atmospheric windows. Figure 2.2 shows the position of the infrared band within the electromagnetic spectrum as well as a five bands subdivision scheme.




Figure 2.2: Electromagnetic spectrum and five bands infrared subdivision scheme

2.2.1 Black body

A black body is a theoretical physical body that absorbs all incident electromagnetic radiation, regardless of wavelength and angle of incidence. A black body also behaves as a perfect emitter. It emits the maximum amount of electromagnetic radiation that an object can emit at any temperature and for any wavelength. Furthermore, the radiance emitted by a black body at constant temperature is the same in all directions. The spectral radiance of a black body is described by Planck's law (2.4) and is depicted in

Figure 2.3.

$$L_{\lambda,bb}(T) = \frac{2hc_o^2}{n^2\lambda^5(\exp(\frac{hc_o}{nk_B\lambda T}) - 1)} \quad (2.4)$$

Here, $L_{\lambda,bb}(T)$ denotes the emitted spectral radiance as a function of the black body's temperature T , h is Planck's constant, c_o is the speed of light in vacuum, n is the refraction index of the medium and k_B is Boltzmann's constant. According to [57], the values of h , c_o and k_B are $6.626\,070\,04 \times 10^{-34}$ J s, $2.997\,924\,58 \times 10^8$ m s $^{-1}$ and $1.380\,648\,52 \times 10^{-23}$ J K $^{-1}$, respectively.

The plot includes the spectral radiance of a black body at -20°C , 25°C (room temperature) and 100°C to illustrate typical values of spectral radiance detected by thermographic cameras. It also includes the spectral radiance of a black body at 5778°C , which approximates the spectral radiance emitted by the sun, to illustrate that the peak of this curve falls within the visible part of the electromagnetic spectrum.

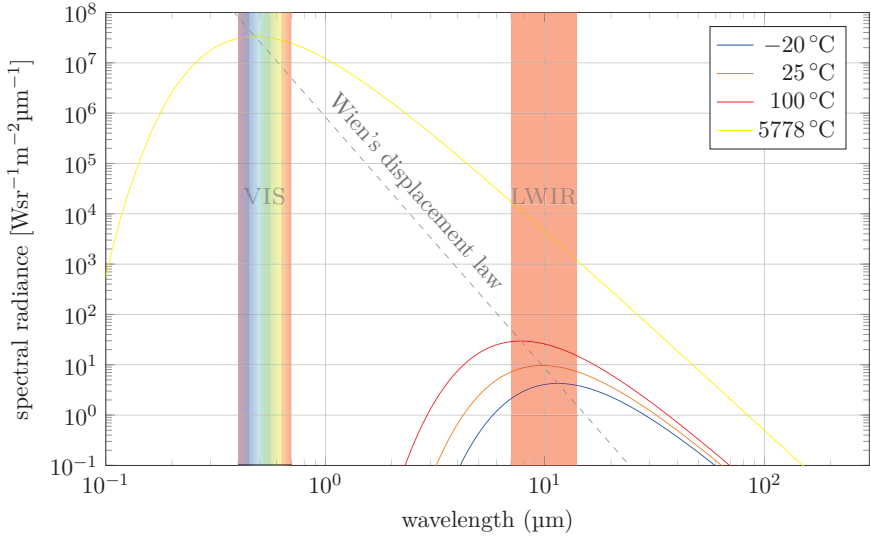


Figure 2.3: Spectral radiance of a black body

The integration of (2.4) for $n = 1$ (i.e. in vacuum) over a hemispherical solid angle¹ centered on the area element dA and over all wavelengths results in (2.5). This equation is known as the Stefan-Boltzmann law and σ denotes the Stefan-Boltzmann constant. According to [57], σ is equal to $5.670\,367 \times 10^{-8}$ W m $^{-2}$ K $^{-4}$. The equation indicates the maximum radiant exitance $M_{bb}(T)$ that an area element dA at temperature T can emit into a hemisphere. Figure 2.3 shows that the area under the curves increases as the

¹ The integration over a hemispherical solid angle is represented by the surface integral $\iint_{2\pi}(\)d\Omega$ and is equivalent to the double integral $\int_0^{2\pi} \int_0^{\pi/2}(\)\sin\theta d\theta d\phi$.

temperature increases.

$$M_{\text{bb}}(T) = \int_0^\infty \iint_{2\pi} L_{\lambda,\text{bb}}(T) \cos\theta d\Omega d\lambda = \sigma T^4 \quad (2.5)$$

Similarly, calculating the derivative of (2.4) with respect to λ and equating it to zero for $n = 1$ (i.e. in vacuum) results in (2.6). This equation is known as Wien's displacement law and b represents Wien's displacement constant. According to [57], b is equal to $2.897\,772\,9 \times 10^3 \text{ m K}$. The equation allows the calculation of the wavelength λ_{max} at which the spectral radiance $L_{\lambda,\text{bb}}(T)$ reaches its maximum as a function of the black body's temperature T . Figure 2.3 shows that the maximum spectral radiance shifts towards shorter wavelengths as the temperature increases.

$$\lambda_{\text{max}} T = b \quad (2.6)$$

2.2.2 Real body

In practice, the amount of radiant power that real bodies emit or absorb depends not only on their temperature but also on local properties such as surface finish and the direction of the emitted or incident electromagnetic radiation. To take into account these factors two new quantities are introduced: emissivity and absorptivity. The emissivity specifies the fraction of the radiant power emitted by a black body that a real body at the same temperature emits. It is a directional quantity because its value varies with the direction of emission. From its definition it follows that the emissivity of a black body is equal to one. The emissivity is represented by the Greek letter ε . The absorptivity, on the other hand, specifies the fraction of the radiant power incident on a real body that gets absorbed by it. It is also a directional quantity because its value depends on the direction of the incident radiation. The absorptivity is represented by the Greek letter α . Since a black body absorbs all incident electromagnetic radiation, its absorptivity is one.

Due to the law of conservation of energy, the radiant power incident on a real body that does not get absorbed by it must either be reflected back into the surroundings or be transmitted through the body. To take into account these effects two additional quantities are introduced: reflectivity and transmissivity. The reflectivity specifies the fraction of radiant power incident on a body that gets reflected by it. It is a bidirectional quantity because its value varies both with the direction of incidence and the direction of reflection. The reflectivity is represented by the Greek letter ρ . Since a black body absorbs all incident electromagnetic radiation, its reflectivity is equal to zero. Similarly, the transmissivity specifies the fraction of radiant power incident on a body that gets transmitted by it. It is also a bidirectional quantity because its value varies both with the

direction of incidence and the direction of transmission. The transmissivity is represented with the Greek letter τ . The transmissivity of a black body is also equal to zero.

Manufacturers of radiation thermometers provide tables with typical values of emissivity for different materials at different temperatures. These values are frequently measured with the surface oriented perpendicular to the radiometer and represent an average emissivity within the operating wavelength range of the radiometer. A comparison of the emissivity of various materials for different spectral ranges can be found in [11].

The terms emissivity, absorptivity, reflectivity and transmissivity are used to describe intensive properties of matter, i.e., properties that do not depend on the amount of material involved, for example, the emissivity of polished copper. The terms emittance, absorptance, reflectance and transmittance follow similar definitions but are used to describe extensive properties, i.e., properties that depend on the amount the material involved, for example, the atmospheric transmittance. All these properties vary with the temperature of the body.

2.2.3 Other theoretical bodies

The processes of emission, absorption, reflection and transmission of electromagnetic radiation are very complex. To facilitate their description, analysis and understanding other types of theoretical bodies are commonly used. A white body, for example, behaves opposite to a black body: it has zero absorptivity and perfect reflectivity for all wavelengths. In nature, no true white bodies are known, but some materials exhibit very high reflectivity within a particular wavelength interval. Similarly, a gray body refers to a body that emits a constant fraction of the radiation emitted by a black body, regardless of wavelength. A transparent body, on the other hand, is one that transmits all radiation incident on it and an opaque body, on the contrary, is a body that transmits no radiation. Radiation incident on an opaque body is either absorbed or reflected. This type of body offers a good approximation to describe a large number of objects found in real applications. A common practice is to assume in addition that the opaque body's surface is gray and Lambertian. A Lambertian surface emits and reflects radiance uniformly in all directions.

The equation $\varepsilon + \rho + \tau = 1$ is often used in the literature to indicate that, under *local thermal equilibrium* (LTE), energy is conserved. This expression only holds for ideal surfaces that are both gray and Lambertian and ε , ρ and τ indicate hemispherical total quantities. The emissivity, reflectivity and transmissivity of real bodies vary with wavelength and direction of propagation of the electromagnetic radiation, as will be discussed later on in this chapter.

2.2.4 Propagation in a medium

When the electromagnetic radiation emitted by an object propagates through a medium other than vacuum, it loses some power due to absorption and out-scattering but also gains some power due to local emission of the medium and in-scattering. The net change in radiant power depends on the properties of the medium along the traversed path. The propagation of radiance through a volume element dV in [LTE](#) is described by (2.7).

$$\begin{aligned}
 \frac{\partial L_\lambda(S, \Omega, t)}{c_o \partial t} + \frac{\partial L_\lambda(S, \Omega, t)}{\partial S} = & \underbrace{-\sigma_{\lambda,a}(S)L_\lambda(S, \Omega, t)}_{\text{loss due to absorption}} \underbrace{-\sigma_{\lambda,s}(S)L_\lambda(S, \Omega, t)}_{\text{loss due to out-scattering}} \\
 & \underbrace{+\sigma_{\lambda,a}(S)L_{\lambda,bb}(S)}_{\text{gain due to emission}} \\
 & \underbrace{+\frac{1}{4\pi} \iint_{4\pi} \sigma_{\lambda,s}(S)L_\lambda(S, \Omega_i, t)P_\lambda(S, \Omega_i, \Omega)d\Omega_i}_{\text{gain due to in-scattering}} \quad (2.7)
 \end{aligned}$$

This formula is known as the *radiative transfer equation* ([RTE](#)). The terms $\sigma_{\lambda,a}(S)$ and $\sigma_{\lambda,s}(S)$ represent the spectral absorption and the scattering coefficient of the medium, respectively. They have units of reciprocal length and their values depend on local properties of the medium including pressure, temperature and concentration of the constituents materials as well as the wavelength of the traversing radiance. $L_\lambda(S, \Omega, t)$ is the spectral radiance in the direction of propagation Ω . $L_{\lambda,bb}(S)$ is the spectral radiance emitted by a black body at the same temperature as the volume element dV . Ω_i is an integration variable that represents the direction of incident radiance. $L_\lambda(S, \Omega_i)$ is the spectral radiance incident on dV from the direction Ω_i . $P_\lambda(S, \Omega_i, \Omega)$ is the scattering phase function, which describes how much of $L_\lambda(S, \Omega_i)$ is scattered in the direction of propagation Ω . Gains due to in-scattering are integrated from all possible directions, that is, over a spherical solid angle². All variables are evaluated at the position S and time t .

When $L_\lambda(S, \Omega, t)$ varies slowly in time, the medium does not cause scattering in a given wavelength interval (i.e. $\sigma_{\lambda,s}(S) = 0$) and the local energy contribution of the medium is negligible when compared to the magnitude of $L_\lambda(S, \Omega)$, the [RTE](#) reduces to (2.8a) and the solution to this differential equation is given by (2.8b), where the direction Ω is omitted for simplicity. This equation is known as Bouguer's law, Beer's law, the

² The integration over a spherical solid angle is represented by the surface integral $\iint_{4\pi} () d\Omega$ and is equivalent to the double integral $\int_0^{2\pi} \int_{-\pi/2}^{\pi/2} () \sin\theta d\theta d\phi$.

Lambert-Beer law or the Bouguer-Lambert-Beer law.

$$\frac{dL_\lambda(S, \Omega)}{dS} = -\sigma_{\lambda,a}(S)L_\lambda(S, \Omega) \quad (2.8a)$$

$$L_\lambda(S) = L_\lambda(0)\exp\left(-\int_0^S \sigma_{\lambda,a}(s)ds\right) \quad (2.8b)$$

When the medium is vacuum, $\sigma_{\lambda,a}$ and $\sigma_{\lambda,s}$ are both zero and L_λ remains constant along the direction given by Ω . This property is known as radiance invariance. The same result can be obtained considering the spectral flux Φ_λ that passes through the two area elements dA_1 and dA_2 depicted in Figure 2.4. Since there are no losses along the path, according to (2.3c), this flux can be expressed either as (2.9a) or as (2.9b), where $d\Omega_1$ is the solid angle subtended by dA_2 when viewed from dA_1 and $d\Omega_2$ is the solid angle subtended by dA_1 when viewed from dA_2 .




Figure 2.4: Radiance invariance

$$d^2\Phi_\lambda = L_{\lambda,1}\cos\theta_1 dA_1 d\Omega_1 \quad (2.9a)$$

$$d^2\Phi_\lambda = L_{\lambda,2}\cos\theta_2 dA_2 d\Omega_2 \quad (2.9b)$$

Since dA_1 and dA_2 are not perpendicular to r , according to (2.1b), $d\Omega_1 = \cos\theta_2 dA_2 / r^2$ and $d\Omega_2 = \cos\theta_1 dA_1 / r^2$. It follows then that $L_{\lambda,1} = L_{\lambda,2} = L_\lambda$. This means that the radiance that passes through dA_1 and dA_2 is the same.

2.2.5 Radiative properties at an interface

To better describe the directional nature of the radiative properties at the surface of a real body, rigorous definitions of emissivity, absorptivity, reflectivity and transmissivity are necessary. These definitions also enable a more accurate description of the relationship between the radiative properties. In this section, only the definitions that apply to opaque bodies are presented. Although not all definitions are used later in this work, they are included nonetheless for consistency. To improve the readability of the definitions, the dependence of the radiative quantities and radiative properties on the temperature of the real body is omitted. Illustrations of these definitions can be found in appendix A.

Directional emissivities: $\varepsilon_\lambda(\theta, \phi)$ and $\varepsilon(\theta, \phi)$

Let $d^2\Phi_{\lambda,bb}(\theta, \phi)$ be the spectral flux emitted by an area element dA of a black body at temperature T within the solid angle $d\Omega$ in the direction (θ, ϕ) . Similarly, let $d^2\Phi_\lambda(\theta, \phi)$ be the spectral flux emitted by an area element dA of a real body at the same temperature T within the same solid angle $d\Omega$ in the same direction (θ, ϕ) . The directional spectral emissivity $\varepsilon_\lambda(\theta, \phi)$ is defined as the ratio of $d^2\Phi_\lambda(\theta, \phi)$ to $d^2\Phi_{\lambda,bb}(\theta, \phi)$ and is given by (2.10a). The definition of the directional total emissivity $\varepsilon(\theta, \phi)$ is an extension of the definition of $\varepsilon_\lambda(\theta, \phi)$ that includes all wavelengths of the electromagnetic spectrum and is given by (2.10b).

$$\varepsilon_\lambda(\theta, \phi) = \frac{d^2\Phi_\lambda(\theta, \phi)}{d^2\Phi_{\lambda,bb}(\theta, \phi)} = \frac{L_\lambda(\theta, \phi)\cos\theta dA d\Omega}{L_{\lambda,bb}\cos\theta dA d\Omega} = \frac{L_\lambda(\theta, \phi)}{L_{\lambda,bb}} \quad (2.10a)$$

$$\varepsilon(\theta, \phi) = \frac{\int_0^\infty L_\lambda(\theta, \phi) d\lambda}{\int_0^\infty L_{\lambda,bb} d\lambda} \quad (2.10b)$$

Hemispherical emissivities: ε_λ and ε

Let $d^2\Phi_{\lambda,bb}$ be the spectral flux emitted by an area element dA of a black body at temperature T in all directions of a hemisphere centered on dA . Similarly, let $d^2\Phi_\lambda$ be the spectral flux emitted by an area element dA of a real body at the same temperature T in all directions of a hemisphere centered on dA . The hemispherical spectral emissivity ε_λ is defined as the ratio of $d^2\Phi_\lambda$ to $d^2\Phi_{\lambda,bb}$ and is given by (2.11a), where the terms dA cancel out. The definition of the hemispherical total emissivity ε is an extension of the definition of ε_λ that includes all wavelengths of the electromagnetic spectrum and is given by (2.11b).

$$\varepsilon_\lambda = \frac{d^2\Phi_\lambda}{d^2\Phi_{\lambda,bb}} = \frac{\iint_{2\pi} L_\lambda(\theta, \phi)\cos\theta d\Omega}{\iint_{2\pi} L_{\lambda,bb}\cos\theta d\Omega} \quad (2.11a)$$

$$\varepsilon = \frac{\int_0^\infty \iint_{2\pi} L_\lambda(\theta, \phi)\cos\theta d\Omega d\lambda}{\int_0^\infty \iint_{2\pi} L_{\lambda,bb}\cos\theta d\Omega d\lambda} \quad (2.11b)$$

The directional spectral emissivity, the directional total emissivity, the hemispherical spectral emissivity and the hemispherical total emissivity are dimensionless quantities.

Directional absorptivities: $\alpha_\lambda(\theta, \phi)$ and $\alpha(\theta, \phi)$

Let $d^2\Phi_{\lambda,i}(\theta_i, \phi_i)$ be the spectral flux incident on an area element dA coming from the direction (θ_i, ϕ_i) that passes normally through an area element dA_i on a hemisphere

centered on dA within an a solid angle $d\Omega$, where $d\Omega$ is the solid angle subtended by dA when viewed from dA_i . On the other hand, let $d^2\Phi_{\lambda,a}(\theta_i, \phi_i)$ be the fraction of $d^2\Phi_{\lambda,i}(\theta_i, \phi_i)$ that gets absorbed by dA . The directional spectral absorptivity $\alpha_\lambda(\theta_i, \phi_i)$ is defined as the ratio of $d^2\Phi_{\lambda,a}(\theta_i, \phi_i)$ to $d^2\Phi_{\lambda,i}(\theta_i, \phi_i)$ and is given by (2.12a). Here, the radiance invariance property was used to express all terms of the equation from the perspective of dA ; $d\Omega_i$ is the solid angle subtended by dA_i when viewed from dA . The definition of the directional total absorptivity $\alpha(\theta, \phi)$ is an extension of the definition of $\alpha_\lambda(\theta_i, \phi_i)$ that includes all wavelengths of the electromagnetic spectrum and is given by (2.12b).

$$\alpha_\lambda(\theta_i, \phi_i) = \frac{d^2\Phi_{\lambda,a}(\theta_i, \phi_i)}{d^2\Phi_{\lambda,i}(\theta_i, \phi_i)} = \frac{d^2\Phi_{\lambda,a}(\theta_i, \phi_i)}{L_{\lambda,i}(\theta_i, \phi_i)\cos\theta_i dA d\Omega_i} \quad (2.12a)$$

$$\alpha(\theta_i, \phi_i) = \frac{\int_0^\infty \alpha_\lambda(\theta_i, \phi_i) L_{\lambda,i}(\theta_i, \phi_i) d\lambda}{\int_0^\infty L_{\lambda,i}(\theta_i, \phi_i) d\lambda} \quad (2.12b)$$

If the area element dA is in **LTE**, the radiative power that it absorbs must be equal to the radiative power that it emits. This implies that $\varepsilon_\lambda(\theta, \phi)$ must be equal to $\alpha_\lambda(\theta, \phi)$. This relation is known as Kirchhoff's law of thermal radiation. The relation $\varepsilon(\theta, \phi) = \alpha(\theta, \phi)$, however, only holds for gray bodies.

Hemispherical absorptivities: α_λ and α

Let $d^2\Phi_\lambda$ be the spectral flux incident on an area element dA coming from all possible directions of a hemisphere centered on dA . And again, let $d^2\Phi_{\lambda,a}$ be the fraction of this spectral flux that gets absorbed by dA . The hemispherical spectral absorptivity α_λ is defined as the ratio of $d^2\Phi_{\lambda,a}$ to $d^2\Phi_\lambda$ and is given by (2.13a). The definition of the hemispherical total absorptivity α is an extension of the definition of α_λ that includes all wavelengths of the electromagnetic spectrum and is given by (2.13b).

$$\alpha_\lambda = \frac{d^2\Phi_{\lambda,a}}{d^2\Phi_\lambda} = \frac{\iint \alpha_\lambda(\theta_i, \phi_i) L_{\lambda,i}(\theta_i, \phi_i) \cos\theta_i d\Omega_i}{\iint_{2\pi} L_{\lambda,i}(\theta_i, \phi_i) \cos\theta_i d\Omega_i} \quad (2.13a)$$

$$\alpha = \frac{\int_0^\infty \iint_{2\pi} \alpha_\lambda(\theta_i, \phi_i) L_{\lambda,i}(\theta_i, \phi_i) \cos\theta_i d\Omega_i d\lambda}{\int_0^\infty \iint_{2\pi} L_{\lambda,i}(\theta_i, \phi_i) \cos\theta_i d\Omega_i d\lambda} \quad (2.13b)$$

The directional spectral absorptivity, the directional total absorptivity, the hemispherical spectral absorptivity and the hemispherical total absorptivity are dimensionless quantities.

Bidirectional reflectivities: $\rho_{\lambda, \text{bi-d}}(\theta_r, \phi_r, \theta_i, \phi_i)$ and $\rho_{\text{bi-d}}(\theta_r, \phi_r, \theta_i, \phi_i)$

Let $dE_{\lambda, i}(\theta_i, \phi_i)$ the spectral irradiance of an area element dA within the solid angle $d\Omega_i$ from the direction (θ_i, ϕ_i) . On the other hand, let $L_{\lambda, r}(\theta_r, \phi_r, \theta_i, \phi_i)$ be the spectral radiance reflected by dA in the direction (θ_r, ϕ_r) as a result of $dE_{\lambda, i}(\theta_i, \phi_i)$. The bidirectional spectral reflectivity $\rho_{\lambda, \text{bi-d}}(\theta_r, \phi_r, \theta_i, \phi_i)$ is defined as the ratio of $L_{\lambda, r}(\theta_r, \phi_r, \theta_i, \phi_i)$ to $dE_{\lambda, i}(\theta_i, \phi_i)$ and is given by (2.14a). The definition of the bidirectional total reflectivity $\rho_{\text{bi-d}}(\theta_r, \phi_r, \theta_i, \phi_i)$ is an extension of the definition of $\rho_{\lambda, \text{bi-d}}(\theta_r, \phi_r, \theta_i, \phi_i)$ that includes all wavelengths of the electromagnetic spectrum and is given by (2.14b).

$$\rho_{\lambda, \text{bi-d}}(\theta_r, \phi_r, \theta_i, \phi_i) = \frac{L_{\lambda, r}(\theta_r, \phi_r, \theta_i, \phi_i)}{dE_{\lambda, i}(\theta_i, \phi_i)} = \frac{L_{\lambda, r}(\theta_r, \phi_r, \theta_i, \phi_i)}{L_{\lambda, i}(\theta_i, \phi_i) \cos \theta_i d\Omega_i} \quad (2.14a)$$

$$\rho_{\text{bi-d}}(\theta_r, \phi_r, \theta_i, \phi_i) = \frac{\int_0^\infty \rho_{\lambda, \text{bi-d}}(\theta_r, \phi_r, \theta_i, \phi_i) L_{\lambda, i}(\theta_i, \phi_i) d\lambda}{\int_0^\infty L_{\lambda, i}(\theta_i, \phi_i) d\lambda} \quad (2.14b)$$

The fact that $\rho_{\lambda, \text{bi-d}}(\theta_r, \phi_r, \theta_i, \phi_i)$ has the term $\cos \theta_i d\Omega_i$ in its denominator allows to express the reflected spectral radiosity $dJ_{\lambda, r}(\theta_i, \phi_i)$ of an area element dA caused by the incident spectral radiance $L_{\lambda, i}(\theta_i, \phi_i)$ as the sum of the reflected radiances $L_{\lambda, r}(\theta_r, \phi_r, \theta_i, \phi_i)$ and also as a fraction of the spectral irradiance $dE_{\lambda, i}(\theta_i, \phi_i)$ according to (2.15a) and (2.15b), respectively.

$$dJ_{\lambda, r}(\theta_i, \phi_i) = \iint_{2\pi} L_{\lambda, r}(\theta_r, \phi_r, \theta_i, \phi_i) \cos \theta_r d\Omega_r \quad (2.15a)$$

$$\begin{aligned} &= L_{\lambda, i}(\theta_i, \phi_i) \cos \theta_i d\Omega_i \iint_{2\pi} \rho_{\lambda, \text{bi-d}}(\theta_r, \phi_r, \theta_i, \phi_i) \cos \theta_r d\Omega_r \\ &= dE_{\lambda, i}(\theta_i, \phi_i) \underbrace{\iint_{2\pi} \rho_{\lambda, \text{bi-d}}(\theta_r, \phi_r, \theta_i, \phi_i) \cos \theta_r d\Omega_r}_{\text{dimensionless}} \end{aligned} \quad (2.15b)$$

The bidirectional reflectivity is symmetric, i.e., $\rho_{\lambda, \text{bi-d}}(\theta_r, \phi_r, \theta_i, \phi_i) = \rho_{\lambda, \text{bi-d}}(\theta_i, \phi_i, \theta_r, \phi_r)$.

Directional-hemispherical reflectivities: $\rho_{\lambda, \text{d-h}}(\theta_i, \phi_i)$ and $\rho_{\text{d-h}}(\theta_i, \phi_i)$

Let $d^2\Phi_{\lambda, i}(\theta_i, \phi_i)$ be the spectral flux incident on an area element dA within the solid angle $d\Omega_i$ from the direction (θ_i, ϕ_i) . Let $d^2\Phi_{\lambda, r}(\theta_i, \phi_i)$ be the fraction of $d^2\Phi_{\lambda, i}(\theta_i, \phi_i)$ that gets reflected in all directions of a hemisphere centered on dA . The directional-hemispherical spectral reflectivity $\rho_{\lambda, \text{d-h}}(\theta_i, \phi_i)$ is defined as the ratio of $d^2\Phi_{\lambda, r}(\theta_i, \phi_i)$ to $d^2\Phi_{\lambda, i}(\theta_i, \phi_i)$ and is given by (2.16a). The definition of the directional-hemispherical total reflectivity $\rho_{\text{d-h}}(\theta_i, \phi_i)$ is an extension of the definition of $\rho_{\lambda, \text{d-h}}(\theta_i, \phi_i)$ that includes

all wavelengths of the electromagnetic spectrum and is given by (2.16b).

$$\begin{aligned}\rho_{\lambda, \text{d-h}}(\theta_i, \phi_i) &= \frac{d^2 \Phi_{\lambda, \text{r}}(\theta_i, \phi_i)}{d^2 \Phi_{\lambda, \text{i}}(\theta_i, \phi_i)} \\ &= \frac{\iint_{2\pi} L_{\lambda, \text{r}}(\theta_r, \phi_r, \theta_i, \phi_i) \cos \theta_r d\Omega_r dA}{L_{\lambda, \text{i}}(\theta_i, \phi_i) \cos \theta_i d\Omega_i dA} \\ &= \iint_{2\pi} \rho_{\lambda, \text{bi-d}}(\theta_r, \phi_r, \theta_i, \phi_i) \cos \theta_r d\Omega_r\end{aligned}\quad (2.16a)$$

$$\rho_{\text{d-h}}(\theta_i, \phi_i) = \frac{\int_0^\infty \rho_{\lambda, \text{d-h}}(\theta_i, \phi_i) L_{\lambda, \text{i}}(\theta_i, \phi_i) d\lambda}{\int_0^\infty L_{\lambda, \text{i}}(\theta_i, \phi_i) d\lambda}\quad (2.16b)$$

Due to the law of conservation of energy, the total radiative power incident on dA must either be absorbed or reflected. This implies that $\alpha_\lambda(\theta, \phi) + \rho_{\lambda, \text{d-h}}(\theta, \phi) = 1$ or, equivalently, $\varepsilon_\lambda(\theta, \phi) + \rho_{\lambda, \text{d-h}}(\theta, \phi) = 1$.

Hemispherical-directional reflectivities: $\rho_{\lambda, \text{h-d}}(\theta_i, \phi_i)$ and $\rho_{\text{h-d}}(\theta_i, \phi_i)$

Let $L_{\lambda, \text{i}}(\theta_i, \phi_i)$ be the spectral radiance incident on an area element dA coming from the direction (θ_i, ϕ_i) . Let $L_{\lambda, \text{r}}(\theta_r, \phi_r)$ be the spectral radiance reflected by the area element dA as a result of irradiating dA from all possible directions of a hemisphere centered on dA . The hemispherical-directional spectral reflectivity $\rho_{\lambda, \text{h-d}}(\theta_r, \phi_r)$ is defined as the ratio of $L_{\lambda, \text{r}}(\theta_r, \phi_r)$ to the average value of $L_{\lambda, \text{i}}(\theta_i, \phi_i)$ over the entire hemisphere and is given by (2.17a). The definition of the directional-hemispherical total reflectivity $\rho_{\text{h-d}}(\theta_r, \phi_r)$ is an extension of the definition of $\rho_{\lambda, \text{h-d}}(\theta_r, \phi_r)$ that includes all wavelengths of the electromagnetic spectrum and is given by (2.17b).

$$\rho_{\lambda, \text{h-d}}(\theta_r, \phi_r) = \frac{L_{\lambda, \text{r}}(\theta_r, \phi_r)}{(1/\pi) \iint_{2\pi} L_{\lambda, \text{i}}(\theta_i, \phi_i) \cos \theta_i d\Omega_i}\quad (2.17a)$$

$$\rho_{\text{h-d}}(\theta_r, \phi_r) = \frac{\int_0^\infty L_{\lambda, \text{r}}(\theta_r, \phi_r) d\lambda}{(1/\pi) \int_0^\infty \iint_{2\pi} L_{\lambda, \text{i}}(\theta_i, \phi_i) \cos \theta_i d\Omega_i d\lambda}\quad (2.17b)$$

It follows that when the incident spectral radiance is isotropic, $\rho_{\lambda, \text{d-h}}(\theta, \phi) = \rho_{\lambda, \text{h-d}}(\theta, \phi)$.

Specular reflectivities: $\rho_{\lambda, \text{s}}(\theta, \phi)$ and $\rho_{\text{s}}(\theta, \phi)$

The definition of the bidirectional reflectivities is not practical to describe specular surfaces, i.e., surfaces that reflect incoming radiance in a single direction according to Fresnel laws of reflection. In such cases, the definitions of specular spectral reflectivity

given by (2.18a) and specular total reflectivity given by (2.18b) are used instead.

$$\rho_{\lambda,s}(\theta, \phi) = \frac{L_{\lambda,r}(\theta, \phi + \pi)}{L_{\lambda,i}(\theta, \phi)} \quad (2.18a)$$

$$\rho_s(\theta, \phi) = \frac{L_r(\theta, \phi + \pi)}{L_i(\theta, \phi)} \quad (2.18b)$$

Due to the law of conservation of energy, the total radiative power incident on dA must either be absorbed or reflected. This implies that $\alpha_\lambda(\theta, \phi) + \rho_{\lambda,s}(\theta, \phi) = 1$ or, equivalently, $\varepsilon_\lambda(\theta, \phi) + \rho_{\lambda,s}(\theta, \phi) = 1$.

The bidirectional spectral reflectivity and the bidirectional total reflectivity have the unit sr^{-1} . The directional-hemispherical spectral reflectivity, hemispherical-directional spectral reflectivity, the directional-hemispherical total reflectivity, the hemispherical-directional total reflectivity, the specular spectral reflectivity and the specular total reflectivity are dimensionless quantities.

2.2.6 Radiative properties of opaque bodies

For optically smooth surfaces the radiative properties described in the previous section can be calculated using electromagnetic wave theory and the optical and electrical properties of the materials involved. A surface is deemed optically smooth when the imperfections or roughness of its texture are much smaller than the wavelength of the incident radiation. In practice, nonideal surface conditions can introduce large deviations from the values predicted by the theory, but the theory helps nonetheless to understand and predict, for example, the directional behavior of the emissivity, absorptivity and reflectivity. Two classes of materials are considered: dielectrics and metals. Figure 2.5 illustrates the directional dependence of the spectral emissivity with respect to the polar angle θ for the dielectric copper oxide (CuO_2) and the metal copper (Cu) at $\lambda = 10 \mu\text{m}$. The curves were plotted using the Fresnel equations presented in Appendix B.2 and the fact that $\varepsilon_\lambda(\theta, \phi) + \rho_{\lambda,s}(\theta, \phi) = 1$ for opaque bodies.

The plots show that dielectrics exhibit high emissivity values whereas metals exhibit very low emissivity values. They also show that these values remain constant for emission angles smaller than approximately 45°C . For angles greater than approximately 60°C , the emissivity drops sharply in the case of dielectrics and increases to some maximum value before dropping to zero in the case of metals. Refractive indices for different materials at different wavelengths can be found in [54].

For more information regarding the topics covered in this section the reader is referred to [32] and [86].

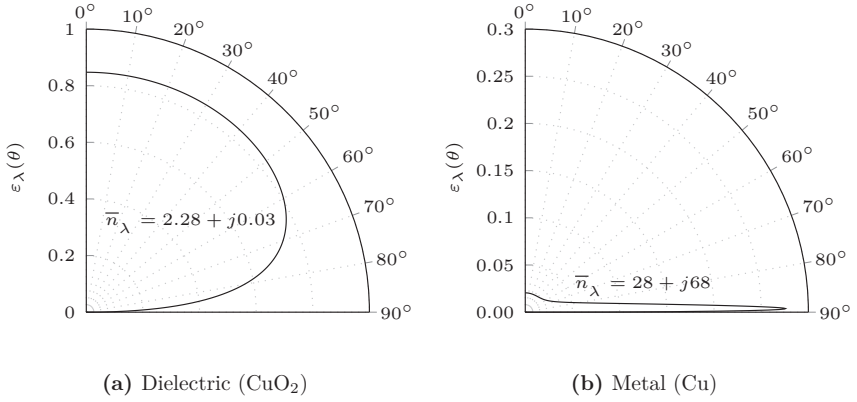


Figure 2.5: Directional spectral emissivities $\epsilon_\lambda(\theta)$

2.3 Optical systems

Optical lenses are needed to redirect the beams of electromagnetic radiation emitted or reflected by a real body towards a detector. The change of direction of the beam is caused by the change of the refractive index in the transmitting medium and is described by Snell's law (see Appendix B.1).

2.3.1 Thin lens

The thin lens is a special type of lens whose thickness along the optical axis is negligible compared to the radii of curvature of the lens surfaces. It is often used to illustrate the process of image formation because optical effects associated with the lens thickness can be neglected (see Figure 2.6).




Figure 2.6: Process of image formation using a symmetric thin lens

The object is assumed to be contained on a plane called the *object plane* and its image is formed on the other side of the lens on a plane called the *image plane*. The *focal point* refers to the point where rays traveling parallel to the optical axis would converge after crossing the lens. The distance f between the lens and the focal point is called the *focal length* and the plane that contains the focal point is called the *focal plane*. All the aforementioned planes are orthogonal to the optical axis. When different rays emitted by the same point on the object converge on a single point on the image plane, the distances shown in Figure 2.6 satisfy the relation given by (2.19). In the case where the object is very far away from the lens (i.e. $D \rightarrow \infty$) the image plane coincides with the focal plane. On the other hand, when (2.19) does not hold, the emitted rays form a circular spot on the image plane known as the *circle of confusion*.

$$\frac{1}{f} = \frac{1}{D} + \frac{1}{D'} \quad (2.19)$$

2.3.2 Depth of field

In practice, the detector on the opposite side of the lens has finite dimensions and the amount of radiation that reaches it is constrained by the diameter of the lens or by a fixed or adjustable opening called *aperture* that is typically located between the lens and the image plane. Under such conditions, an image appears as sharp as possible on the image plane when the circle of confusion does not exceed the size of the detector. This scenario is depicted in Figure 2.7, where, for simplicity, the diameter of the aperture equals the diameter of the lens. These conditions impose a limit on the nearest and farthest distance between the lens and the object which are given by (2.20a) and (2.20b). The depth range delimited by D_F and D_N is known as the *depth of field (DOF)* and is given by (2.20c). In all three equations, c is the diameter of the circle of confusion and $N = f/d$ is the ratio of the focal length f to the aperture diameter d . N is known as the *f-number* or the *f-stop*. The distance D for which $D_F \rightarrow \infty$ is called the hyperfocal distance.




Figure 2.7: Depth of field

$$D_N = \frac{Df^2}{f^2 + cN(D - f)} \quad (2.20a)$$

$$D_F = \frac{Df^2}{f^2 - cN(D - f)} \quad (2.20b)$$

$$DOF = \frac{2Df^2cN(D - f)}{f^4 - (cN(D - f))^2} \quad (2.20c)$$

2.3.3 Field of view

The *field of view* (FOV) describes the extent of the world seen by a detector through the camera's optical system. Because the majority of detectors arrays have a rectangular shape, the FOV is typically expressed as the product of two angles that denote the *horizontal field of view* (HFOV) and *vertical field of view* (VFOV) (see Figure 2.8).




Figure 2.8: Field of view

The HFOV and VFOV are related to the focal length f and to the detector array's width w and height h , by (2.21a) and (2.21b) respectively. Furthermore, in remote sensing, the FOV subtended by a single detector element in a detector array is called *instantaneous field of view* (IFOV).

$$HFOV = 2 \tan^{-1} \left(\frac{w}{2f} \right) \quad (2.21a)$$

$$VFOV = 2 \tan^{-1} \left(\frac{h}{2f} \right) \quad (2.21b)$$

2.3.4 Optical aberrations

Real lenses introduce optical aberrations that cause image blurring and distortion. These aberrations fall into two categories: chromatic and monochromatic. Chromatic aberrations are caused by the wavelength dependence of the lens' refractive index. This results in rays of different wavelengths being brought to focus at different distances from the lens and the optical axis. Monochromatic aberrations, also known as geometric aberrations, are caused by the shape of the lens and its alignment in the optical system. In this case, rays emitted by an object converge at different distances from the lens and the optical axis. The main types of monochromatic aberrations are: spherical aberration, coma, astigmatism, field curvature and distortion. The first four types of aberrations can be minimized through proper optical lens design. The last type, distortion, is easier and cheaper to correct via image processing algorithms.

For more information regarding optical systems the reader is referred to [29] and [72].

2.3.5 Pinhole camera model

The pinhole camera model is a simple mathematical model that describes the correspondence between points in the real world and their projections on the image plane in a pinhole camera. This is illustrated in Figure 2.9.




Figure 2.9: Pinhole camera model

Points in the real world are projected onto the image plane using a series of projective transformations that map the point from one coordinate system onto another. These transformations are given by (2.22a), (2.22b) and (2.22c).

$$\begin{bmatrix} X_C & Y_C & Z_C \end{bmatrix}^T = \begin{bmatrix} \mathbf{R}_{3 \times 3} & \mathbf{t}_{3 \times 1} \end{bmatrix} \begin{bmatrix} X_W & Y_W & Z_W & 1 \end{bmatrix}^T \quad (2.22a)$$

$$\begin{bmatrix} x & y & 1 \end{bmatrix}^T = \frac{1}{Z_C} \cdot \mathbf{K}'_{3 \times 3} \begin{bmatrix} X_C & Y_C & Z_C \end{bmatrix}^T \quad (2.22b)$$

$$\begin{bmatrix} u & v & 1 \end{bmatrix}^T = \mathbf{K}''_{3 \times 3} \begin{bmatrix} x & y & 1 \end{bmatrix}^T \quad (2.22c)$$

where

$$\left[\mathbf{R}_{3 \times 3} \mid \mathbf{t}_{3 \times 1} \right] = \begin{bmatrix} R_{11} & R_{12} & R_{13} & t_{11} \\ R_{21} & R_{22} & R_{23} & t_{21} \\ R_{31} & R_{32} & R_{33} & t_{31} \end{bmatrix} \quad (2.22d)$$

$$\mathbf{K}'_{3 \times 3} = \begin{bmatrix} f_x & \alpha & 0 \\ 0 & f_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (2.22e)$$

$$\mathbf{K}''_{3 \times 3} = \begin{bmatrix} k_u & 0 & c_u \\ 0 & k_v & c_v \\ 0 & 0 & 1 \end{bmatrix} \quad (2.22f)$$

Equation (2.22a) describes the relationship between two sets of coordinates that represent the same point P. $[X_W \ Y_W \ Z_W \ 1]^T$ indicates the position of P with respect to an arbitrary world coordinate system. $[X_C \ Y_C \ Z_C]^T$ indicates the position of P with respect to the camera's coordinate system with origin at the camera's optical center and whose z-axis is aligned with the camera's optical axis. The augmented matrix $\left[\mathbf{R}_{3 \times 3} \mid \mathbf{t}_{3 \times 1} \right]$ is called the *extrinsic camera matrix* and its components, the rotation matrix $\mathbf{R}_{3 \times 3}$ and the translation vector $\mathbf{t}_{3 \times 1}$, describe the geometric transformation between the world and camera coordinate systems. Equation (2.22b) projects the point P (expressed in terms of the camera coordinate system) onto the image plane. Here, f_x and f_y represent scaling factors and α represents a skew coefficient. Usually, $\alpha = 0$ because the pixels have a rectangular shape. Furthermore, unless the lens introduces unintentional or intentional distortions, $f_x = f_y = f$, where f represents the distance between the image plane and the camera's optical center. Finally, (2.22c) transforms P' , the projection of point P, to pixel coordinates $[u \ v]^T$. Here, k_u and k_v represent the inverse of the detector element's dimensions (in pixels per unit length) and $[c_u \ c_v]^T$ denotes the position (in pixels and with respect to the image corner) of the point where the optical axis intersects the image plane.

The product of the matrices $\mathbf{K}'_{3 \times 3}$ and $\mathbf{K}''_{3 \times 3}$ given by (2.23a) is called the *intrinsic camera matrix* and its elements are known as the *intrinsic parameters* of the camera. Here, $f_u = k_u f_x$, $f_v = k_v f_y$ and $\alpha' = k_u \alpha$.

$$\mathbf{K}_{3 \times 3} = \mathbf{K}''_{3 \times 3} \mathbf{K}'_{3 \times 3} = \begin{bmatrix} f_u & \alpha' & c_u \\ 0 & f_v & c_v \\ 0 & 0 & 1 \end{bmatrix} \quad (2.23a)$$

To avoid working with inverted images, the image plane is often assumed to be located between the object and the camera's optical center. This plane is called the *virtual image*

plane. It is at the same distance from the camera's optical center as the image plane and it is parallel to it.

The pinhole camera model describes an ideal camera, where the projection of points onto the image plane is rectilinear. Images produced by real cameras, however, are subject to distortions introduced by the optical system. The two most common forms of distortions are radial and decentering distortions. The former is caused by the shape of the lens and the position of the aperture within the optical system. The latter occurs when the optical axis is not exactly perpendicular to the image plane, for example, due to alignment errors during the lens assembly. As a result of these distortions, a point in space, its projection onto the image plane and the camera's optical center are not always collinear. The model most widely used to describe these distortions is the Brown-Conrady model. Equations (2.24a) and (2.24b) describes the offsets caused by the radial and decentering distortions based on this model. The distortion parameters $[k_1, k_2, p_1, p_2, k_3]$, commonly given in this order, are considered part of the camera's intrinsic parameters.

$$\delta_x(x, y) = \underbrace{x(k_1 r^2 + k_2 r^4 + k_3 r^6)}_{\text{radial distortion}} + \underbrace{(p_1(3x^2 + y^2) + 2p_2 xy)}_{\text{decentering distortion}} \quad (2.24a)$$

$$\delta_y(x, y) = \underbrace{y(k_1 r^2 + k_2 r^4 + k_3 r^6)}_{\text{radial distortion}} + \underbrace{(p_2(x^2 + 3y^2) + 2p_1 xy)}_{\text{decentering distortion}} \quad (2.24b)$$

$$r^2 = x^2 + y^2 \quad (2.24c)$$

Equation (2.25) relates the position of the distorted projected point $[x_d \ y_d \ 1]^T$ that results from adding these offsets to the ideal projected point $[x \ y \ 1]^T$.

$$\begin{bmatrix} x_d & y_d & 1 \end{bmatrix}^T = \begin{bmatrix} x & y & 1 \end{bmatrix}^T + \begin{bmatrix} \delta_x(x, y) & \delta_y(x, y) & 0 \end{bmatrix}^T \quad (2.25)$$

For more information regarding the pinhole camera model and the Brown-Conrady model the reader is referred to [52].

2.4 Summary

This chapter explains the main theoretical concepts needed to understand the fundamentals of infrared thermography as well as the terminology employed in this field. Section 2.1, radiometric quantities, presents a summary of the main radiometric quantities used to describe the flux of radiant electromagnetic power. Section 2.2, electromagnetic radiation, explains the relationship between the temperature of an object's surface and

the amount of electromagnetic radiation that it emits and how this radiation propagates in a medium. It also explains various concepts that frequently come up in the field of infrared thermography like black bodies and real bodies as well as the main properties associated with them: emissivity, absorptivity, reflectivity and transmissivity. In order to describe the directional nature of these properties and to understand how exactly they relate to each other, rigorous definitions are provided. Section 2.3, optical systems, briefly explains the concepts of focal length, image plane, depth of field and field of view which are used to describe most imaging systems. This section includes a description of the pinhole camera model, a mathematical model that relates points in space with their projections on the image plane. This model is used to describe the depth and thermographic cameras of the 3D thermographic system presented in this work.

Chapter 3

2D Thermal imaging

The objective of a 2D thermal imaging system is to produce 2D images that depict the surface temperature distribution of the observed object. The generated image is called a thermogram or simply a thermal image and the device that produces these images is known as thermal imager, thermal camera, thermographic camera or infrared camera. The works of [21], [45] and [40] provide good overviews of current applications of thermal cameras.

This chapter uses the concepts introduced in the previous chapter to explain the principle of operation of a conventional thermographic camera and describes the main issues that affect the fidelity of conventional thermal images.

3.1 Radiometric chain

The measurement process starts with the emission of thermal radiation by the observed object. This radiation is not necessarily the only radiation that reaches the detector of the thermographic camera. Other sources of radiation like surrounding objects and the camera housing also contribute to the radiant flux measured by the detector. Moreover, the radiation coming from the object's surface suffers gains and losses as it propagates through the transmission medium and the camera lens (see section 2.2.4). The ensemble of all these phenomena is known as the *radiometric chain* and is depicted in Figure 3.1.

For certain spectral bands, when the propagation medium is clean air and the distance between the object and the thermographic camera is only a few meters, the effects of atmospheric attenuation, scattering and emission can be neglected. Figure 3.2 shows the atmospheric spectral transmittance at short distances. Similarly, if the optical lens




Figure 3.1: Radiometric chain

used in thermal imaging system has very high transmittance, the effects of attenuation, scattering and emission within the lens can also be neglected.

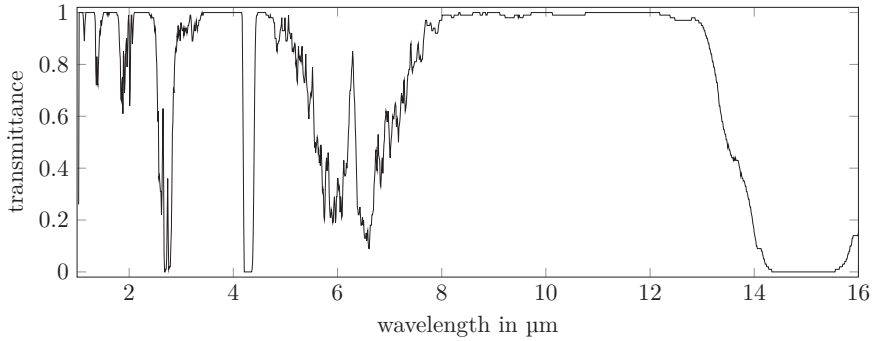


Figure 3.2: Atmospheric spectral transmittance based on the HITRAN database [6] at 25 °C, 1 atm, 50 % relative humidity and for a path length of 10 m

Under such conditions, the spectral radiance $L_{\lambda,\text{obj}}(\theta, \phi)$ leaving an area element dA of an opaque body at temperature T_{obj} that reaches the detector in the thermographic camera is given by (3.1), where $L_{\lambda,\text{sur}}(\theta_i, \phi_i)$ is the spectral radiance incident on dA from the direction (θ_i, ϕ_i) radiated by objects in the surroundings.

$$L_{\lambda,\text{obj}}(\theta, \phi) = \underbrace{\varepsilon_{\lambda}(\theta, \phi) L_{\lambda,\text{bb}}(T_{\text{obj}})}_{\text{radiance emitted by object}} + \underbrace{\int_{2\pi} \rho_{\lambda,\text{bi-d}}(\theta, \phi, \theta_i, \phi_i) L_{\lambda,\text{sur}}(\theta_i, \phi_i) \cos(\theta_i) d\Omega_i}_{\text{radiance reflected by object}} \quad (3.1)$$

Furthermore, if the object under inspection has high emissivity (and therefore low reflectivity) or if T_{obj} is much higher than the temperature of the surroundings, the second

summand in (3.1) can be neglected. This is a very common assumption in thermal imaging, however, this assumption is only valid when the aforementioned conditions hold.

Assuming that the spectral radiance of the object $L_{\lambda,\text{obj}}$ and spectral radiance of the camera housing $L_{\lambda,\text{cam}}$ are uniform within the solid angles subtended by them when viewed from the detector, the irradiance E_{det} measured by the detector is then given by (3.2).

$$E_{\text{det}} = \underbrace{\Omega_a \int_{\lambda_1}^{\lambda_2} L_{\lambda,\text{obj}} d\lambda}_{\text{irradiance from object}} + \underbrace{(\pi - \Omega_a) \int_{\lambda_1}^{\lambda_2} L_{\lambda,\text{cam}} d\lambda}_{\text{irradiance from camera housing}} \quad (3.2)$$

Here, the limits λ_1 and λ_2 represent the operating spectral range of the detector and Ω_a is the projected solid angle subtended by the camera aperture. The value of Ω_a depends on the position of the detector with respect to the lens and the optical axis but is always smaller than its hemispherical complement $(\pi - \Omega_a)^1$. This means that the temperature of the camera housing has a strong influence on the total irradiance measured by the detector, in particular, when temperature of the object is not much greater than the temperature of the camera housing.

3.2 Thermographic camera

The thermographic camera measures the thermal radiation leaving an object's surface in the direction of the camera and produces images that depict the object's surface temperature distribution. The thermal radiation measured by each element of the camera's detector array is converted to a temperature value using the theoretical concepts and assumptions discussed so far and taking into account the emissivity of the object's surface, which is given as input parameter to the camera. The color assigned to each image pixel is determined based on the temperature estimated from the measurement of the corresponding detector, a predefined color scale and the scale limits, which can be set manually or computed automatically according to some predefined criterion.

3.2.1 Basic components

The main components of a thermographic camera are depicted in the block diagram in Figure 3.3.

¹ The hemispherical solid angle is 2π but the projected hemispherical solid angle $\iint_{2\pi} \cos\theta d\Omega$ equals π .



Figure 3.3: Main building blocks of a thermographic camera

Optics

Thermographic cameras require an optical lens to focus the incoming beams of electromagnetic radiation onto the detector. Common glass materials used in conventional cameras optics, have very poor transmittance in the infrared spectrum. For this reason, other types of materials are used in thermal imaging systems. Cameras operating in the [LWIR](#) band, for example, are typically equipped with lenses made of germanium and covered with an antireflex coating to improve its transmittance. Furthermore, thermographic cameras are designed to have a low f-number to let in as much radiation as possible. This practice helps to improve the *signal to noise ratio* (SNR) albeit at the expense of having a reduced [DOF](#).

Detector

In order to transform electromagnetic radiation into a useful electrical signal that can be amplified, filtered, digitized and digitally processed, a detector is required. For the measurement of infrared radiation, two classes of detectors can be found: photon detectors and thermal detectors. In photon (or quantum) detectors, incident photons with sufficient energy change the concentration of free charge carriers on the detector element. Depending on the type of photon detector, this change may alter the conductivity of the detector element or induce a photoelectric current through it. Photon detectors require substantial cooling (typically down to 77 K) in order to reduce the concentration of free charge carriers that originate from thermal excitation at room temperature. Examples of photon detectors are photoconductors, photodiodes and *quantum well infrared photodetectors* (QWIP). In thermal detectors, on the other hand, the absorbed radiation changes the temperature of the detector element. In a similar manner, depending on the type of thermal detector, this temperature change may alter the electrical resistance of the detector's material or induce a voltage across it. Examples of thermal detectors are thermopiles, pyroelectric detectors and bolometers. In addition, detectors differ from each other in terms of their spectral responsivity, response time, noise level and size. The work of [\[74\]](#) elaborates on recent progress in infrared detector technologies.

The detector array in today's infrared cameras is arranged in a (typically rectangular) matrix called the *focal plane array* (FPA). Photon detector FPA cameras offer better performance than the thermal detector FPA cameras in terms of thermal and temporal resolution, but are also bulkier and much more expensive because of the need of a

cooling system. At present, the majority of thermographic cameras are equipped with a microbolometer **FPA** because their performance meets the demands of most practical applications. The microbolometer **FPA** does not require cooling but in most cases its temperature is thermoelectrically stabilized by means of a Peltier element.

Each detector element of an **FPA** has a slightly different sensitivity (gain) and bias point (offset). To obtain the same signal value from all detectors when they are irradiated by the same radiant power, gain and offset correction values must be calculated for and applied to the signal of each detector element. This process is called *nonuniformity correction* (**NUC**). The **NUC** can be carried out by placing a gray body of high emissivity at uniform temperature in front of the **FPA** and measuring the detector's signal for two different temperature values. This is known as a two-point **NUC**. The procedure is normally carried out by the manufacturer and the calibration data are stored in the firmware of the camera.

Signal conditioning and processing

The electrical signal produced by the detector array must be read out from the matrix, amplified, filtered and digitized before it can be digitally processed. Dedicated electronic circuits care for these tasks. Once in digital form, the signal from each detector is corrected for changes in the operating conditions and then converted into a temperature value.

The detectors in the array measure not only the radiation from the observed object but also the radiation emitted by the camera housing and the optics as indicated by (3.2). Since the temperature of the camera housing may change over time, signal corrections need to be carried out on a regular basis. This is especially true for microbolometer **FPA** cameras, whose responsivity strongly depends on the temperature of the detectors. For this reason, most microbolometer **FPA** cameras are equipped with a shutter to carry out a one point offset **NUC** during normal camera operation. The shutter has a uniform temperature and is moved in front of detector for a brief period of time. This technique provides a reasonable correction in the vicinity of the calibration point.

The conversion of the detector's signal into a temperature value is carried out using data obtained from the radiometric calibration of the camera. During the radiometric calibration, a black body is placed in front of the thermographic camera such that it covers its entire **FOV**. The distance between the black body and the camera is kept short in order to have almost ideal atmospheric transmittance. The relation between the black body's temperature and the detector's signal is determined for different temperature values of the black body. The data obtained during the calibration take into account the spectral responsivity of the detector and the spectral transmittance of the lens. The

accuracy of the calibration can be further improved when the camera temperature is also taken into consideration. Factory calibrated cameras store the calibration data in the firmware. Because the calibration is done for a black body, during normal operation, the emissivity of the real body must be entered as parameter to obtain the best possible estimation of its surface temperature. A more detailed description of the calibration steps of microbolometer array based thermal infrared cameras is given in [12].

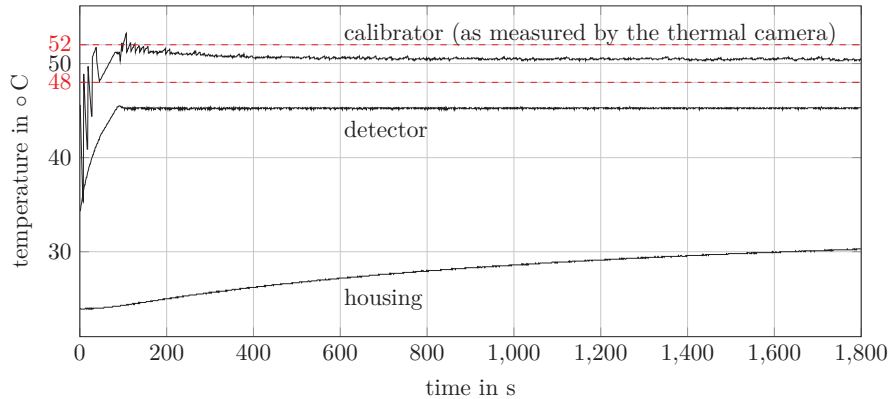
3.2.2 Characterization

According to [86] the main parameters that describe the characteristics and performance of a thermographic camera can be divided into two groups: objective and subjective parameters. The first group includes design parameters and other parameters that can be calculated or measured through a series of experiments. Detector type, size and pitch, spectral range, temperature range(s), temperature accuracy, thermal resolution, frame rate, integration time, FOV, focal length and f-number belong to this group (see Table C.3 in Appendix C). The second group includes parameters that describe the ability of a human observer to detect and recognize temperature differences in a thermal image. These parameters are closely related to the objective parameters that describe thermal and spatial resolution but also take other factors into consideration, for example, the choice of the color palette. The *minimum resolvable temperature difference* (MRTD) and the *minimum detectable temperature difference* (MDTD) belong to this group. A more thorough description of these and other parameters can be found in [11] and [86].

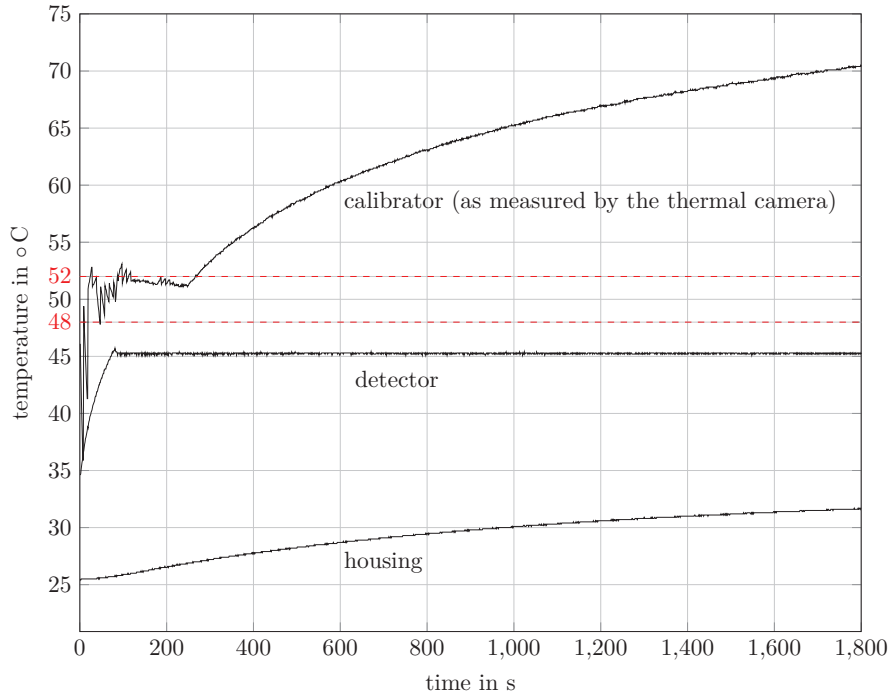
3.2.3 Switch-on behavior

As stated in (3.2), the temperature of the camera housing has a strong influence on the radiative power measured by the detectors. In the case of microbolometer FPA cameras, the temperature of the camera housing also affects the responsivity of the detectors. In order to get accurate temperature measurements, the camera has to be in thermal equilibrium. It takes the camera several minutes to reach thermal equilibrium after it is switched on. Furthermore, the camera's thermal equilibrium might be affected by sudden changes in the ambient temperature. To guarantee the measurement lies within the accuracy range of the camera, a NUC operation must be carried out regularly. The more stable the operation conditions, the less frequent a NUC operation must be carried out. Figure 3.4a illustrates a long-term temperature measurement of a FLUKE 4180 precision infrared calibrator set to 50 °C as recorded by an Optris PI 450 thermal camera. The temperature of the detector chip stabilizes approximately two minutes after starting the thermal camera. The measurement of the calibrator's temperature made by

the detector at the center of the thermal camera's FPA falls within the range $50^{\circ}\text{C} \pm 2^{\circ}\text{C}$ about one minute later. The internal housing temperature takes much longer to stabilize and periodic NUC operations are needed to maintain the accuracy of the measurement. Figure 3.4b shows what happens to the temperature measurement if the NUC operation is disabled after the temperature measurement falls within the accuracy range of the thermal camera, but the camera housing has not reached thermal equilibrium.



(a) NUC operation periodically executed



(b) NUC operation disabled 4 minutes after starting the camera

Figure 3.4: Temperature measurements during the camera's warm up phase

In this case, a 5 °C increase in the camera housing temperature causes a 20 °C error in the estimation of the calibrator's temperature.

3.3 Thermal image

The thermal image produced by the thermographic camera depicts the surface temperature distribution of the observed object. Different shades of gray or colors from a color palette are used to distinguish the different temperature values. The accuracy of the temperature measurements that make up the thermal image and the ability to recognize an object and its features depend on the characteristics of the thermographic camera and on external factors. This section illustrates the effects that these factors have on the accuracy and discernibility of the thermal images.

3.3.1 Accuracy

In addition to the thermographic camera's own accuracy limitations, several external factors affect the accuracy of the temperature measurements carried out with a thermographic camera. Some of these factors are inherent to the material of the object being inspected; others can be controlled by the person carrying out the inspection.

Emissivity

According to (3.1), the emissivity $\varepsilon_\lambda(\theta, \phi)$ of a surface has a direct influence on the temperature value measured by the thermographic camera. This parameter is not measured by the camera; it must be input manually by the operator via the camera's user interface. In practice, information regarding the radiative properties of a material is limited. Approximate emissivity values for materials commonly encountered in inspection tasks can be found in books on infrared thermometry as well as in documentation provided by manufacturers of infrared thermal sensors and cameras. These are average values estimated for specific wavelength ranges and observed orthogonally to the surface. However, as shown in Figure 2.5, the emissivity of a material has a directional nature and this behavior must be taken into account. Figure 3.5 shows how the temperature of the FLUKE 4180 precision infrared calibrator as measured by an Optris PI 450 infrared camera decreases as the angle θ between the surface normal and the camera optical axis increases. The calibrator's disc has an emissivity of 0.95, its temperature is set to 100 °C, the room temperature is roughly 23 °C and the thermal camera is located approximately 50 cm away from the calibrator. The temperature value measured by the thermographic camera decreases because the surface's directional-hemispherical reflectivity increases

as its directional emissivity decreases and because the temperature of the surroundings is much smaller than that of the calibration disc. As will be explained later, part of the temperature drop is also due to the fact that less pixels of the camera's FPA are irradiated by the calibration disc and more are irradiated by the background which is at lower temperature. However, the temperature deviation caused by this effect is much smaller than the deviation caused by the drop in emissivity.

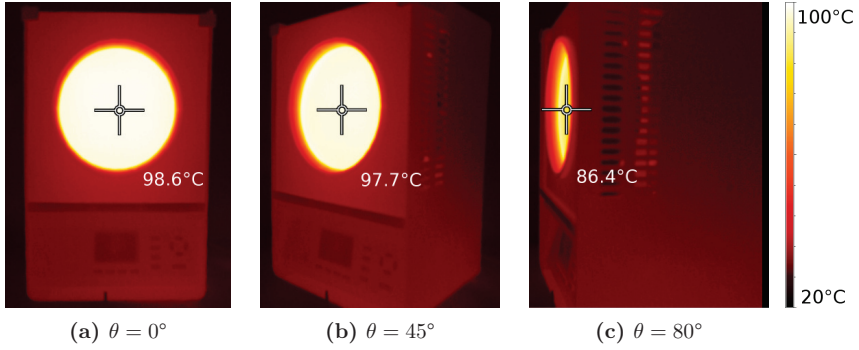


Figure 3.5: Dependence of temperature measurement on angle of observation

Reflectivity

According to (3.1), the reflectivity $\rho_\lambda(\theta, \phi, \theta_i, \phi_i)$ also has a direct influence on the temperature value measured by the thermographic camera. When an object does not exhibit high emissivity and its temperature is not much higher than that of its surroundings, the second summand in (3.1) introduces considerable errors in the temperature estimation. Furthermore, according to section 2.2.5, the lower the emissivity of the object, the higher its reflectivity and, consequently, the greater the measurement errors. Figure 3.6 shows the reflection of the infrared calibrator on plates of different materials. The plates are $10\text{ cm} \times 10\text{ cm}$ in size and are made of: a) glass, b) copper and c) weathered titanium zinc, respectively. The calibrator's temperature was set to 50°C and the temperature of the plates was the same as the room temperature, approximately 23°C . A cardboard with a circular hole of 67 mm in diameter was placed directly in front of the calibrator to produce a reflection that is easy to recognize. The glass plate produces a faint specular reflection while the copper plate behaves almost like a mirror. The weathered titanium zinc plate produces a strong but diffuse reflection. The intensity and direction of the reflections is governed by $\rho_\lambda(\theta, \phi, \theta_i, \phi_i)$. In this example, the reflections are easy to recognize; the temperature profile of the reflected object is known, the reflections are at most slightly blurred and the flat geometry of the plates does not cause geometric distortions of the reflection. This is not always the case. In practice, most surfaces are curved and not optically smooth. For these reasons, it is often difficult to recognize reflections.




Figure 3.6: Reflections on plates of different material

Focus

As with conventional color cameras, objects that lie outside the thermal camera's **DOF** look blurred in the thermal image. In this case, the temperature measurement made by one pixel (i.e. one detector element) gets averaged with the temperature measurements made by the neighboring pixels. This can be appreciated in Figure 3.7. The figure shows two images of the infrared calibrator: one correctly focused and one out of focus. A resolution pattern made of cardboard was placed in front the calibrator. The temperature of the calibration disc was set to 50 °C. The temperature of resolution pattern was slightly above the room temperature (approximately 23 °C) because it was warmed up by the calibration disc. The temperature value measured at the center of a square hole in the resolution pattern in the unfocused image is significantly lower than the value measured at the same point in the focused image.




Figure 3.7: Effect of focusing on temperature measurement

Distance to object

The distance d between thermographic camera and object affects the smallest surface element for which an accurate temperature measurement can be made. To illustrate this behavior, a cardboard mask with a 4 mm \times 4 mm opening was placed in front of the infrared calibrator and measurements of the hottest spot in the thermal image were

made for two different values of d . The temperature of the calibration disc was set to 50 °C. The room temperature was approximately 23 °C. Figure 3.8 shows a drop in the temperature reading as a result of increasing the distance d to the calibrator from 50 cm to 100 cm. According to [63], the IFOV of the Optris PI 450 is 1.1 mm at 50 cm and 2.3 mm at 100 cm. However, the smallest area for which the camera can accurately estimate its temperature value depends not only on the IFOV but it is also affected by diffraction of the optics and the fact that the projection of the area might not exactly coincide with a single detector element. To guarantee the most accurate results, the smallest area should be at least three times the size of the IFOV.




Figure 3.8: Effect of distance to object on temperature measurement

Movement

The electrical signal produced by a detector element in an FPA requires a certain amount of time to reach a new steady state after the incident radiation changes. For microbolometer based FPAs, typical response times are in the order of 10 ms. If the irradiance of the detector does not remain constant during this time interval, for example, due to motion of the thermal camera or the observed object, the measurement gets corrupted. In addition, most thermographic cameras read out the electrical signals from the FPA one row at a time. If the camera or object movement is too fast, objects in the thermal image looked skewed. These two effects are illustrated in Figure 3.9b.




Figure 3.9: Effect of camera movement on temperature measurement

Background radiation

The total amount of radiation that enters the thermal camera also affects the accuracy of temperature measurements. Figure 3.10 presents six different measurement setups in which the temperature near the center of the calibration disc is measured with the thermal camera.

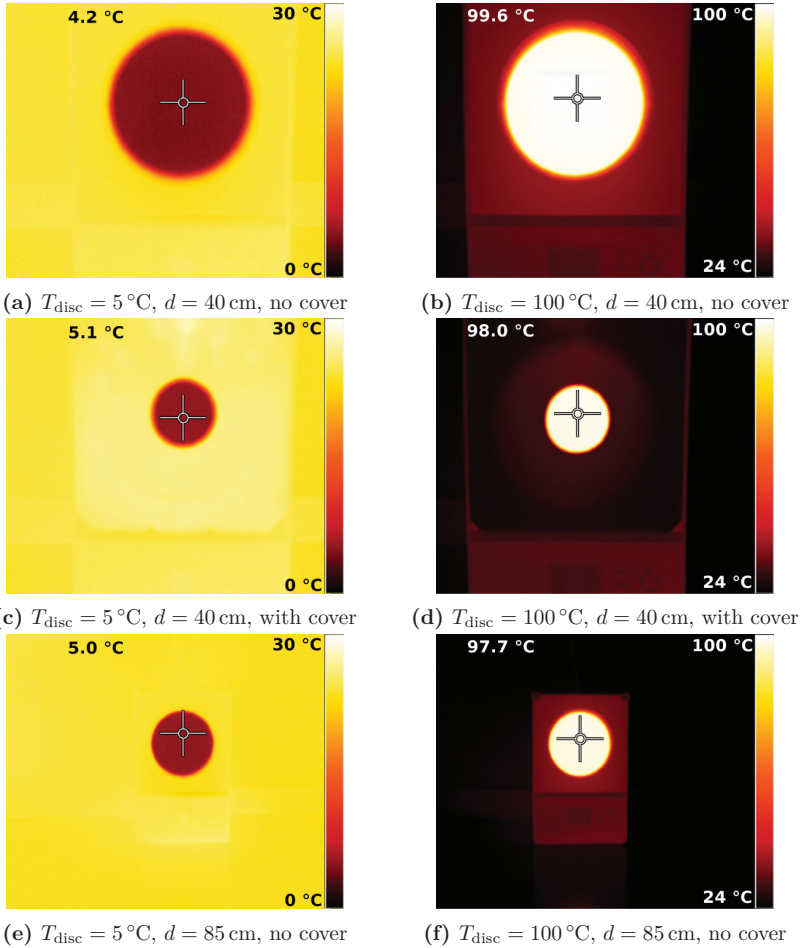


Figure 3.10: Effect of thermal background radiation on temperature measurement

The temperature of the calibrator disc was set to 5°C and 100°C for the setups on the left and right column, respectively. The images in the first and second row were taken from a distance of 40 cm but the latter were captured shortly after placing a cardboard plate with a circular hole of 67 mm in diameter at room temperature in front of the calibration disc. Immediately after inserting the cardboard plate, the temperature measurement increased when the calibration disc temperature was 5°C and decreased

when it was 100 °C. The images in the third row were taken from a distance of 85 cm. From this distance, the number of pixels depicting the disc and background temperatures is roughly the same as those in the second row, that is, the total amount of radiation that enters the thermal camera is approximately the same. The temperature values obtained in these cases were very close to those obtained in the scenarios shown in the second row. It is clear that the total amount of radiation that enters the camera affects the accuracy of the temperature measurements. In these experiments, the angle of observation was close to zero, the camera was appropriately focused and the **IFOV** was much smaller than the area of the calibration disc.

3.3.2 Discernibility

The choice of the color palette and the presence of large enough temperature differences in the observed scene have a strong influence on the discernibility of object features in the thermal image. This is illustrated in Figure 3.11, where four different measurement setups of the same scene are presented. In the top and bottom rows, the difference ΔT between the maximum and minimum temperature in the scene is approximately 10 °C and 2 °C, respectively.

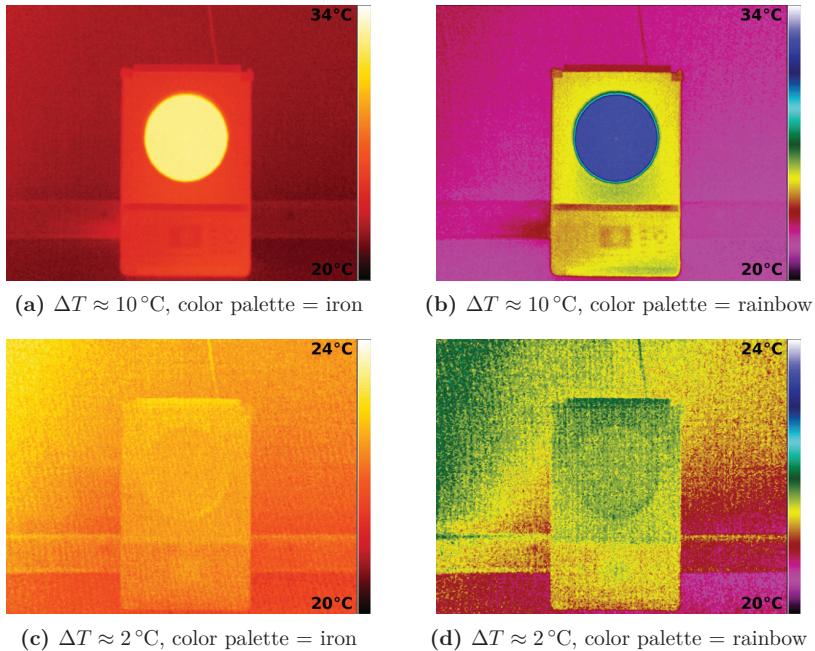


Figure 3.11: Discernibility of objects in thermal images

The images on the left column are false colored with an 'iron' palette and those on the right column with a 'rainbow' palette. In the top row, the 'rainbow' palette, with its frequently alternating colors, helps highlight temperature differences in the thermal image; object features, such as the control buttons and display, are easier to spot than with the 'iron' palette. In the bottom row, the temperature differences are almost zero, yet some shapes can be appreciated due to the different emissivities of the surfaces and the high thermal sensitivity of the camera.

3.4 Summary

This chapter explains the principle of operation of a conventional thermographic camera. It helps understand when the simplifications and assumptions that are frequently made in conventional thermography are justified. It also describes the main factors that affect the fidelity of conventional thermal images. Section 3.1, radiometric chain, explains how the thermal radiation emitted by different sources combine before reaching the detector in the thermographic camera. Section 3.2, thermographic camera, describes the basic components of a thermographic camera and summarizes the main parameters used to characterize it. This section also discusses the camera's switch-on behavior. It shows the strong influence that the temperature of the camera housing has on the accuracy of temperature measurements and the importance of carrying out a [NUC](#) when the temperature of the camera housing deviates too much from the last calibration point. Section 3.3, thermal image, shows, by means of simple experiments, additional factors that affect the accuracy of the temperature measurements and the ability of the user to discern features in the thermal image.

Chapter 4

3D thermal imaging: overview

The objective of 3D thermal imaging is to produce texturized 3D models that depict the surface temperature distribution of the observed object. The resulting texturized 3D models are called 3D thermograms and the device that produces these 3D thermograms is known as 3D thermal imaging system or 3D thermographic system. There are many ways to calibrate and design a 3D thermographic system. The advantages and disadvantages of the different approaches will be discussed in chapters 5 and 6. This chapter introduces the main building blocks of the 3D thermal imaging system developed and tested in this work.

4.1 System

The 3D thermal imaging system used in this work consists of an Optris PI 450 infrared camera, a PrimeSense Carmine 1.09 depth-sensing camera and a laptop with a Intel® Core™ i7-4980HQ *central processing unit* (CPU) and an Nvidia® GeForce® GTX 980M *graphics processing unit* (GPU). The thermographic camera produces thermal images with a size of 382×288 pixels at a rate of 40 fps. It operates in the spectral range $7.5 \mu\text{m}$ to $13 \mu\text{m}$, has a FOV of $62^\circ \times 49^\circ$, an accuracy of $\pm 2^\circ\text{C}$ or $\pm 2\%$ of the reading (whichever is greater) and has been radiometrically calibrated at the factory. The depth-sensing camera generates depth images in normal operation with a size of 640×480 pixels at a rate of 30 fps. It has a FOV of $57.5^\circ \times 45^\circ$, an operating range of 0.35 m to 1.4 m and a spatial resolution of 1 mm at a distance of 0.5 m. The sensor detects surfaces beyond 1.4 m but the precision and the spatial resolution of the depth measurements decrease significantly. Both sensors are connected to the laptop via USB 2.0. More information regarding the technical specifications of the sensors can be found in Appendix C.

The infrared camera and the depth sensor are rigidly attached to an acrylic frame to preserve their relative position and orientation (see Figure 4.1). With this configuration the sensors share almost the same FOV and registration errors due to the parallax effect area reduced. The acrylic frame can be attached to a camera handle or mounted on a tripod.




Figure 4.1: Thermal camera and depth sensor rigidly attached to acrylic frame

The thermal camera and the depth sensor’s near-infrared camera have been characterized via the pinhole camera model. The parameters of the model as well as the geometric transformation between the cameras coordinate systems were estimated using the OpenCV Camera Calibration Toolbox and a heated cardboard target with an asymmetric grid with 36 laser-cut circular holes (diameter = 50 mm, spacing = 100 mm). These parameters are used to correct the optical distortions introduced by the camera lens and to determine the correspondence between points in the depth and thermal images.

The 3D model is generated from the depth images using version 1.8 of the Point Cloud Library’s open source implementation of the Kinect Fusion algorithm [35, 81]. The algorithm uses a volumetric array to store information needed to reconstruct the surface of objects. Each element in the volumetric array, called a voxel, stores the signed distance from its centroid to the object’s surface along the line that goes through the optical center of the depth sensor. The distance is truncated to a predefined region near the surface. The pose of the volumetric array is defined with respect to the initial pose of the depth sensor and is configurable. Subsequent poses of the depth sensor are then expressed with respect to the volumetric array’s frame of reference. The depth sensor pose is tracked using the *iterative closest point (ICP)* algorithm. The algorithm uses the latest depth image and a less noisy synthetic point cloud generated from the data in the volumetric array to determine the camera displacement and rotation between depth images. The code outputs a point cloud and renders a 2D view of the 3D model from the current depth sensor’s perspective. A more detailed description of the algorithm is given in section 6.1. Both the point cloud and the rendering can be texturized with color

information provided by the depth sensor's built-in RGB camera. The color information is stored in a separate volumetric array.

The original implementation of the open source Kinect Fusion algorithm was complemented with an additional volumetric array that stores information about the object's surface temperature and the measuring conditions under which the temperature measurements are taken. The 3D thermogram is generated by texturizing the 3D model with the temperature information captured with the thermographic camera. Additional textures based on the measuring conditions can also be applied to the 3D model. These textures provide the user with live feedback on whether temperature measurements were taken under good measuring conditions or not. An algorithm has also been designed to avoid incorrect temperature-vertex assignments caused by the parallax effect for surfaces within the **FOV** of the depth sensor. A more detailed description of the temperature-vertex assignment algorithm and the additional confidence textures is given in sections 6.2 and 6.3.

The 3D model is updated with every new depth image. Because the images produced by the sensors are not synchronized, the pose of the thermal camera at the time it delivers a new thermal image is estimated by interpolating the pose of the depth sensor for depth images generated before and after the new thermal image. This is described in more detail in sections 5.3 and 6.2.

The thermographic camera is switched on at least 20 minutes before scanning an object. A **NUC** procedure is carried out before every scan and the camera's automatic **NUC** procedure remains disabled during the scan.

The system can be operated in real time with a voxel count of $480 \times 480 \times 480$ voxels but the spatial resolution can be improved in off-line mode by increasing the voxel count up to $640 \times 640 \times 640$ voxels. At this volume resolution, for voxels with edge length equal to 2 mm, the maximum volume size would be $1.28 \text{ m} \times 1.28 \text{ m} \times 1.28 \text{ m}$. Algorithms to extend the operating volume and to correct drift errors in the generation of 3D models have been proposed in [88–90].

The software is built on top of the open source Robot Operating System (ROS) framework [70]. ROS facilitates running the data acquisition and data processing nodes on separate computers. This allows, for example, to mount the sensors and a lightweight computer on a copter and keep the processing laptop or workstation on ground. ROS visualization tool RViz is used to display in real time the input images (depth, color and thermal), the textured point clouds, the textured renderings, the position of the sensor with respect to the 3D model, the boundaries of the volume containing the 3D model,

the preconfigured scanning range and the path followed by the sensors during the scan as shown in Figure 4.2.

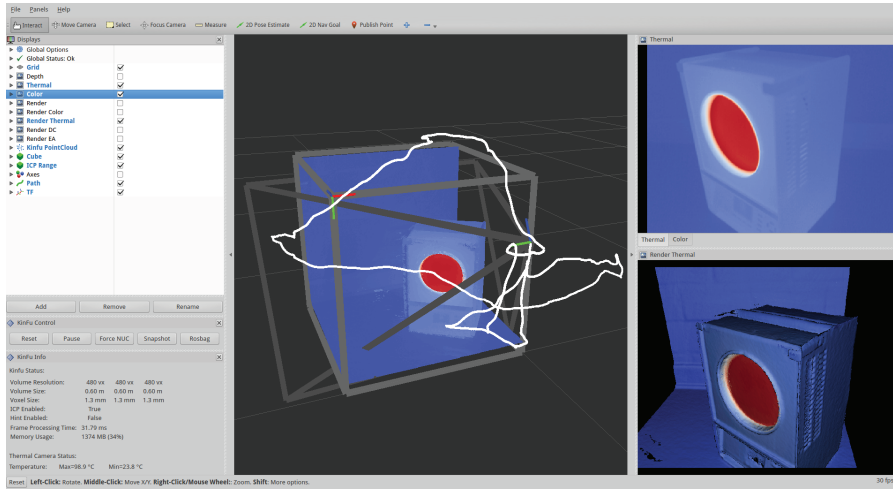


Figure 4.2: RViz screenshot

As it is, this 3D thermal imaging system can be used to scan objects whose surface temperature, position and shape do not vary appreciably during the scan.

4.2 Summary

This chapter introduces the 3D thermal imaging system used in this work. It presents the main characteristics of the hardware components and briefly explains how the system is calibrated, how it generates 3D thermograms in real time, under which conditions the system is typically operated and how the results are visualized.

Chapter 5

3D thermal imaging: calibration

There are many aspects involved in the generation of high-fidelity 3D thermograms. This chapter elaborates on the most important aspects concerning the calibration of the 3D thermal imaging system. The current state of the art and the contributions presented in this work are described in the following sections.

5.1 Radiometric calibration

5.1.1 Overview

The *International Temperature Scale of 1990* ([ITS-90](#)) is an international standard that specifies methods and equipment for making precise and highly reproducible temperature measurements [\[69\]](#). For a temperature measurement to be internationally accepted, it must be traceable back to this scale.

In the context of infrared thermography, radiometric calibration refers to the process by which a traceable [ITS-90](#) temperature is attributed to the output of a thermographic camera in such a way that the measurement uncertainties are understood [\[49\]](#). The description of the calibration procedure and other aspects of radiation thermometry including terminology, usage, testing, accuracy levels and uncertainty analysis are covered by additional standards. A summary of these standards as well as a brief description of their scope and applicability are presented in [\[47\]](#). Different aspects involved in the calibration of thermal cameras are discussed in [\[46\]](#) and [\[12\]](#).

A thermographic camera that has been radiometrically calibrated can be used as a contactless temperature measuring device when the measurement conditions under which

the radiometric calibration was carried out stay the same. Even if the measurement conditions vary slightly, it is possible to adjust some parameters in the user interface of the thermographic camera, like the emissivity of the observed surface or the transmittance of the medium, and still obtain accurate estimates of the object's surface temperature.

Often, the radiometric calibration of the thermographic cameras is carried out at the factory for the typical temperature ranges and measurement conditions under which the camera will be operated. The results of the calibration are included in the device's firmware or are provided with the accompanying software. The Optris PI 450, for example, includes calibration data for the ranges -20°C to 100°C , 0°C to 250°C and 150°C to 900°C . Sometimes, it is desired to optimize the accuracy of the camera for a specific temperature range or for particular measurement conditions. In medical applications, for example, high accuracy is desired in the temperature range between 35°C and 42°C . Ad-hoc calibration methods have been proposed for this purpose [79].

Because the 3D thermal imaging system in this work is not restricted to a particular application, no ad-hoc radiometric calibration was carried out. However, as it was briefly discussed in section 3.3.1, the value of the temperature measurements is affected by the total amount of radiation that enters the thermographic camera. This behavior is not well documented in the literature and can affect the fidelity of 3D thermograms, since during the scan the total amount of radiation that enters the thermographic camera is constantly changing. A series of experiments was carried out in order to understand this behavior and quantify the magnitude of the measurement error. The result of these experiments are presented in this section.

5.1.2 Related work

A review of the literature returned only one work that addresses the fact that temperature measurements made by the thermographic camera of a particular object in a scene are affected by the thermal radiation emitted by other objects in the same scene [91]. In his book, the author defines the term *scene influence factor (SIF)* as the change of measured temperature of a target at the center of the image, which is maintained at a temperature approximately equal to the mean of the camera's operating temperature range, as the temperature of the area surrounding the target is changed from being the minimum to the maximum temperature covered by the range when the target has a diameter equal to $1/4$ of the width of the image. The author proposes the use of SIF as an indicator of possible errors if accurate temperature measurements are required when viewing a scene that may include some high- or low-temperature objects. The calculation of SIF will return some non-zero value, since the phenomenon is real, but this value is

not very informative. It does not indicate, for example, what is the maximum measurement error that can be expected nor under which circumstances this does happen. Does the size of the objects in the image affect the measurements? Moreover, SIF does not indicate how the temperature measured by other pixels is affected. Unfortunately, the author does not accompany his definition with experimental data.

5.1.3 Experiments

In order to understand how the temperature measurement of an individual pixel is affected by the temperature measured by other pixels in the thermal image, three experiments were conducted. Figure 5.1 shows the experimental setup. The thermographic camera was turned on at least 30 min before carrying out the experiments to let the temperature of its housing stabilize. Furthermore, a NUC was carried out before any sequence of images was captured. Images were captured at a frame rate of 40 fps.




Figure 5.1: Experimental setup

In the first experiment, the Optris PI 450 was placed in front of the calibrator at a distance d of approximately 22 cm. The temperature of the calibration disc was varied from 30 °C to 120 °C in 10 °C increments. At each temperature setting, cardboard patterns at room temperature (approx. 24 °C) with circular holes of different sizes were placed in front of the infrared calibrator and forty thermal images were captured. For each pattern and temperature setting, the capture procedure took less than 2 seconds to minimize the warming up of the patterns. The images were averaged to reduce image noise. Figure 5.2 shows the patterns utilized in the experiment and Figure 5.3 shows the corresponding averaged thermal images when the temperature of the calibration disc was 120 °C.

In the second experiment, the temperature of the calibration disc was set to 120 °C and the camera was placed in front of the calibrator at a distance d of 15 cm, 25 cm, 35 cm, 45 cm and 55 cm. At each distance, the cardboard patterns at room temperature were




Figure 5.2: Color images of test patterns with holes of different radii r




Figure 5.3: Thermal images of test patterns right after inserting them in front of the infrared calibrator. $T_{\text{pattern}} \approx 24^\circ\text{C}$, $T_{\text{calibrator}} = 120^\circ\text{C}$.

inserted in front of the infrared calibrator and forty thermal images were captured. For each pattern, the capture procedure took less than 2 seconds to minimize the warming up of the patterns. The images were averaged to reduce image noise. The thermal images corresponding to each pattern at the different distances are shown in Figure 5.4.

In the third experiment, the camera was moved very close to the infrared calibrator such that its entire FOV was covered by the calibration disc. The temperature of the calibration disc was set to 120°C . Again, forty thermal images were captured and averaged to reduce image noise. The camera remained in front of the calibrator only during the capture procedure, which took less than 2 seconds. This was done intentionally so to minimize the warming up of the camera housing due to the heat produced by the calibration disc. The images were averaged to reduce signal noise. Figure 5.5 shows the

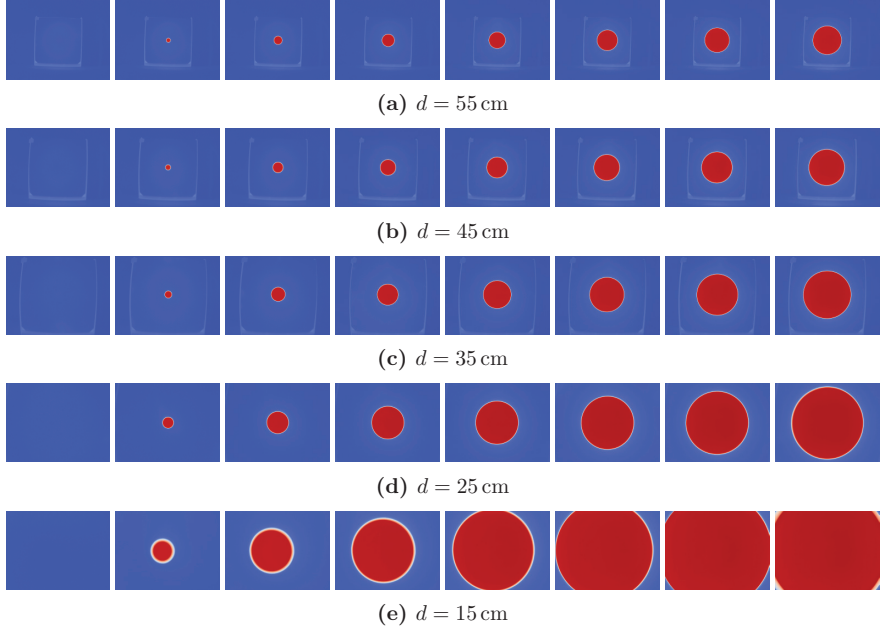


Figure 5.4: Thermal images of test patterns right after inserting them in front of the infrared calibrator taken at different distances d . $T_{\text{pattern}} \approx 24^\circ\text{C}$, $T_{\text{calibrator}} = 120^\circ\text{C}$.

averaged thermal images for different focus settings. In the captions, the variable d denotes the distance from the camera's optical center to the center of the calibration disc and the variables D_F and D_N are defined in section 2.3.2. At this distance, the DOF is very narrow and only the region near the center of the calibrator is in focus. It is not clear what causes the temperature gradients at the corners of the image. Moving the camera slightly to the left, to the right, up or down, while keeping the calibration disc within the FOV of the camera, produces the same pattern. A plausible explanation would be an over- and under compensation of the vignetting effect (see natural vignetting in [11]).




Figure 5.5: Thermal images of the calibration disc when it covers the entire FOV of the thermographic camera. $T_{\text{calibrator}} = 120^\circ\text{C}$

5.1.4 Analysis

Figure 5.6 shows the temperature value measured by the pixel that sees the center of the calibration disc in the first experiment. The hole radii are expressed in pixels as seen by the camera at a distance of 22 cm. The ratio 10 mm to 70 mm corresponding to the smallest and largest hole is not exactly the same when measured in pixels due to distortions introduced by the lens. It can be observed that the temperature value increases as more pixels are irradiated by the calibration disc. It can also be observed that the difference between the temperature measured with the patterns that have the smallest and largest hole increases as the temperature of the calibration disc increases. This trend can be better appreciated in Figure 5.7. The magnitude of the error appears to increase linearly with the temperature of the calibrator ($R^2 = 98.3\%$).




Figure 5.6: Temperature measured at the center of the calibration disc for patterns with different hole sizes and for different temperatures of the calibrator ($d = 22$ cm)




Figure 5.7: Difference between the maximum and minimum temperature measurements for each temperature setting of the calibrator

Figure 5.8 shows the temperature value measured by four pixel that see the center of the calibration disc in the second experiment. It can be observed that the curves produced with data captured at different distances overlap each other. This suggests that the distance to the calibrator plays no role as long as the same amount of thermal radiation reaches the same detectors. The plot converges to a value of approximately 121.2°C as more pixels get irradiated by the calibration disc. This is the value measured at the center of the calibration disc in the third experiment, when the disc center is within the camera's DOF.






Figure 5.8: Temperature measured at the center of the calibrator for patterns with different hole sizes and at different distances d from the calibrator ($T_{\text{calibrator}} = 120^{\circ}\text{C}$)

The graph shows that for a background temperature of 24°C , depending on how much of the surface at 120°C is seen by the camera, the measurement error can be as large as 7°C . This error is much larger than the accuracy of $\pm 2^{\circ}\text{C}$ or 2% of the reading specified for the camera and, based on the trend exposed in Figure 5.7, it is expected to be even larger at higher temperatures. Furthermore, the shape of the curve indicates that the error sensitivity is greater when the calibration disc is depicted by fewer pixels in the thermal image.


Figure 5.9 shows temperature measured by each pixel along a radial path starting at the center of the calibration disc for three different temperature settings. The plot were generated using the averaged thermal images from the first experiment. Each curve corresponds to a test pattern with different hole size.




(a) Path when hole radius is 70 mm



(b) Temperature measured along path
($T_{\text{calibrator}} = 40^\circ\text{C}$)



(c) Temperature measured along path
($T_{\text{calibrator}} = 80^\circ\text{C}$)



(d) Temperature measured along path
($T_{\text{calibrator}} = 120^\circ\text{C}$)

Figure 5.9: Temperature profile from the center of the calibration disc to the upper right corner of the thermal image

On the one hand, even though the actual temperature of both the calibration disc and the cardboard pattern is uniform, at high temperatures and large hole radii, it becomes evident that the temperature measurements corresponding to the calibration disc start decreasing and those corresponding to the pattern start increasing long before the boundary between the calibration disc and the pattern is reached. This means that the influence that surrounding pixels have on an individual pixel decreases with distance. On the other hand, even though the detector element at the center of thermographic camera's FPA and its neighbors are impinged by exactly the same amount of radiation when the hole radius is 10 mm as when it is 70 mm, they indicate different temperatures. This means that the total amount of radiation that enters the camera also influences the temperature measurements and not just the thermal radiation that impinges on the neighboring detectors.

5.1.5 Discussion

The Optris PI 450 uses an uncooled microbolometer FPA to measure thermal radiation. Changes in the electrical signal produced by each element in the array results from changes in its electric resistance which depends on the detector's temperature. The changes are measured with respect to an operating point with resistance R_0 at temperature T_0 . The electrical signal is mapped to temperature values using calibration curves obtained from the radiometric calibration. Because microbolometer arrays work under nearly vacuum conditions [11], changes in the detector's temperature occur only due to radiation and conduction. Ideally, the operating conditions would remain constant and changes in the detector's temperature would only result from changes in the amount of radiation it is subject to. The experiments show, however, that this assumption does not hold. The temperature value measured by a particular detector is also influenced by the radiation detected by other elements in the array. This suggests that some amount of the incident energy is being transferred between the detectors affecting their temperatures. The most reasonable explanation is that this energy is conducted by the common substrate that all detectors are attached to. This would cause temperature gradients in the substrate that depend on the intensity and distribution of the incident radiation. These temperature gradients would explain the slopes near the edge of the calibration disc in Figure 5.9. Similarly, the increase in the temperature value measured at the center of the calibration disc would be explained by a rise in the temperature of the substrate.

The measurement errors can only be clearly appreciated when a large number of pixels at significant different temperatures are involved. The individual contributions of each pixel are too small to be directly measured. Without further information about the inner structure of the thermographic camera, in particular, of the microbolometer array, it is

difficult to precisely characterize the temperature gradients that form in the sensor. For instance, the gradient may be symmetric at the center of the detector array but this might not necessarily be the case at the edges. In theory, information about how energy is transferred between the individual detectors could be used to correct via software the measurement errors it introduces.

The behavior observed in the experiments is not particular to the Optiris PI 450. The second experiment was also conducted with a VarioCAM hr head 600, a high-end thermographic camera, and similar results were obtained (see Appendix D).

In 3D thermal imaging, this behavior not only affects the accuracy of the thermal texture but also the uncertainty of the individual temperature measurements, since the distribution of the temperature of objects in the scene is constantly changing. Moreover, according to Figure 5.8, the measurement error as the camera is moved closer to or farther from object or as the object enters or leaves the FOV is expected to be greater for objects depicted by fewer pixels.

5.2 Geometric calibration

5.2.1 Overview

As will be discussed in more detail in section 6.2, knowing in advance the parameters of the mathematical model that represents the depth and thermographic cameras as well as the six parameters that relate the position and orientation of one of the cameras with respect to the other greatly simplifies the fusion of depth and thermal data.

In the field of computer vision, the procedure used to determine the parameters of the camera model and the pose of the camera with respect to a reference coordinate system is known as geometric calibration or camera calibration. In particular, the estimation of the parameters of the camera model is referred to as intrinsic calibration and the estimation of parameters that describe the camera pose relative to the reference frame is called extrinsic calibration. Both, the model parameters and the pose parameters can be estimated using the same data generated during the calibration procedure. The mathematical model commonly used to describe the camera is the pinhole camera model introduced in section 2.3.5.

The most accurate procedures to carry out a camera calibration require the use of a reference object. This object is referred to as calibration target. It contains multiple features that can be seen by the camera and whose positions with respect to an arbitrary reference coordinate system are well known. During the calibration procedure

the target is oriented in different directions and images of it are captured from different positions. These images are then processed with a camera calibration software. The *mean reprojection error* **MRE** is used as a qualitative measure of the accuracy of the calibration procedure. This value is expressed in pixels and indicates the average distance between the actual projection of the target's features onto the image plane and the projection calculated using the pin hole camera model and the estimated parameters. The calibration software attempts to minimize this value. The OpenCV camera calibration function calculates the **MRE** as the *root mean squared error* (**RMSE**) of the distances (5.1a) while other works including [82, 93] calculate it as the *mean absolute error* (**MAE**) of the distances (5.1b). For the same data set, the two metrics satisfy the inequality 5.1c. In these equations, M denotes the number of images, N represents the number of features per image, p denotes the coordinates of a detected feature and p' denotes the coordinates of the corresponding projected feature. More information on the various calibration procedures and algorithms can be found in [52].

$$\text{MRE}_{\text{RMSE}} = \sqrt{\frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N \|p[m, n] - p'[m, n]\|^2} \quad (5.1a)$$

$$\text{MRE}_{\text{MAE}} = \frac{1}{MN} \sum_{m=1}^M \sum_{n=1}^N \|p[m, n] - p'[m, n]\| \quad (5.1b)$$

$$\text{MRE}_{\text{MAE}} \leq \text{MRE}_{\text{RMSE}} \leq \sqrt{MN} \text{MRE}_{\text{MAE}} \quad (5.1c)$$

Camera calibration procedures were originally developed to calibrate conventional color cameras. The calibration of these cameras is often carried out using a printed chessboard pattern attached to a flat rigid surface. The features to be tracked are given by the corners of the squares. Such targets are easy to fabricate and the degree of accuracy that can be achieved with them is good enough for most practical applications. Unfortunately, this type of target cannot be used to calibrate the thermographic camera because its surface has a uniform temperature and the difference in emissivity of the white and black squares does not provide enough contrast to accurately detect their corners.

The target required to calibrate the 3D thermographic system must provide features that can be clearly appreciated by both the thermographic camera and the near-infrared camera in the depth sensor. This section discusses targets proposed in the literature for the intrinsic calibration of the thermographic camera as well as those proposed for its extrinsic calibration relative to a depth sensor. This section also introduces an improved calibration target design, presents the results obtained with it, proposes new improvements based on an analysis of the results and uncovers an inherent limitation of the calibration procedure.

5.2.2 Related work

The challenge of designing good calibration targets for thermographic cameras lies in producing sharp features in the thermal image at well-defined positions. In any material, changes in temperature due to thermal conductivity tend to be gradual. As a consequence, it is difficult to estimate the position of thermal features with high accuracy. This problem was circumvented in [60] by placing a metal wire grid in front of a plastic block and heating it up with hot air. The grid can be seen in the thermal image due to the difference in emissivity of both materials. The metal wire basically reflects the cooler temperature of its surroundings. Because the wires are thin and the pixel resolution of thermographic cameras in general is low, this calibration target can only be used for calibrating thermographic cameras at short distances. Furthermore, the metal grid must be handled with great care to avoid deforming it.

Other works also exploit the very low emissivity characteristic of common metals like copper and aluminum in the design of calibration targets. In [30], for example, a thermographic camera was calibrated using a *printed circuit board* (PCB) with a chessboard pattern. Two advantages of this approach are the high degree of accuracy with which the pattern can be fabricated and the rigidity of the PCB. As with the previous calibration target, the PCB was heated up with a hair dryer before imaging to improve the contrast of the thermal images. A similar approach was employed in [75], but in this case the squares of one color of a printed chessboard pattern were covered with aluminum foil. This calibration target is easier and cheaper to manufacture but is less accurate than the one made with a PCB. An alternative approach that also uses a chessboard pattern was presented in [1]. In this work, the pattern is printed on a vinyl film which is then attached to an aluminum plate. In addition, the plate is heated by a flexible silicone heater to help improve the contrast of the thermal image. This avoids the problem of having to repeatedly heat up the target after each new capture. The work presented in [48] uses a similar target but achieves the desired contrast by reflecting the radiation from the sky instead of heating up the target.

The calibration target proposed in [14] also makes use of an aluminum plate but in this case the features are created by drilling blind holes into it. Multiple internal reflections inside the holes make them emit much more radiation than the aluminum surface. The plate is also heated to improve the contrast of the thermal image. A shortcoming with this approach is that the amount of radiation that the holes emit varies strongly with the angle of emission. Holes that are not aligned with the optical axis of the thermal camera will appear less sharp than those that are. This adds uncertainty to the estimation of their centroids when changing the orientation of the target during the calibration procedure.

A common disadvantage of targets implemented with low emissivity materials is that they are highly reflective. Care must be taken during the calibration procedure to avoid the reflection of hot spots that might interfere with the detection of the target's features. An alternative target design that does not require the use of low emissivity materials was presented in [42]. The target consists of a matrix of incandescent light bulbs embedded in a large wood plank. Because large calibration targets can be built this way, this design can be useful when calibrating a thermographic camera with a large nearest focus distance D_N (see section 2.3.2). Nevertheless, the same author argues in a subsequent publication that more accurate calibration results can be achieved with targets based on low emissivity materials [43]. Moreover, this design is not suitable for calibrating thermographic cameras at short distances because the irregular thermal gradients within the light bulbs prevent an accurate detection of their centroids.

Finally, calibration targets based on perforated boards have been proposed in [7, 71, 80, 82]. This type of target offers several advantages over the previously discussed designs. The target can be made of a non-reflecting material like cardboard and the holes can be cut with great accuracy using a laser cutter. Moreover, the background can be chosen in such a way that it can be seen by cameras operating in different regions of the electromagnetic spectrum. For instance, a white cardboard with a black hot background could be seen simultaneously by conventional, near-infrared and thermographic cameras. The problem with the implementations that have been proposed so far is that the board has to be thick to remain rigid but the target's thickness affects the accuracy with which the centroid of the holes can be detected when the target is tilted.

5.2.3 Calibration procedure

The calibration target utilized in this work is made of a light-colored three-layer cardboard plane with a 4 by 9 asymmetric grid of laser-cut circular holes, a black painted metal plate and a heating foil (see Figure 5.10). The manufacturing accuracy of the laser cutter is approximately 0.1 mm. The distance between the hole centers is 100 mm and the holes in the front, middle and back cardboard layers have a diameter of 50 mm, 55 mm and 60 mm, respectively. Each cardboard layer is 1 mm thick and the layers are glued together to provide rigidity to the target. Thus, only the front cardboard layer is visible for viewing angles smaller than 68° and the target appears to have a thickness of only 1 mm. The heating foil is attached to the back of the metal plate and both are mounted behind the cardboard plane. The thermal conductivity of the metal plate cares for the even distribution of the heat produced by the foil. This setup provides a background at a uniform temperature of approximately 60°C . The plate is not in direct contact with the cardboard to avoid temperature gradients due to thermal conduction.

Moreover, the plate is painted black to avoid reflections and to increase contrast in the color and near-infrared images. The target is held in a lightweight frame for additional structural stability.





Figure 5.10: Calibration target


During the calibration procedure, the target was mounted on a tripod and placed in front of the cameras at a distance of approximately 80 cm. The focus of the thermographic camera was adjusted accordingly. For the chosen focus setting, the [DOF](#) of the thermographic camera extends approximately from 60 cm to 100 cm. It was assumed that the target also lay within the depth sensor's [DOF](#), since its operating range extends from 35 cm to 140 cm. In addition, the target was illuminated with near-infrared light emitted by the LED rings of OptiTrack S250e cameras to improve image contrast.

The target was manually moved, rotated and tilted within the [FOV](#) and the [DOF](#) of both cameras, for the most part, trying to have as many features as possible near the corners and edges of the common [FOV](#). For each target pose, sequences of 40 thermal images and 40 near-infrared images were taken at frame rates of 40 fps and 30 fps, respectively. A sequence pair was discarded if the uncertainty in the estimation of any feature in either sequence was greater than 0.12 pixels, an achievable precision that is smaller than the expected [MRE](#) value. The images corresponding to each pose and modality were averaged to reduce noise. A total of 30 targets poses were captured (see [Figure 5.11](#)).

The images were processed with the OpenCV camera calibration library. In particular, the function `findCirclesGridDefault` was used to estimate the coordinates of the hole centroids in both sets of images and the function `calibrateCamera` was used to estimate the intrinsic parameters. The `calibrateCamera` function is based on the calibration method proposed in [\[95\]](#) and the MATLAB Calibration Toolbox [\[10\]](#). The camera calibration routine was run 10000 times. Each time, a different combination of 20 images was used. (There are over 30 million different ways of combining 20 images selected from a pool of 30 images.) Due to the bell-shaped distribution of the estimated intrinsic parameters estimation, the average values were accepted as the most likely result of the intrinsic calibration and were used in the extrinsic calibration (see [5.2.4](#)).



(a) Near-infrared



(b) Thermal infrared

Figure 5.11: Calibration images

The pose of the thermographic and infrared cameras with respect to the calibration target was estimated using the function *solvePnP* which solves the Perspective-n-Point problem. The pose of the thermographic camera with respect to the near-infrared camera was then computed using a simple geometric transformation. One set of translation and rotation vectors was obtained for each target pose. The average translation and rotation

vectors were accepted as the result of the extrinsic calibration. The average rotation vector was calculated using the algorithm proposed in [50].

5.2.4 Experimental results and analysis

The results of the geometric calibration are summarized in Table 5.1. Here, the terms $MRE_{NIR \rightarrow Thermal}$ and $MRE_{Thermal \rightarrow NIR}$ denote mean reprojection errors calculated using the geometric transformation between the cameras. $MRE_{NIR \rightarrow Thermal}$ is calculated by first expressing the target features coordinates with respect to the near-infrared camera coordinate system, then projecting those points onto the thermal image and finally finding the average distance between these projections and the features detected on the thermal image. Conversely, $MRE_{Thermal \rightarrow NIR}$ is calculated by first expressing the target features coordinates with respect to the thermographic camera coordinate system, then projecting those points onto the near-infrared image and finally finding the average distance between these projections and the features detected in the near-infrared image. The precision of the results is given for one standard deviation from the mean.

	MRE	f_u (px)	f_v (px)	c_u (px)	c_v (px)
Primsense Carmine 1.09	0.204 ± 0.008	567.7 ± 0.7	567.7 ± 0.7	321.3 ± 1.7	242.4 ± 0.7
Optris PI 450	0.153 ± 0.005	371.4 ± 0.6	372.1 ± 0.6	205.5 ± 1.3	150.5 ± 0.5

	$k'_1 (\times 10^{-3})$	$k'_2 (\times 10^{-3})$	$k'_3 (\times 10^{-3})$	$p'_1 (\times 10^{-6})$	$p'_2 (\times 10^{-6})$
Primsense Carmine 1.09	-45 ± 5	80 ± 35	-109 ± 75	-651 ± 373	2360 ± 925
Optris PI 450	-439 ± 3	290 ± 16	-114 ± 23	454 ± 289	-257 ± 286

(a) Intrinsic calibration of each camera

$MRE_{NIR \rightarrow Thermal}$	$MRE_{Thermal \rightarrow NIR}$	t_x (mm)	t_y (mm)	t_z (mm)	r_x ($^\circ$)	r_y ($^\circ$)	r_z ($^\circ$)
0.141 ± 0.032	0.201 ± 0.063	48.5 ± 1.1	53.3 ± 1.5	11.5 ± 0.3	-0.5	0.4	-0.1

(b) Extrinsic calibration between cameras

Table 5.1: Calibration results

OpenCV uses normalized image coordinates for the estimation of the distortion coefficients [62]. As a consequence, the distortion coefficients k'_1 , k'_2 , k'_3 , p'_1 and p'_2 used by OpenCV differ in value from those that would be estimated using the definitions given in section 2.3.5. The normalized image coordinates x' and y' are calculated using (5.2a) and the pixel coordinates are calculated using (5.2b).

$$\begin{bmatrix} x' & y' & 1 \end{bmatrix}^T = \frac{1}{Z_C} \cdot \begin{bmatrix} X_C & Y_C & Z_C \end{bmatrix}^T \quad (5.2a)$$

$$\begin{bmatrix} u & v & 1 \end{bmatrix}^T = \mathbf{K}_{3 \times 3} \begin{bmatrix} x'' & y'' & 1 \end{bmatrix}^T \quad (5.2b)$$

Here, x'' and y'' result from applying the Brown-Conrady distortion model to x' and y' . In the particular case when $f_u = f_v = f$ and no higher order coefficients are used, the following equivalences hold: $k_1 = k'_1/f^2$, $k_2 = k'_2/f^4$, $k_3 = k'_3/f^6$, $p_1 = p'_2/f^2$ and $p_2 = p'_1/f^2$.

Figure 5.12 shows histograms of the intrinsic parameters based on the 10000 calibration runs that were carried out with different combination of images. The plots confirm that the results of the calibration procedure tend to concentrate around specific values and illustrate the distribution of the estimated parameters.

The small values of the mean reprojection errors and the low uncertainty (i.e. small standard deviation) with which some of the intrinsic parameters have been estimated can lead to believe that these have also been estimated with great accuracy. This is not necessarily correct. Table 5.2 shows by how much the values of intrinsic parameters would have to deviate from their estimated value in order to cause a increment of 0.1 pixels in the MRE. For instance, even though the parameter c_v corresponding to the thermographic camera has been estimated with high precision, its true value could be off by 16 pixels and the MRE would only change by 0.1 pixels.

	Δf_u	Δf_v	Δc_u	Δc_v	$\Delta k'_1$	$\Delta k'_2$	$\Delta k'_3$	$\Delta p'_1$	$\Delta p'_2$
Primsense Carmine 1.09	0.7 %	0.6 %	3.2 %	4 %	62 %	95 %	210 %	795 %	220 %
Optiris PI 450	1.0 %	1.0 %	8.2 %	11 %	7.3 %	29 %	186 %	1450 %	2450 %

Table 5.2: Percentage change in the individual parameters from their estimated value that causes an increase of 0.1 pixels in the MRE

The MRE depends not only on the intrinsic parameters but also on the extrinsic parameters. Thus, some errors in the estimation of the intrinsic parameters can be compensated by corresponding errors in the estimation of the extrinsic parameters. In particular, large errors in the estimation of the parameters c_u and c_v can be compensated by relatively small errors in the estimation of the orientation parameters. This explains why the value of c_v in the thermographic camera model can deviate by as much as 11 % (i.e. 16.5 pixels) from the estimated value without having much effect on the MRE. Such error would be compensated by a rotation of the camera of 2.2° around its x-axis according to the camera orientation depicted in Figure 2.9. In the case of the depth sensor, a deviation of 4 % (i.e. 9.7 pixels) from the estimated value could be compensated by a rotation of the camera of approximately 0.8° around its x-axis. A mathematical analysis of this behavior is presented in Appendix E.

The low sensitivity of the MRE with respect to changes in the distortion parameters has other causes. On the one hand, small coefficient values cause small distortions values that do not contribute much to the computation of the MRE, hence, they are difficult to estimate with high precision. This explains why the MRE is so insensitive to changes

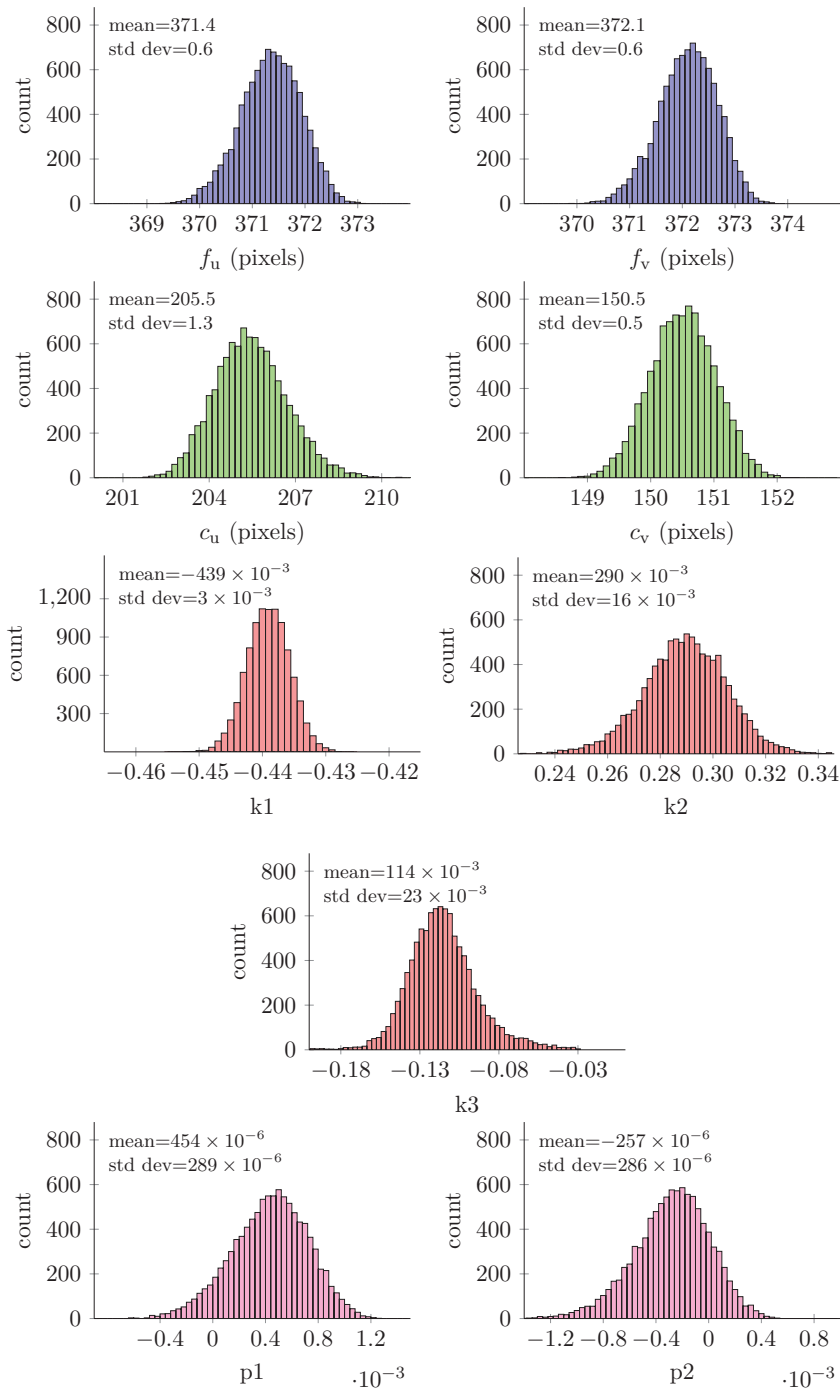


Figure 5.12: Histograms of intrinsic parameters of thermographic camera

in p'_1 and p'_2 in both cameras. It also explains why the uncertainty in the estimation of the radial distortion parameters is greater in the case of the near-infrared camera than in the case of the thermographic camera. On the other hand, the largest values that the normalized image coordinates x' and y' can take within the FOV of a camera is $\tan(\text{HFOV}/2)$ and $\tan(\text{VFOV}/2)$, respectively. This means, that for the thermographic and depth cameras the normalized radius $r' = \sqrt{x'^2 + y'^2}$ is smaller than one. This in turn implies that for the same coefficient value and the same distance to the image center, the higher the order of the radial distortion coefficient, the smaller its contribution to the MRE. This also explains why higher order radial distortion coefficients are estimated with greater uncertainty. Figure 5.13 shows the magnitude of the distortion caused by each individual distortion coefficient. The values were calculated along the straight path that goes from the image center at $[c_u, c_v]$ to the farthest corner of the cameras' FOV. The camera coordinates were first normalized using (5.2a) and the normalized errors were scaled back to pixels by multiplying them by f_u and f_v . The tangential distortion does not exhibit radial symmetry but it is zero at the image center and maximum at the farthest image corner.

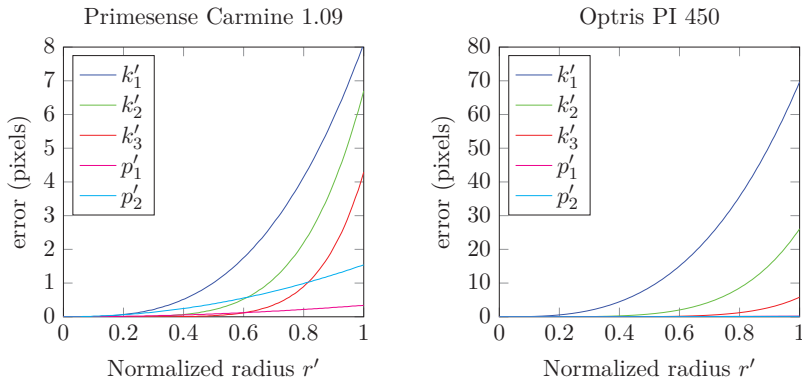


Figure 5.13: Error contribution caused by each individual distortion coefficient

Furthermore, the uncertainty in the estimation of the coefficients depends on the position of the features in the image. The uncertainty becomes larger when most of the features are located near the image center where the distortion is smaller. Table 5.3 shows two alternative sets of intrinsic parameters. These were estimated using subsets of the features detected in each image. The results in the second row were obtained using the subset consisting of the 18 features that are farthest from the image center. The results in the third row were obtained using the subset consisting of the 18 features that are closest to the image center. The original results using all 36 features are presented again in the first row for quick reference. A comparison of the parameters in the second and third rows, which were estimated using the same number of features, clearly shows

greater uncertainty in the estimation of all distortion coefficients, in particular k'_3 . At the same time, it can be observed that despite the large uncertainty of the estimated coefficients the **MRE** gets smaller when the features are all near the image center. This is explained by the fact that errors due to distortion are small near the image center.

	MRE	f_u (px)	f_v (px)	c_u (px)	c_v (px)
All features	0.204 ± 0.008	567.7 ± 0.7	567.7 ± 0.7	321.3 ± 1.7	242.4 ± 0.7
Features near image borders	0.177 ± 0.008	567.5 ± 0.9	567.5 ± 0.9	320.3 ± 1.7	242.9 ± 0.7
Features near image center	0.137 ± 0.007	567.1 ± 1.1	567.6 ± 1.1	324.5 ± 2.7	247.1 ± 1.6

	$k'_1 (\times 10^{-3})$	$k'_2 (\times 10^{-3})$	$k'_3 (\times 10^{-3})$	$p'_1 (\times 10^{-6})$	$p'_2 (\times 10^{-6})$
All features	-45 ± 5	80 ± 35	-109 ± 75	-651 ± 375	2360 ± 925
Features near image borders	-35 ± 10	50 ± 59	-87 ± 106	-636 ± 384	1900 ± 929
Features near image center	-88 ± 16	689 ± 367	-3570 ± 2480	2100 ± 848	3670 ± 1560

(a) Primesense Carmine 1.09

	MRE	f_u (px)	f_v (px)	c_u (px)	c_v (px)
All features	0.153 ± 0.005	371.4 ± 0.6	372.1 ± 0.6	205.5 ± 1.3	150.5 ± 0.5
Features near image borders	0.146 ± 0.005	372.1 ± 0.7	373.0 ± 0.6	205.1 ± 1.4	151.9 ± 0.9
Features near image center	0.104 ± 0.004	369.4 ± 0.9	370.2 ± 0.9	204.3 ± 1.8	150.3 ± 1.0

	$k'_1 (\times 10^{-3})$	$k'_2 (\times 10^{-3})$	$k'_3 (\times 10^{-3})$	$p'_1 (\times 10^{-6})$	$p'_2 (\times 10^{-6})$
All features	-439 ± 3	290 ± 16	-114 ± 23	454 ± 289	-257 ± 286
Features near image borders	-443 ± 5	306 ± 22	-130 ± 31	661 ± 414	-233 ± 387
Features near image center	-443 ± 17	307 ± 365	-30 ± 2390	-677 ± 686	612 ± 512

(b) Optris PI 450

Table 5.3: Intrinsic calibration using the same images but different subsets of features

Finally, the accuracy of the intrinsic calibration also affects the accuracy of the extrinsic calibration between the thermographic camera and the near-infrared camera. The errors in the estimation of the cameras poses with respect to the calibration target (which compensate the errors in the estimation of the c_u and c_v parameters) are absorbed by the geometric transformation between both sensors without having much impact on the computation of $\text{MRE}_{\text{NIR} \rightarrow \text{Thermal}}$ and $\text{MRE}_{\text{Thermal} \rightarrow \text{NIR}}$. For instance, if the extrinsic parameters were recalculated using c_v equal to 134 pixels instead of 150.5 pixels for the thermal camera, the new relative orientation would be $[r_x = 2.0^\circ, r_y = 0.4^\circ, r_z = -0.1^\circ]$ and the new values of $\text{MRE}_{\text{NIR} \rightarrow \text{Thermal}}$ and $\text{MRE}_{\text{Thermal} \rightarrow \text{NIR}}$ would be 0.24 and 0.34 pixels, respectively. The mean reprojection errors would increase by less than 0.15 pixels but the value of r_x would change by 2.5° . This is a significant change in orientation considering that deviations of more than 0.2° from the correct value already produce noticeable mismatches between the 3D model and its thermal texture.

5.2.5 Discussion

The results of the camera calibration show that low reprojection errors and low uncertainty in the estimation of the most significant intrinsic parameters can be achieved with the proposed target. For comparison, the reprojection error reported in [82] for a thermographic camera with similar image resolution is 0.324 ± 0.011 pixels. This reprojection error was calculated using the MRE_{MAE} metric and is twice as large as the one obtained here even though, according to (5.1c), $\text{MRE}_{\text{MAE}} \leq \text{MRE}_{\text{RMSE}}$. A strict comparison, however, is not possible because the reprojection error depends not only on the image resolution and the quality of the target but also on other factors like the thermal sensitivity of the camera and the position of the features in the images. Furthermore, the analysis of the results shows that a small MRE and a low uncertainty in the estimation of the intrinsic parameters alone are not enough to guarantee the accuracy of the calibration results. The main problem is the estimation of the parameters c_u and c_v since the error they contribute to the MRE can be largely compensated by accepting a different camera orientation. According to the analysis presented in Appendix E, the range of possible values that c_u and c_v can take without having a relevant effect on the MRE is very large for cameras with low image resolution. Unfortunately, this problem cannot be circumvented by using other types or better calibration targets. The range of values that c_u and c_v can take is determined by the scaling factors f_u and f_v , the image resolution and the accuracy of the feature detector which in part depends on the accuracy of the detectors. For this reason, these and the extrinsic parameters between cameras had to be manually fine-tuned. This was done by generating a 3D model of the calibration target and renewing its entire thermal texture with every new thermal image. The parameters were adjusted via a configuration window until the texture matched the geometry of the 3D model. The matching process was carried out from various perspectives. Figure 5.14 shows renderings of the 3D thermograms before and after fine-tuning the intrinsic and extrinsic parameters. Section 6.4 describes in more detail how these thermal renderings are generated.

The analysis of the experimental results also indicates that the uncertainty in the estimation of the distortion parameters can be reduced by making sure that many features are as close as possible to the image borders. This suggests that a calibration target based on corners could potentially help achieve slightly more precise results since the edges of a shape can get closer to the image borders than its centroids. Another aspect to consider is that the accuracy with which an algorithm can detect the centroid of a shape decreases as the shape gets distorted near the image borders. Unfortunately, reducing the size of the shape to lessen these effects would at the same time decrease the accuracy with which the centroid can be detected. It remains to be tested whether a




Figure 5.14: Accuracy of the registration before and after fine-tuning the intrinsic and extrinsic parameters

corner detection algorithm can detect distorted corners more accurately than a centroid detection algorithm can detect the centroids of distorted shapes.

Finally, it is recommended to run multiple calibration procedures using different combination of images to figure out the parameter values that occur with highest probability and to determine the uncertainty with which they were estimated. Also, it is recommended that the brightest and darkest color in the near-infrared image (the hottest and coldest temperature in the case of the thermographic camera) belong to the target and not some reflecting surface in the background so that the entire quantization range is used to represent only the target.

5.3 Temporal calibration

5.3.1 Overview

To be able to correctly match and combine data from multiple sensors, it is important to accompany the data with timestamps that indicate the exact time when they are captured. If the timestamps are generated by independent clocks, they need to be synchronized with each other and with the processing unit they are connected to. Furthermore, if the data are to be processed by a commercial computer with a multitasking operating system, the random delays with which the operating system processes the data must be taken into account. This section addresses the problem of the temporal calibration of the different components of the 3D thermal imaging system.

5.3.2 Related work

Because most 3D thermographic systems proposed so far do not acquire data while the sensors are in motion, little attention has been paid to the need to synchronize data acquired with multiple independent sensors. An exception to this is the work presented in [85]. Here, the temporal offset between the depth and thermal images timestamps is estimated by tracking the random cyclic motion of an object with both cameras. Because the timestamps are also subject to non-deterministic delays introduced by the host computer, they are corrected in a subsequent step using a clock synchronization algorithm. The problem with this approach is that the offset calibration procedure has to be repeated every time one of the sensors is reconnected to the host computer as the temporal offset between the cameras may change.

5.3.3 Proposed approach

The depth and thermal cameras used in the 3D thermal imaging system assign timestamps to each generated image but the clocks that generate the timestamps are not synchronized and the sensors do not allow for external synchronization, so it is necessary to synchronize the clocks via software. For this purpose, a simple approach based on a median filter is used. This technique is used in computer networks to synchronize clocks with timing jitter [55] and is already implemented in the depth camera driver [37] but not in the driver of the thermographic camera [51].

Let $t'_s[i]$ and $t_s[i]$ denote the exact time when data frame i is generated by sensor s , the former given by the sensor's internal clock and the latter by the clock of the computer the sensor is connected to. In addition, let $t''_s[i]$ be the time when the computer starts processing data frame i according its internal clock. Then, $t'_s[i]$ and $t''_s[i]$ are related to $t_s[i]$ according to (5.3a) and (5.3b).

$$t'_s[i] = t_s[i] + t_{\text{offset},s} \quad (5.3a)$$

$$t''_s[i] = t_s[i] + \underbrace{t_{\text{delay_avg},s} + t_{\text{delay_rnd},s}[i]}_{\text{total delay}} \quad (5.3b)$$

Here, $t_{\text{offset},s}$ indicates the time difference between the internal clock of the computer and that of the sensor, $t_{\text{delay_avg},s}$ is the average delay between the generation of a data frame and the moment when the computer starts processing it and $t_{\text{delay_rnd},s}[i]$ represents timing jitter. It is assumed that offset drift due to relative clock skew is negligible during the scan period. It is also assumed that the timing jitter of the sensor's clock is negligible compared $t_{\text{delay_rnd},s}[i]$ because timing in the sensor is managed by a

dedicated microcontroller whereas timing in a computer is managed by a multitasking operating system.

In order to reduce random timing errors due to jitter, in particular impulse noise, a median filter of length N is used. The filter is applied to $\Delta t_s[i]$, the difference between $t''_s[i]$ and $t'_s[i]$. The filtered time $t'''_s[i]$ corresponding to frame i of sensors s can then be calculated by adding $t'_s[i]$ and the filter's output $\Delta \tilde{t}_s[i]$. These steps are described by (5.4a), (5.4b) and (5.4c).

$$\Delta t_s[i] = t''_s[i] - t'_s[i] = t_{s,\text{delay_avg}} - t_{s,\text{offset}} + t_{\text{delay_rnd},s}[i] \quad (5.4a)$$

$$\Delta \tilde{t}_s[i] = \text{median}_N(\Delta t_s[i]) = t_{s,\text{delay_avg}} - t_{s,\text{offset}} + \tilde{t}_{\text{delay_rnd},s}[i] \quad (5.4b)$$

$$t'''_s[i] = t'_s[i] + \Delta \tilde{t}_s[i] = t_s[i] + t_{\text{delay_avg},s} + \tilde{t}_{\text{delay_rnd},s}[i] \quad (5.4c)$$

The computer clock is used as a reference for the purpose of synchronizing the sensor clocks; there is no real need to determine the exact value of $t_{\text{delay_avg},s}$ for each sensor. As long as the average delay during the scan is not too large, it is enough to make sure that both sensors experience the same average delay. This can be accomplished by adding the term $\Delta t_{\text{delay_avg}}$ to the timestamp of one of the sensors such that $t_{\text{delay_avg},s1} = t_{\text{delay_avg},s2} + \Delta t_{\text{delay_avg}}$. Thus, the timestamps $\hat{t}_{\text{depth}}[i]$ and $\hat{t}_{\text{thermal}}[i]$ assigned to depth and thermal images can be calculated using (5.5a) and (5.5b).

$$\begin{aligned} \Delta t_{\text{delay_avg}} &= t_{\text{delay_avg},\text{depth}} - t_{\text{delay_avg},\text{thermal}} \\ \hat{t}_{\text{depth}}[i] &= t'''_{\text{depth}}[i] = t'_{\text{depth}}[i] + \Delta \tilde{t}_{\text{depth}}[i] \end{aligned} \quad (5.5a)$$

$$\hat{t}_{\text{thermal}}[i] = t'''_{\text{thermal}}[i] + \Delta t_{\text{delay_avg}} = t'_{\text{thermal}}[i] + \Delta \tilde{t}_{\text{thermal}}[i] + \Delta t_{\text{delay_avg}} \quad (5.5b)$$

Then, the problem of synchronizing the sensors is reduced to minimizing $\tilde{t}_{\text{delay_rnd},s}[i]$ for each sensor and estimating the correct value of $\Delta t_{\text{delay_avg}}$.

5.3.4 Experimental results

Figure 5.15 shows the difference between $t''_s[i]$ and the time values that result from applying a linear regression to the t''_s data set for both sensors. The timestamps were acquired under normal operation conditions of the 3D thermal imaging system (CPU load $\approx 30\%$).

Figure 5.16, on the other hand, shows the difference between $t'''_s[i]$ and the time values that result from applying a linear regression to the t'''_s data set. Again, the data were acquired under normal operation conditions of the 3D thermal imaging system. A median filter of length $N = 15$ that is part of the depth sensor driver was enabled. The same




Figure 5.15: Variability of image timestamps before filtering

filter with $N = 21$ was added to the driver of the thermal camera. These values of N are equivalent to a filter length of approx. 0.5 s for frame rates of 30 fps and 40 fps, respectively.

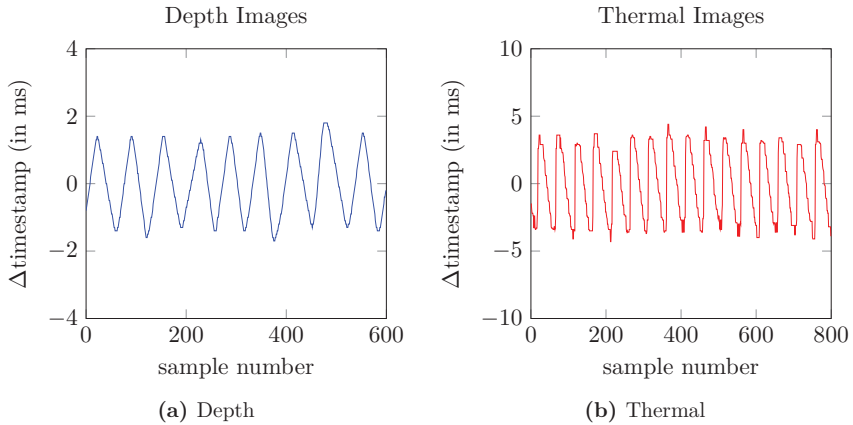


Figure 5.16: Variability of image timestamps after filtering

The parameter $\Delta t_{\text{delay_avg}}$ in (5.5b) can be configured manually. For this purpose, a 3D thermogram of the infrared calibrator was generated and its thermal texture completely renewed with every new thermal image. The system was rotated back and forth roughly around the optical center of the thermographic camera as shown in 5.17. The parameter $\Delta t_{\text{delay_avg}}$ was adjusted via a configuration window until the 3D model and the thermal texture remained aligned during the rotation motion.

Misalignment effects that occur during the motion of the system when the value of $\Delta t_{\text{delay_avg}}$ is not appropriately set are depicted in Figure 5.18. The snapshots were



Figure 5.17: Motion of 3D thermal imaging system to calibrate $\Delta t_{\text{delay_avg}}$ (top view)

taken while scanning the infrared calibrator by rotating the system anti-clockwise (i.e. the calibrator in the rendering moves to the right).

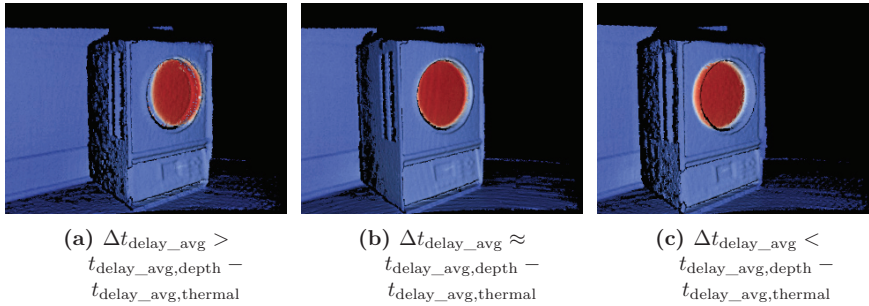


Figure 5.18: Effects of inaccurate calibration of $\Delta t_{\text{delay_avg}}$

5.3.5 Discussion

The variability of the filtered image timestamps in Figure 5.16 is not zero but is reasonably small and the impulse noise has been eliminated. Under these conditions and at scan velocities of 0.3 m/s and 18°/s, a 6 ms error in the timestamp of a thermal image causes a spatial offset of approx. 1.8 mm due to translation and 1.9 mm due to rotation when the object's surface is at a distance of 1 m. These values are smaller than the **IFOV** of the thermal camera at the same distance which is about 3 mm, so errors due to the asynchronicity of data are hard to notice in the resulting 3D thermogram. A mismatch will be observed, however, if the scan direction is abruptly changed. If higher scanning speeds or greater scanning distances are required, not only does the amount of filtering need to be increased, but the performance of the depth and thermal sensors would need to be improved as well.

Because of the assumption that the sensors' internal clocks are very reliable, that is, the timing jitter and the offset drift during operation are negligible, it would suffice to

find out the offset between the clocks to be able to synchronize the data generated by sensors. The advantage of this solution is that the synchronization is not affected by the timing jitter introduced by the processing unit when it is used as reference clock. The disadvantage of this solution, however, is that the offset would need to be recalculated every time the sensors are disconnected from the computer or the computer is turned off. The proposed approach works with the difference of the average processing delays which remains constant as long as the same computer is used under approximately the same CPU load.

An additional temporal calibration may be required if the acquisition and processing of data are not carried out by the same computer, for example in aerial inspections, where the sensors are attached to a lightweight computer on a copter and the data processing is done on ground. In such scenarios, the clocks of the two computers need to be synchronized. For this purpose, the Linux application *chrony*¹ can be used.

5.4 Summary

This chapter elaborates on the three key aspects concerning the calibration of the 3D thermal imaging system.

Section 5.1, radiometric calibration, presents a series of original experiments based on test patterns and an infrared calibrator that reveals how the temperature measurement associated with a particular detector in the camera's FPA is affected by the radiation that impinges on other detectors in the array. Only one mention of the existence of this phenomenon was found in the literature but no data to characterize it and no explanation as to its cause were provided. The experiments showed that the measurement deviations increased as the temperature difference between the surface of interest and the background increased and also as the ratio of pixels irradiated by the surface of interest to those irradiated by the background decreased. An analysis of the results suggests that the measurement deviations are caused by an energy exchange between the individual elements in the detector array and the substrate they lie on. In principle, knowledge about the geometric structure of the thermographic camera's detector array and how energy is transferred between the individual detectors and the substrate could be used to design an algorithm to compensate for these measurement deviations. In applications where accurate temperature measurements are necessary or where the range of surface temperatures to be inspected is large, this phenomenon may significantly affect the fidelity and appearance of 3D thermograms because the total amount of radiation that

¹ <https://chrony.tuxfamily.org/>

enters the camera continuously changes as the scanning system is moved around the object.

Section 5.2, geometric calibration, analyzes the results of the intrinsic and extrinsic calibration of the thermographic camera and of the near-infrared camera in the depth sensor. The calibration of the cameras was carried out with an improved calibration target proposed in this work. Despite the low MRE and the very low uncertainty in the estimation of the dominant parameters of both cameras achieved with the new calibration target, it was necessary to manually fine-tune the value of some parameters in order to properly register depth and thermal data. The analysis revealed that errors introduced by the inaccurate estimation of the camera orientation relative to the calibration target can be compensated by the inaccurate estimation of the c_u and c_v parameters. This is true regardless of the type, quality, size and pose of the calibration target. The accuracy of the estimation depends mainly on the image resolution and the accuracy of the feature detector. As a consequence, the calculation of the orientation of one camera relative to the other using the current approach is also inaccurate and the manual fine-tuning of some of the intrinsic and extrinsic parameters might be required.

Section 5.3, temporal calibration, provides a detailed analysis of the problem of synchronizing depth and thermal data when the sensors do not allow for external synchronization. A simple solution based on a median filter and a one-time manual offset adjustment is proposed and evaluated. This approach effectively eliminates impulse noise that would otherwise cause noticeable data fusion errors. The residual noise goes unnoticed under the proposed scanning speeds. This approach also does not require to estimate the offset between the two sensors every time the computer is restarted or the sensors are connected to the computer.

Chapter 6

3D thermal imaging: generation

Once the sensors have been calibrated, the 3D thermal imaging system is ready to acquire and process depth and thermal data. This chapter elaborates on the most important aspects concerning the generation of 3D thermograms in real time. The current state of the art and the contributions presented in this work are described in the following sections.

6.1 3D model generation

6.1.1 Overview

The reconstruction of accurate 3D models is an essential step in the generation of high-fidelity 3D thermograms. The choice of sensors is important because they determine the application range and performance limits that can be expected from the system. One of the sensors is invariably a thermal camera. The other sensor must enable the reconstruction of a 3D model of an object and the localization of the thermal camera with respect to the 3D model. In this section, the reconstruction technique utilized in this work and two minor modifications made to it will be described in greater detail. Also, the behavior of the reconstruction system under different configurations and operating conditions is evaluated. The main modifications to texturize the 3D model with thermal data are discussed later in sections 6.2 and 6.3.

6.1.2 Related work

Besides the thermographic camera, which is mainly used for noninvasively measuring the surface temperature of an object, most 3D thermographic systems proposed in the

literature include one or more additional sensors or actuators to capture the geometry of the scanned object. The choice of sensor determines to a large extent the method with which the 3D model is reconstructed. It also determines the freedom of movement and the degree of maneuverability that the system will have. Greater freedom of movement and greater maneuverability allow better measurement conditions during the scan. Other aspects to be considered are cost and overall performance. The following paragraphs summarize the sensors and methods proposed so far for the reconstruction of the 3D model and briefly explain their advantages and disadvantages.

The work of [60] proposes the reconstruction of 3D models by capturing the silhouette of an object with a single thermographic camera as the object is rotated on a turntable. This approach has the advantage that it does not require additional sensors but fails to recognize and reconstruct concavities in the object. Also, because it requires a turntable, its practical use is limited to a lab environment, that is, it cannot be used to scan large objects such as buildings or industrial machines.

The reconstruction of 3D models proposed in [36], [27] and [68] is carried out using stereo vision systems mounted on a tripod, the former with a pair of color cameras and the latter two with a pair of thermal cameras. The first approach has the potential to produce good 3D models if the object is well illuminated, the object's surface has a rich visual texture and high resolution cameras are used. However, thermographic inspections are sometimes carried out under low light conditions and often surfaces do not have a rich visual texture. Furthermore, high resolution cameras with high frame rates are expensive and demand significant computational power to process the high resolution images. The second approach is not affected by illumination conditions but requires a second thermographic camera, which is also expensive. Moreover, the quality of the generated 3D models (or disparity maps) is poor because thermal images have low resolution and because thermal gradients vary rather slowly so that most surfaces do not exhibit a rich thermal texture. In both cases, the cameras are mounted on a tripod and the 3D model is captured from a single perspective but, in principle, if lightweight cameras are used, stereo vision could offer complete freedom of movement and good maneuverability.

3D models in [87] and [28] are generated using computer vision techniques known as *Structure from Motion* (SfM) and *Multi-View Stereo* (MVS). This approach requires only a single high resolution color camera and offers great freedom of movement during the scan. But, as is the case with stereo vision, this approach only works well when the objects are well illuminated and their surfaces are feature rich. Even then, the point clouds that it produces tend to be sparse. Another disadvantage of this particular approach is that the 3D models it produces have no scale information.

In [4, 13, 23, 41] and [9], *terrestrial laser scanners* (TLS) are used to acquire geometric information of monuments and buildings. The advantage of using TLS is that they can generate point clouds of large volumes that are much denser and more accurate than those created with the strategies explained before. However, TLS must remain motionless while scanning their surroundings and because they are heavy and bulky, they need to be mounted on a tripod or robot. These properties restrict their maneuverability and freedom of movement during the scan. TLS are also very expensive.

Another strategy that has been used to reconstruct models for 3D thermal imaging is structured light. Here, geometric information is derived from the deformations that a known pattern suffers when it is projected onto the surface of an object. This technique overcomes many of the limitations of stereo vision: the projected pattern provides the features that are needed to gather depth information, the scale is known and it can be used under no light conditions, which is important when carrying out inspections at night or in low lit environments. In [26, 65, 76, 92] a pattern is projected onto the object's surface using a bulky external projector which restricts the maneuverability of the system, but in [80] and [84] the projector is part of the depth sensor making the system lightweight, compact and therefore easy to maneuver. A metrological evaluation of the sensors used in these works is presented in [24]. The system proposed in [84] goes one step further and integrates the Kinect Fusion algorithm [35] in order to generate 3D thermograms in real time as the sensors are moved around the object. An open source version of the Kinect Fusion algorithm has been made available by The Point Cloud Library [81] and the algorithm is being improved independently by other research groups [61, 66, 88].

A more detailed description of these and other techniques for generating 3D models can be found in [94].

6.1.3 Original approach and modifications

In the 3D thermal imaging system used in this work, the reconstruction of the 3D model is carried out using the open source implementation of the Kinect Fusion algorithm [81]. In this approach, the surface of an object is modeled by means of a volumetric array, where each element of the array, called a voxel, is assigned a value based on the distance between its centroid and the object surface. This distance is measured along the line that goes through the voxel's centroid and the depth camera's optical center at the moment a new depth image is captured. The value that is assigned to the voxel results from truncating and normalizing the measured distance according to a predefined threshold. It is positive when the voxel is in front of the surface and negative when it is

behind. This value is averaged with the previous value assigned to the voxel to reduce the measurement uncertainty. Only voxels in front of the object or whose distance to the object's surface is smaller than the threshold are updated. Because of the way surfaces are described, the volumetric array is known as **TSDF** volume, where **TSDF** stands for *truncated signed distance function*. This representation is depicted for simplicity in two dimensions in Figure 6.1.

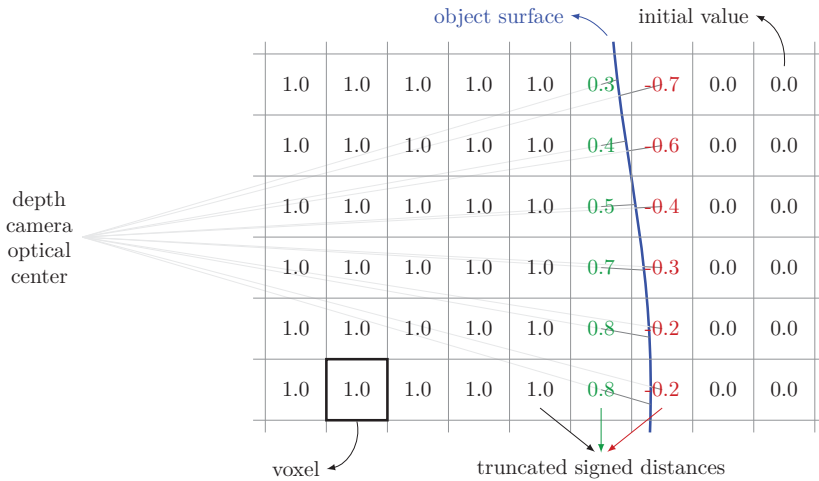


Figure 6.1: TSDF based surface description

The maximum voxel count is constrained by the GPU's memory size and, in the case of real-time operation, by the GPU's processing power. This parameter and the voxel size determine the largest scan volume that can be achieved with the current implementation. Algorithms to extend the scan volume and to correct drift errors in the generation of 3D models have been proposed in [88–90].

The volumetric array can take the shape of a cube or a rectangular cuboid. The vertex and edges of one of its corners are used as the origin and axes of the 3D model coordinate system. At the beginning of a scan, preconfigured values are used to define the position and orientation of the volume with respect to the depth camera. Thereafter, new poses of the sensor are expressed with respect to the volume's frame of reference. The code has been modified to show the boundaries of the scan volume and the FOV of the depth camera (see Figure 6.2). These modifications allow the user to keep the inspected object within the operating range of the depth camera during the scanning process.

The pose of the depth camera is estimated using the ICP algorithm and point clouds derived from the current 3D model and the latest depth image. ICP is an iterative procedure that applies a series of translational and rotational transformations to one of the




Figure 6.2: Volume, FOV and coordinate systems at the beginning of a scan

point clouds until it aligns with the other within a configurable tolerance. This procedure was independently developed by [8, 15, 16]. The product of these transformations describes the relative position and orientation of the original point clouds and can be used to determine the new pose of the depth camera with respect to the 3D model. This allows tracking the trajectory of the depth sensor during the 3D model generation. The new pose of the depth camera is used to integrate the geometric information of the latest depth image with the current 3D model. The point cloud derived from the 3D model is obtained via a procedure called ray casting which was first introduced in [5]. A change in sign of the **TSDF** value as the rays traverse the volume (starting from the current pose of the depth camera) indicates the presence of a surface. Its exact position can be obtained via interpolation. This process is repeated every time a new depth image arrives. Algorithms 1, 2 and 3 describe briefly how the open source Kinect Fusion algorithm generates and updates the 3D model in real time with each new incoming depth image. In these algorithms and in subsequent sections, the following additional definitions are used:

- **depth image:** a two-dimensional array that stores the depth data of points on the object's surface seen by the depth camera. Element $[u, v]$ in the array stores the Z_C coordinate in Figure 2.9 of the point seen by pixel $[u, v]$ of the depth camera.
- **vertex:** a point in the 3D model that represents a point on the object's surface.
- **vertex map:** a two-dimensional array that stores the coordinates of the points on the 3D model seen by the depth camera from its current pose. Element $[u, v]$ in the array stores the X_C , Y_C and Z_C coordinates in Figure 2.9 of the point seen by pixel $[u, v]$ of the depth camera. The vertex map provides the same information as

the depth map but in a more convenient format for certain types of operations like geometric transformations. In essence, both the depth and vertex maps describe a point cloud.

- **normal map:** a two-dimensional array that stores the components of the normal vectors at the points on the 3D model indicated by the vertex map. Basically, element $[u, v]$ in the array stores the components n_x , n_y and n_z of the normal vector corresponding to vertex $[u, v]$ in the vertex map.

Algorithm 1 3D model generation

```

1: procedure OnNewDepthImage
2:    $\mathbf{I}_{\text{depth,filtered}} \leftarrow \text{filter}(\mathbf{I}_{\text{depth}})$  ▷ filter depth image
3:    $\mathbf{M}_v, \mathbf{M}_n \leftarrow \text{generateMaps}(\mathbf{I}_{\text{depth,filtered}})$  ▷ generate vertex and normal maps
4:   if isFirst( $\mathbf{I}_{\text{depth}}$ ) then
5:      $\mathbf{R}_{\text{depth}}, \mathbf{t}_{\text{depth}} \leftarrow \mathbf{R}_{\text{init}}, \mathbf{t}_{\text{init}}$  ▷ initialize camera pose
6:      $\mathbf{V}_{\text{tsdf}} \leftarrow \text{UpdateVolume}(\mathbf{V}_{\text{tsdf}}, \mathbf{I}_{\text{depth}}, \mathbf{R}_{\text{depth}}, \mathbf{t}_{\text{depth}})$  ▷ update TSDF volume
7:      $\mathbf{M}'_v, \mathbf{M}'_n \leftarrow \mathbf{M}_v, \mathbf{M}_n$  ▷ save vertex and normal maps
8:   else
9:      $\mathbf{R}_{\text{depth}}, \mathbf{t}_{\text{depth}} \leftarrow \text{ICP}(\mathbf{M}_v, \mathbf{M}_n, \mathbf{M}'_v, \mathbf{M}'_n, \mathbf{R}'_{\text{depth}}, \mathbf{t}'_{\text{depth}})$  ▷ calculate camera pose
10:    if  $\|\mathbf{t}_{\text{depth}} - \mathbf{t}'_{\text{depth}}\|_2 < t_{\text{thld}}$  and  $\angle(\mathbf{R}_{\text{depth}}, \mathbf{R}'_{\text{depth}}) < \theta_{\text{thld}}$  then
11:       $\mathbf{V}_{\text{tsdf}} \leftarrow \text{UpdateVolume}(\mathbf{V}_{\text{tsdf}}, \mathbf{I}_{\text{depth}}, \mathbf{R}_{\text{depth}}, \mathbf{t}_{\text{depth}})$  ▷ update TSDF volume
12:       $\mathbf{M}'_v, \mathbf{M}'_n \leftarrow \text{RayCast}(\mathbf{V}_{\text{tsdf}}, \mathbf{R}_{\text{depth}}, \mathbf{t}_{\text{depth}})$  ▷ generate vertex and normal maps
13:       $\mathbf{R}'_{\text{depth}}, \mathbf{t}'_{\text{depth}} \leftarrow \mathbf{R}_{\text{depth}}, \mathbf{t}_{\text{depth}}$  ▷ save current camera pose

```

Notes:

- ▷ \mathbf{V}_{tsdf} : TSDF volume. Each voxel contains the truncated signed distance from its centroid to the object's surface.
 - ▷ $\mathbf{I}_{\text{depth}}$: Depth image. Each pixel contains the distance to the point on the object observed by the depth camera.
 - ▷ \mathbf{M}_v : Vertex map. Each element contains the coordinates of a point on the object seen by the depth camera.
 - ▷ \mathbf{M}_n : Normal map. Each element contains a unit vector orthogonal to the object's surface at the point given by the corresponding element in \mathbf{M}_v .
 - ▷ \mathbf{R}_{ini} : Initial orientation of the depth camera. Manually configured.
 - ▷ $\mathbf{R}_{\text{depth}}$: Current orientation of the depth camera.
 - ▷ $\mathbf{R}'_{\text{depth}}$: Previous orientation of the depth camera.
 - ▷ \mathbf{t}_{ini} : Initial position of the depth camera. Manually configured.
 - ▷ $\mathbf{t}_{\text{depth}}$: Current position of the depth camera.
 - ▷ $\mathbf{t}'_{\text{depth}}$: Previous position of the depth camera.
-

The existing implementation has been modified such that the **TSDF** volume gets updated only when the translation and rotation of the depth camera between consecutive depth frames are within the ranges 0.1 cm to 1 cm and 0.06° to 0.6°, respectively. These values were obtained empirically. The modification helps reduce the accumulation of errors in the 3D model, in particular, when the sensor is abruptly moved or when the pose change estimate is abnormally large due to noisy data. This, in turn, helps reduce the error in the subsequent estimation of the depth camera pose. The modification also sets a limit to the maximum linear and angular velocity with which the system can be moved to approx. 0.3 m/s and 17°/s. The accuracy with which the **ICP** algorithm estimates the camera pose change depends on the accuracy of the depth sensor, the geometry and density of the generated point clouds and the magnitude of the translation and rotation between the point clouds. More accurate depth measurements, complex

Algorithm 2 TSDF volume update

```

1: function UpdateVolume(  $\mathbf{V}_{\text{tsdf}}$ ,  $\mathbf{I}_{\text{depth}}$ ,  $\mathbf{R}_{\text{depth}}$ ,  $\mathbf{t}_{\text{depth}}$  )
2:   for each voxel in  $\mathbf{V}_{\text{tsdf}}$  with indices  $i, j$  and  $k$  do
3:      $\mathbf{c} \leftarrow \text{getCentroid}(i, j, k)$  ▷ calculate position of voxel's centroid
4:      $d_c \leftarrow \|\mathbf{c} - \mathbf{t}_{\text{depth}}\|_2$  ▷ distance from centroid to depth camera
5:      $u, v \leftarrow \text{projectToDepth}(\mathbf{c})$  ▷ project centroid onto depth image
6:      $d_o \leftarrow \mathbf{I}_{\text{depth}}[u, v]$  ▷ distance to object's surface seen by pixel  $[u, v]$ 
7:      $d \leftarrow d_o - d_c$  ▷ signed distance between depth camera and object's surface
8:      $d \leftarrow \max(-1, \min(1, d/d_{\text{max}}))$  ▷ normalize and truncate  $d$  according to fixed  $d_{\text{max}}$ 
9:      $d', w \leftarrow \mathbf{V}_{\text{tsdf}}[i, j, k]$  ▷ get voxel's TSDF value and weight
10:     $\mathbf{V}_{\text{tsdf}}[i, j, k] \leftarrow \text{weightedAverage}(d, d', w)$  ▷ update voxel's TSDF value and weight

```

Notes:

- ▷ \mathbf{V}_{tsdf} : TSDF volume. Each voxel contains the truncated signed distance from its centroid to the object's surface.
 - ▷ $\mathbf{I}_{\text{depth}}$: Depth image. Each pixel contains the distance to the point on the object observed by the pixel.
 - ▷ $\mathbf{R}_{\text{depth}}$: Current orientation of the depth camera.
 - ▷ $\mathbf{t}_{\text{depth}}$: Current position of the depth camera.
-

Algorithm 3 Ray cast

```

1: function RayCast(  $\mathbf{V}_{\text{tsdf}}$ ,  $\mathbf{R}_{\text{depth}}$ ,  $\mathbf{t}_{\text{depth}}$  )
2:    $\mathbf{r}_{\text{ini}} \leftarrow \mathbf{t}_{\text{depth}}$  ▷ ray origin
3:   for each depth detector element with indices  $u$  and  $v$  do
4:      $\mathbf{r} \leftarrow \text{getPosition}(u, v)$  ▷ get pixel position
5:      $\mathbf{r}_{\text{dir}} \leftarrow (\mathbf{r} - \mathbf{t}_{\text{depth}}) / \|\mathbf{r} - \mathbf{t}_{\text{depth}}\|_2$  ▷ normalized ray direction
6:      $\mathbf{r}_{\text{in}}, \mathbf{r}_{\text{out}} \leftarrow \text{getIntersections}(\mathbf{r}_{\text{dir}})$  ▷ intersections of ray and volume faces
7:      $t_{\text{entry}} \leftarrow \|\mathbf{r}_{\text{entry}} - \mathbf{r}_{\text{ini}}\|_2$  ▷ ray length when entering volume
8:      $t_{\text{exit}} \leftarrow \|\mathbf{r}_{\text{exit}} - \mathbf{r}_{\text{ini}}\|_2$  ▷ ray length when exiting volume
9:      $t_{\text{curr}} \leftarrow t_{\text{entry}}$  ▷ current ray length
10:     $\mathbf{r}_{\text{curr}} \leftarrow \mathbf{r}_{\text{ini}} + t_{\text{curr}} \mathbf{r}_{\text{dir}}$  ▷ current ray
11:     $i, j, k \leftarrow \text{getVoxelIndices}(\mathbf{r}_{\text{curr}})$  ▷ voxel pointed at by current ray
12:     $d_{\text{curr}} \leftarrow \mathbf{V}_{\text{tsdf}}[i, j, k]$  ▷ get TSDF value for current voxel
13:    while  $t_{\text{curr}} \leq t_{\text{exit}}$  do
14:       $\mathbf{r}_{\text{next}} \leftarrow \mathbf{r}_{\text{ini}} + (t_{\text{curr}} + t_{\text{step}}) \mathbf{r}_{\text{dir}}$  ▷ increase ray length by fixed amount  $t_{\text{step}}$ 
15:       $i, j, k \leftarrow \text{getVoxelIndices}(\mathbf{r}_{\text{next}})$  ▷ voxel pointed at by next ray
16:       $d_{\text{next}} \leftarrow \mathbf{V}_{\text{tsdf}}[i, j, k]$  ▷ get TSDF value for current voxel
17:      if  $d_{\text{curr}} > 0$  and  $d_{\text{next}} < 0$  then ▷ zero crossing
18:         $t_{\text{int}} \leftarrow \text{interpolate}(\mathbf{r}_{\text{next}}, \mathbf{r}_{\text{curr}}, t_{\text{step}})$ 
19:         $\mathbf{v} \leftarrow \mathbf{r}_{\text{curr}} + t_{\text{int}} \mathbf{r}_{\text{dir}}$  ▷ calculate vertex position
20:         $\mathbf{n} \leftarrow \text{calculateNormal}(\mathbf{v})$  ▷ calculate normal vector at position  $\mathbf{v}$ 
21:         $\mathbf{M}_v[u, v] \leftarrow \mathbf{v}$  ▷ add vertex to vertex map
22:         $\mathbf{M}_n[u, v] \leftarrow \mathbf{n}$  ▷ add normal to normal map
23:        break ▷ stop searching for vertex along this ray
24:       $t_{\text{curr}} \leftarrow t_{\text{curr}} + t_{\text{step}}$ 
25:       $\mathbf{r}_{\text{curr}} \leftarrow \mathbf{r}_{\text{next}}$ 
26:       $d_{\text{curr}} \leftarrow d_{\text{next}}$ 

```

Notes:

- ▷ \mathbf{V}_{tsdf} : TSDF volume. Each voxel contains the truncated signed distance from its centroid to the object's surface.
 - ▷ $\mathbf{R}_{\text{depth}}$: Current orientation of the depth camera.
 - ▷ $\mathbf{t}_{\text{depth}}$: Current position of the depth camera.
 - ▷ \mathbf{M}_v : Vertex map. Each element contains the coordinates of a point on the object seen by the depth camera.
 - ▷ \mathbf{M}_n : Normal map. Each element contains a unit vector orthogonal to the object's surface at the point given by the corresponding element in \mathbf{M}_v .
-

geometries, denser point clouds and smaller translations and rotations result in better accuracy. The actual accuracy of the algorithm is difficult to measure, but its effects can be appreciated in the resulting 3D model.

6.1.4 Experimental results

Two experiments were conducted to understand the behavior of the open source implementation of the Kinect Fusion algorithm under different configurations and operating conditions.

In the first experiment, the impact of the voxel size on the quality of the 3D model was investigated. To keep depth measurement errors as low as possible, the infrared calibrator was scanned from a distance of approximately 50 cm. 3D models were generated using the same data and three different voxel sizes. The voxel size was changed by varying the size of the scan volume while keeping the voxel count constant at $800 \times 800 \times 800$ voxels¹. The size of the scan volume was modified offline after the raw data had been acquired. The resulting 3D models for voxels with edge lengths equal to 4.5 mm, 1.5 mm and 0.5 mm are shown in Figure 6.3.




Figure 6.3: Effect of voxel size on quality of 3D model. Voxel edge length is: a) 4.5 mm, b) 1.5 mm and c) 0.5 mm

In the second experiment, the effect of increased depth measurement errors and the effectiveness of the code modification in reducing the accumulation of errors in the 3D model were investigated. For this purpose, the infrared calibrator was scanned from a distance of approximately 100 cm. Using the same data, two 3D models were generated: one with the original code and one with the modified code. As a reference, a third 3D model was generated from a distance of approximately 50 cm. In all cases, the voxel

¹ Voxel counts larger than $640 \times 640 \times 640$ voxels are possible when the thermal volume is disabled.

edge length was 1.6 mm and the voxel count was $480 \times 480 \times 480$ voxels. The results are depicted in Figure 6.4.




Figure 6.4: Effect of increased depth measurement errors on quality of 3D model: a) $d = 100$ cm with unconditional data integration, b) $d = 100$ cm with conditional data integration and c) $d = 50$ cm with conditional data integration

At all times, a part of the object with a geometry that can be reliably tracked by the ICP algorithm was within the FOV of the depth camera.

6.1.5 Discussion

The first experiment demonstrates the effect that the voxel size has on the quality of the 3D model and suggests that there is a particular voxel size for which best results are achieved. Larger voxel sizes produce excessively smoothed 3D models and small details are lost. Smaller voxel sizes, on the contrary, result in noisy and distorted 3D models. Most of the uncertainty in the estimation of the position of a point on the object's surface is no longer contained within the corresponding voxels in the 3D model and does not get averaged during the TSDF volume update. As a consequence, the ICP algorithm operates on noisier point clouds and the pose of the depth sensor relative to the 3D model is estimated with higher uncertainty. The effects get compounded and manifest themselves as noise and distortions in the 3D model.

The second experiment demonstrates that the conditional integration of depth data based on the magnitude of the depth camera's pose change between consecutive depth frames significantly improves the performance of the Kinect Fusion algorithm when the camera pose is estimated with higher uncertainty. As in the first experiment, higher camera pose uncertainty introduces errors in the 3D model during the volume update

which then causes more errors in the estimation of subsequent camera poses. The proposed modification reduces the compounding of error. The downside of restricting the pose change is that the scan velocity is restricted as well. For the selected thresholds, the maximum scan velocities are 0.3 m/s and 18 °/s.

An comparative analysis of the quality of 3D models generated with the Kinect Fusion algorithm and with a high-end laser scanner can be found in [53].

6.2 Data fusion

6.2.1 Overview

This step in the creation of 3D thermograms concerns itself with determining which temperature value measured by each individual pixel of the thermographic camera corresponds to which vertex of the 3D model. There are two general approaches to fuse thermal and geometric data.

One conceivable way to determine the correspondences between temperature measurements and vertices on the 3D model is to identify common features present in both modalities. The main advantage of this approach is that it does not require prior knowledge of the pose of the thermal camera with respect to the object. The problem, however, is that often there is no direct relation between features in the 3D model and features in the thermal image. For example, a flat surface, like a heated floor, does not exhibit any salient geometric features yet the heating pipelines or cables underneath the floor can be clearly appreciated in a thermal image. Conversely, an edge that can be easily identified in the 3D model may be hard to discern in the associated thermal image if the background or the surfaces that form it are at the same temperature. The directional emissivity of the material might provide some contrast but it might not be sufficient to reliably identify the edge (see rim of the calibration disc in Figure 3.11). Furthermore, due to the low contrast and low resolution characteristics of thermal images, the number of features that could be reliably detected is very small, especially if large groups of pixels in the thermal image are at approximately the same temperature. Last but not least, feature detection is computationally expensive and therefore not very practical when tens of thousands of features need to be detected in real time. For these reasons, this approach is not the most effective way to map thermal information onto the 3D model in real time.

A simpler and by far more effective approach to figure out the correspondences between vertices in the 3D model and pixels in the thermal image is to use the estimated pose

of the depth sensor with respect to the 3D model and the known geometric relation between the depth sensor and the thermographic camera to calculate the pose of the latter with respect to the 3D model. Once the pose of the thermographic camera with respect to the 3D model is known, vertices of the 3D model can be projected onto the thermal image using the pinhole camera model to determine their corresponding temperature values (see section 2.3.5). This approach, however, also introduces some challenges. On the one hand, thermal data are not necessarily generated at the same time as range data, since the sensors may operate at different frame rates and their clocks may not be synchronized. When the system is moving, the pose of the thermal camera at the moment a new thermal image is captured may not be the same as the pose of the depth sensor when the last depth image was generated. If this is not taken into account, temperature values are incorrectly assigned to the 3D model with varying offsets producing a misaligned thermal texture. On the other hand, the depth sensor and the thermographic camera do not share the same FOV and are therefore subject to the parallax effect. Not all points on the object's surface seen by the depth sensor are also seen by the thermographic camera. If this is not taken into account, incorrect temperature values are assigned to vertices in the 3D model corresponding to these points.

6.2.2 Related work

A review of the literature shows that the general approach to map thermal data onto a 3D model is to estimate the pose of the thermographic camera with respect to the 3D model and then project the vertices or polygons that form the model onto the thermal image to obtain the corresponding temperature values. The approaches encountered in the different works differ mainly in the way the pose of the thermographic camera is estimated. The following paragraphs briefly describe and evaluate some of the ideas that have been proposed in the literature.

The simplest approach is to manually select control points in the thermal image, associate them with corresponding points in the 3D model and then use existing algorithms to solve the Perspective- n -Point problem, i.e., the problem of estimating the pose of a calibrated camera given a set of n points in world coordinates and their corresponding 2D projections in the image. This approach was used in [44] to map thermal images on 3D models of building facades. This approach is very effective for quickly texturizing large flat surfaces like building facades but it can be very time consuming when many thermal images are involved. Furthermore, the accuracy of the projections completely depends on the availability of well-defined features common to both data sets. Also, with this approach, the 3D thermogram cannot be generated in real time.

In [31] and [34], the mapping process is automated by using the *scale-invariant feature transform* (SIFT) algorithm to automatically detect features in a sequence of thermal images of a building facade. The images are taken with a thermographic camera mounted on a mobile platform. The features are used to reconstruct the path followed by the camera which is also recorded separately by a GPS receiver. Both paths are registered to determine the approximate pose trajectory followed by the thermographic camera in global coordinates. Finally, the point cloud formed by the detected features is registered with the GPS referenced 3D model to refine the camera pose estimates. This approach to generate 3D thermograms exhibits three major shortcomings that limit its application range and performance. First, since the pose estimation relies on the use of GPS, it can not be used to generate 3D thermograms indoors or of small objects at close range. An *indoor positioning system* (IPS) could be used instead of the GPS but it would require the installation of such system. Second, because both registration steps must be done after the data have been captured, the approach is not suitable for real-time scanning. Third, the approach relies twice on features extracted automatically from thermal images which are, in general, unreliable. The reader is referred to the work of [28], where the unreliability of SIFT and other feature detection algorithms applied thermal images is discussed.

Other approaches to automatize the fusion of thermal and geometric data are presented in [23], [41] and [2]. These approaches are based on the detection and matching of features common to the thermal images and the 3D model. In [23], the features are interest points in both thermal and range images detected with the Harris and Stephens operator. In [41], the features are lines in the thermal image detected with a line segment detector and edges in the 3D model detected with *principal component analysis* (PCA). In [2], the features are detected using the Hough transform of edges in preprocessed 3D height map and infrared images. These approaches do not require additional GPS data but they still depend on the existence of well-defined features that are common to the thermal image and the 3D model, which are neither common nor reliable.

The work of [67] presents an alternative approach to directly map multispectral images onto a 3D model. The method relies on the extraction of a depth map from the model geometry whose pixels maintain an exact correspondence with the vertices of the 3D model. The spectral image to be mapped is translated, rotated and scaled until the mutual information between it and the depth map is maximized. This is equivalent to finding the pose of the spectral camera from which the image was taken but limiting the camera movement to three degrees of freedom. For this reason, the method can only be used to texture flat surfaces in 3D space, that is, it is not suitable for mapping thermal images onto complex geometries. Moreover, though this method does not require the extraction of features from the spectral image, it does rely on the assumption that both

images share some amount of information, which is not always the case for thermal images.

The aforementioned approaches are best suited for scenarios in which thermal and geometric data are not captured at the same time or when the 3D model is generated with CAD software. However, when the thermal and geometric data can be acquired simultaneously, it is easier and more reliable to determine the pose of the thermographic camera with respect to the 3D model indirectly via the sensor used to acquire range information. In this case, the pose of the thermographic camera relative to the range sensing element is determined in advance through some calibration procedure and remains fixed during the data acquisition process. This approach has been used with laser scanners in [19, 59], with depth cameras in [56, 80], with *time of flight* (ToF) cameras in [73] and with color cameras in [25, 28, 92]. In the case of color cameras, range information is obtained via stereo vision, projected structured light or using the SfM algorithm.

Because most 3D thermographic systems proposed so far do not acquire thermal and geometric data while the sensors are in motion, little attention has been paid to the fact that the thermographic camera and the depth sensor might not acquire data synchronously. As a consequence, the rigid transformation between the sensors cannot be used directly to estimate the pose of the thermographic camera with respect to the 3D model. A review of the literature returned only one work that addresses this problem. In the 3D thermographic system presented in [85], a cubic spline is fitted to four consecutive poses of the depth sensor, two of them estimated before and two after the moment a new thermal image is captured. This spline is used to estimate the pose of the depth camera at the time a new thermal image is captured. The pose of the thermographic camera is then calculated using the known geometric transformation relating the two cameras.

Also, very few works have handled the problem of occlusion. A review of the literature related to 3D thermography returned only two works that address this problem. In [19], no temperature value is assigned to a point P in the 3D model, if another point P' is within distance t of the ray that connects P and the optical center O of the thermographic camera. No information is provided regarding the performance of this procedure. In [83], ray casting is used to map temperature values to voxels in the volumetric array and avoid occlusions at the same time. Basically, the volumetric array containing the 3D model is traversed by virtual rays starting at each pixel of the thermal image and going through the camera's optical center until the conditions that describe the presence of a surface are met. When these conditions are found, the temperature value is assigned to the voxel. Thus, other surfaces that lie further back in the path of the ray are ignored. The main shortcoming of this approach is that ray casting is very time consuming.

6.2.3 Proposed approach

If the geometric transformation between the depth sensor and the thermographic camera is known, the problem of figuring out the pose of the thermographic camera with respect to the 3D model can be addressed by interpolating the pose of the depth sensor between the poses estimated from the depth images captured before and after the newly generated thermal image and then applying the rigid geometric transformation between the sensors. This scenario is depicted in Figure 6.5. The gray colored sensors depict the pose of the system when a depth image is captured; the red colored sensors depict the pose of the system when a thermal image is captured. The solid on top represents the depth sensor and the solid at the bottom corresponds to the thermographic camera. Because this approach requires knowing the pose that the depth sensor will have after a new thermal image has been captured, the assignment of thermal information to the 3D model is delayed until a new depth image has arrived. Furthermore, because the images are preprocessed by independent execution threads, to take into account random preprocessing delays, the thermal and depth images are stored in temporary buffers until they are ready to be integrated into the 3D model. Presently, the only preprocessing carried out is image rectification to minimize the effects of radial and tangential distortion.




Figure 6.5: Interpolation of thermal camera pose from depth camera poses

In the current implementation, the position of the thermal camera is calculated by means of linear interpolation and its orientation using spherical linear interpolation (also known as [SLERP](#)). These types of interpolation are simple to implement and work well for constant linear and angular velocities, which may be assumed for short time intervals if the scan trajectory is smooth (i.e., no abrupt changes of direction).

Let $\mathbf{P}_{\text{depth}}(t[i-1])$ and $\mathbf{P}_{\text{depth}}(t[i])$ be the poses of the depth camera with respect to the origin of the 3D model when the depth images $\mathbf{I}_{\text{depth}}[i-1]$ and $\mathbf{I}_{\text{depth}}[i]$ are captured, respectively. In addition, let $\mathbf{P}_{\text{depth}}(t[k])$ be the pose of the depth camera when the thermal image $\mathbf{I}_{\text{thermal}}[k]$ is captured, such that, $t[i-1] \leq t[k] \leq t[i]$. If the pose $\mathbf{P}(t)$ is expressed in terms of the position vector $\mathbf{p}(t)$ and the quaternion $\mathbf{q}(t)$, then, $\mathbf{P}_{\text{depth}}(t[k])$ can be calculated using (6.1b) and (6.1c). The pose of the thermal camera can then be calculated using the fixed geometric transformation between the sensors $\mathbf{T}_{\text{d} \rightarrow \text{t}}$ and

(6.1d).

$$\alpha = (t[k] - t[i - 1]) / (t[i] - t[i - 1]) \quad (6.1a)$$

$$\mathbf{p}_{\text{depth}}(t[k]) = \alpha \mathbf{p}_{\text{depth}}(t[i]) + (1 - \alpha) \mathbf{p}_{\text{depth}}(t[i - 1]) \quad (6.1b)$$

$$\mathbf{q}_{\text{depth}}(t[k]) = \mathbf{q}_{\text{depth}}(t[i - 1]) [\mathbf{q}_{\text{depth}}(t[i - 1])^{-1} \mathbf{q}_{\text{depth}}(t[i])]^\alpha \quad (6.1c)$$

$$\mathbf{P}_{\text{thermal}}(t[k]) = \mathbf{T}_{d \rightarrow t} \mathbf{P}_{\text{depth}}(t[k]) \quad (6.1d)$$

Once the pose of the thermal camera with respect to the origin of the 3D model is known, vertices of the 3D model can be projected onto the thermal image using the equations presented in section 2.3.5 to determine their temperature value. Bilinear interpolation is used when the projection of the vertices falls between pixels.

The parallax problem is depicted in Figure 6.6. It introduces significant fusion errors, in particular, when scanning objects at close range. In the figure, vertices \mathbf{v}_1 , \mathbf{v}_2 and \mathbf{v}_3 are seen by the depth sensor but only \mathbf{v}_1 is seen by the thermal camera. If the parallax effect is not taken into account, the temperature assigned to the vertex \mathbf{v}_1 would also be incorrectly assigned to \mathbf{v}_2 and \mathbf{v}_3 , since the coordinates of both vertices project onto the same detector element in the thermal camera's FPA.




Figure 6.6: Parallax effect

In order to avoid parallax effects during data fusion, the vertex map containing the vertices of the 3D model seen from a given pose should be generated from the perspective of the thermal camera. However, the ray casting procedure used to generate these vertices is computationally expensive and very time consuming. In its current implementation, the proposed system has no time to execute two ray cast procedures between two consecutive depth frames: one from the perspective of the depth camera to generate the 3D model and one from the perspective of the thermographic camera to assign thermal information to the 3D model. Moreover, in the open source implementation of the Kinect Fusion algorithm the values assigned to the voxels in the TSDF volume are calculated based on distances along the path between the voxels and the depth camera. Ray casting

from the perspective of the thermographic camera, that is, following a different path, would cause unwanted artifacts.

In the proposed solution, the vertex map \mathbf{M}_v generated via ray casting from the perspective of the depth camera is filtered before the temperature assignment procedure is carried out. When multiple vertices project onto the same pixel in the thermal image, the filter keeps the vertex closest to the thermal camera and discards the rest. In the example depicted in Figure 6.6, the filter would discard vertices \mathbf{v}_2 and \mathbf{v}_3 and return only vertex \mathbf{v}_1 . Algorithm 4 provides a brief description of this filter.

Algorithm 4 Parallax filtering

```

1: procedure eliminateParallax
2:   for each vertex in  $\mathbf{M}'_v$  with indices  $u$  and  $v$  do
3:      $\mathbf{M}'_v[u, v] \leftarrow \text{NaN}$   $\triangleright$  initialize element of filtered vertex map to invalid value
4:      $\mathbf{M}'_n[u, v] \leftarrow \text{NaN}$   $\triangleright$  initialize element of filtered normal map to invalid value
5:   for each element  $[u, v]$  in  $\mathbf{D}$  do
6:      $\mathbf{D}[u, v] \leftarrow \text{MAX\_FLOAT}$   $\triangleright$  initialize distance of distance map to largest value
7:   for each element  $[u, v]$  in  $\mathbf{M}_v$  do
8:      $d \leftarrow \|\mathbf{M}_v[u, v] - \mathbf{t}_{\text{thermal}}\|_2$   $\triangleright$  distance from vertex to thermal camera
9:      $u', v' \leftarrow \text{projectToThermal}(\mathbf{M}_v[u, v])$   $\triangleright$  project vertex onto thermal image
10:     $\mathbf{D}[u', v'] \leftarrow \min(\mathbf{D}[u', v'], d)$   $\triangleright$  minimum distance from vertex to thermal camera
11:   for each element  $[u, v]$  in  $\mathbf{M}_v$  do
12:      $d \leftarrow \|\mathbf{M}_v[u, v] - \mathbf{t}_{\text{thermal}}\|_2$   $\triangleright$  distance from vertex to thermal camera
13:      $u', v' \leftarrow \text{projectToThermal}(\mathbf{M}_v[u, v])$   $\triangleright$  project vertex onto thermal image
14:     if  $|\mathbf{D}[u', v'] - d| < d_{\text{tol}}$  then
15:        $\mathbf{M}'_v[u, v] \leftarrow \mathbf{M}_v[u, v]$   $\triangleright$  copy vertex from input to output vertex map
16:        $\mathbf{M}'_n[u, v] \leftarrow \mathbf{M}_n[u, v]$   $\triangleright$  copy normal from input to output normal map

```

Notes:

- $\triangleright \mathbf{M}_v[u, v]$: Element $[u, v]$ of the unfiltered vertex map. It is generated by Algorithm 3. It contains the coordinates of the point seen by the depth camera detector element $[u, v]$.
 - $\triangleright \mathbf{M}_n[u, v]$: Element $[u, v]$ of the unfiltered normal map. It is generated by Algorithm 3. It contains the normal vector associated with vertex $\mathbf{M}_v[u, v]$.
 - $\triangleright \mathbf{M}'_v[u, v]$: Element $[u, v]$ of the filtered vertex map. It contains the coordinates of the point seen by both the depth camera detector element $[u, v]$ and the thermal camera.
 - $\triangleright \mathbf{M}'_n[u, v]$: Element $[u, v]$ of the filtered normal map. It contains the normal vector associated with vertex $\mathbf{M}'_v[u, v]$.
 - $\triangleright \mathbf{D}[u', v']$: Element $[u', v']$ of temporary array. It contains the minimum distance between the thermal camera and a vertex in the set of vertices that project onto the thermal camera pixel $[u', v']$.
 - $\triangleright d_{\text{tol}}$: Tolerance
-

Temperature values and information regarding the conditions under which the measurements were taken are stored in a volumetric array called the thermal volume. The array has the same dimensions and voxel count as the **TSDF** volume. The procedure to update the thermal volume is described briefly in Algorithm 5. Basically, the centroid of each voxel in the thermal volume gets projected onto the filtered vertex map. If the distance between the corresponding vertex and the voxel's centroid is smaller than some threshold, the voxel's information is considered for update. This algorithm is based on the algorithm used to assign color to the 3D model in the open source Kinect Fusion implementation but, before the voxel contents are updated, the measurement conditions are validated. This is described in more detail in section 6.3. The information corresponding

to each voxel is stored in a 32 bit word, where the upper 16 bits hold the temperature value and the lower 8 bits the measurement conditions. The remaining bits are reserved for future use.

Algorithm 5 Thermal volume update

```

1: procedure updateThermalVolume
2:   for each voxel in  $\mathbf{V}_{\text{thermal}}$  with indices  $i, j$  and  $k$  do
3:      $\mathbf{c} \leftarrow \text{getCentroid}(i, j, k)$  ▷ calculate position of voxel's centroid
4:      $u, v \leftarrow \text{projectToDepth}(\mathbf{c})$  ▷ project centroid onto depth image
5:      $\mathbf{v} \leftarrow \mathbf{M}'_v[u, v]$  ▷ get coordinates of vertex seen by pixel  $[u, v]$ 
6:      $\mathbf{n} \leftarrow \mathbf{M}'_n[u, v]$  ▷ get normal corresponding to vertex  $\mathbf{v}$ 
7:     if  $\|\mathbf{c} - \mathbf{v}\|_2 < d_{\text{thld}}$  then
8:        $u', v' \leftarrow \text{projectToThermal}(\mathbf{v})$  ▷ project vertex onto thermal image
9:        $d, \theta \leftarrow \text{calcCurrConditions}(\mathbf{v}, \mathbf{n}, \mathbf{t}_{\text{thermal}})$  ▷ deduce current meas. conditions
10:       $d_q, \theta_q \leftarrow \text{discretizeConditions}(d, \theta)$  ▷ discretize current meas. conditions
11:       $d'_q, \theta'_q \leftarrow \text{getMeasConditions}(\mathbf{V}_{\text{thermal}}[i, j, k])$  ▷ get previous meas. conditions
12:       $\text{update} \leftarrow \text{checkConditions}(d_q, \theta_q, d'_q, \theta'_q)$  ▷ check whether to update voxel info
13:      if  $\text{update}$  then
14:         $\mathbf{V}_{\text{thermal}}[i, j, k] \leftarrow \mathbf{I}_{\text{thermal}}[u', v'], d_q, \theta_q$  ▷ update thermal volume

```

Notes:

- ▷ $\mathbf{V}_{\text{thermal}}[i, j, k]$: Voxel $[i, j, k]$ in thermal volume. It contains a temperature value and information regarding the measurement conditions under which the temperature value was acquired.
 - ▷ $\mathbf{I}_{\text{thermal}}[u', v']$: Pixel $[u', v']$ in thermal image. It contains the temperature value measured by the thermal camera detector element $[u', v']$.
 - ▷ $\mathbf{M}'_v[u, v]$: Element $[u, v]$ of the filtered vertex map returned by Algorithm 4.
 - ▷ d_{thld} : Threshold
 - ▷ d : Distance from thermal camera to vertex of temperature measurement to be assigned to this voxel
 - ▷ d' : Distance from thermal camera to vertex of temperature measurement previously assigned to this voxel
 - ▷ θ : Emission angle of temperature measurement to be assigned to this voxel
 - ▷ θ' : Emission angle of temperature measurement previously assigned to this voxel
-

6.2.4 Analytical and experimental results

The depth camera utilized in this work operates at a frame rate of 30 fps. At this frame rate and at a linear scan velocity of 0.3 m/s, the maximum registration offset e_{max} incurred in assuming a linear trajectory (i.e. linear interpolation) between two points A and B instead of a curved trajectory of radius $r = 1$ m centered on the scanned object is less than 0.01 mm (see Figure 6.7). Reducing the radius of curvature to $r = 0.5$ m would increase e_{max} to only 0.02 mm. Increasing the scan velocity to 0.6 m/s would increase e_{max} to only 0.05 mm. In comparison, the maximum registration offset caused by the uncertainty of the timestamps under the same scan conditions can be 3.7 mm. Likewise, a 10% error in the estimation of the thermal camera pose with respect to its previous pose can result in a registration offset of up to 1 mm due to translation and an extra 1 mm due to rotation, assuming linear and rotational velocities of 0.3 m/s and 18°/s, respectively. See Appendix F for more details.




Figure 6.7: Linear interpolation error in a curved trajectory

In addition, an experiment was conducted to test the functionality and performance of the parallax filtering algorithm. The sensors were placed in front of the infrared calibrator at a distance of approx. 50 cm and roughly 25 cm above the ground in order to replicate the conditions illustrated by v_1 and v_3 in Figure 6.6. When no parallax filter is used, vertices representing part of the wall right above the calibrator are incorrectly assigned temperature values corresponding to the calibrator's front panel as indicated by the white arrow in Figure 6.8a. When the parallax filter is used, no temperature information is assigned to these vertices as long as they remain occluded (see Figure 6.8b). Without code optimization, the algorithm takes less than 7 ms to run.




Figure 6.8: Parallax effect in 3D thermogram

6.2.5 Discussion

The accuracy of the data fusion depends mainly on three factors: the accuracy of the intrinsic and extrinsic calibration of the sensors, the accuracy of the timestamps associated with the depth and thermal images and the accuracy with which the pose of the thermographic camera can be estimated with respect to the 3D model. In the proposed approach, the pose of the thermographic camera is obtained via interpolation from the poses of the depth sensor before and after a new thermal image is acquired. A simple analysis of this approach reveals that it is far more important to generate accurate timestamps and to accurately estimate the pose of the depth camera with respect to the 3D model than to use sophisticated interpolation methods. The analysis also shows that,

for a volume size of $0.6\text{ m} \times 0.6\text{ m} \times 0.6\text{ m}$, a voxel count of $480 \times 480 \times 480$ voxels and scan velocities of 0.3 m/s and $18^\circ/\text{s}$, the magnitude of the individual registration offsets is about as large as the volume resolution. Ideally, the combined error should be smaller than the voxel size in order to achieve the most accurate thermal texture for the given operating conditions. This explains random data fusion errors in the thermal texture. The errors are more noticeable when the camera is moved parallel to the object's surface instead of perpendicular to it.

The parallax filter works well only as long as the depth camera is able to see and recognize the points closest to the thermal camera that are involved in the parallax effect. If in Figure 6.6 the depth camera were not able to see v_1 because it is outside its FOV or it were not able to recognize it due to the angle of observation or because it is too close to the edge, its temperature would be incorrectly assigned to the vertices v_2 and v_3 .

6.3 Data validation

6.3.1 Overview

After figuring out which temperature value measured by the thermal camera corresponds to which vertex on the 3D model, the thermal data need to be validated before they get integrated into the 3D model. As the 3D thermographic system is moved around the object, most of the observed surfaces get scanned multiple times from different perspectives. However, as explained in section 3.3, the temperature measurement corresponding to a point on a surface varies with the background temperature distribution and the position of the thermal camera relative to the point. The system needs to know when to replace previous temperature measurements with new values measured from different perspectives. In addition, the user of the system needs to know which surfaces have not been appropriately scanned as well as the reason(s) why they were not appropriately scanned so he or she can rescan the surface under right measurement conditions. This step has been omitted in the majority of 3D thermal imaging systems proposed so far largely because many of these systems generate 3D thermograms from a single perspective. Nevertheless, this step is important for the generation of accurate 3D thermograms.

6.3.2 Related work

A review of the literature related to 3D thermography revealed that few works assess the conditions under which temperature measurement are taken before mapping them onto the 3D model. Moreover, none of the works uses this information to provide feedback to

the user regarding the quality of the data being captured. Although it is not explicitly stated, most likely old values are simply overwritten with new ones. On the other hand, in the few works where the measurement conditions are taken into account, the proposed approaches do not provide the best possible results.

The 3D thermographic system presented in [60], for example, uses weighted averaging to combine temperature values corresponding to the same point on the object's surface but measured from different perspectives. The weights are proportional to the cosine of the emission angle, which can be derived from the geometry of the object and the position of the thermographic camera relative to it. This approach may produce smooth looking 3D thermogram, but it does not provide the most accurate thermal texture. Averaging helps improve the precision of a measurement when the measurement conditions remains the same. During the scan, however, the accuracy of the measurement decreases as the emission angle increases. As a consequence, the average of the measurements will be less accurate than the measurement taken for the smallest emission angle.

In the work of [31], the facade of a large building is captured by means of an ordered sequence of overlapping thermal images which are then mapped onto a 3D model. Because the images are taken from oblique perspectives, the spatial resolution of the thermal texture is not constant. The authors address this problem by overwriting the texture obtained from a thermal image with the higher resolution texture obtained from the subsequent thermal image where these two overlap.

The work of [84] also proposes the use of a weighted average to combine temperature values corresponding to the same point on the object's surface but measured from different perspectives. In this case, the computation of the weights is based on the observation and emission angles, the distance to the object's surface, the scan velocity and the validity of the radiometric model. While this approach takes into consideration more factors that affect the accuracy of the temperature measurements than the one proposed in [60], it actually reduces the accuracy of the thermal texture for the same reason explained before.

6.3.3 Proposed approach

The proposed approach takes into account the measurement conditions that can be inferred by the system to decide whether to replace the temperature measurements made so far with new data from the latest thermal image. The current implementation evaluates two measurement conditions: the distance to object's surface and the emission

angle. These conditions are calculated using (6.2a) and (6.2b), respectively.

$$d = \|\mathbf{v} - \mathbf{t}_{\text{thermal}}\|_2 \quad (6.2a)$$

$$\theta = \cos^{-1} \left(\frac{(\mathbf{v} - \mathbf{t}_{\text{thermal}}) \cdot \mathbf{n}}{\|\mathbf{v} - \mathbf{t}_{\text{thermal}}\|_2 \cdot \|\mathbf{n}\|_2} \right) \quad (6.2b)$$

Here $\mathbf{t}_{\text{thermal}}$ is the position of the thermal camera, \mathbf{v} represents a point on the object surface and \mathbf{n} is the normal vector associated with that point.

The data validation function is embedded in the data fusion procedure and is carried out only for voxels that are suitable for an update according to section 6.2. The function is described briefly by Algorithm 6.

Algorithm 6 Check measurement conditions

```

1: function checkConditions( $d, \theta, d', \theta'$ )
2:    $v' \leftarrow (d_{\min} \leq d' \leq d_{\max})$  and  $(\theta_{\min} \leq \theta' \leq \theta_{\max})$            ▷ validity of previous conditions
3:    $v \leftarrow (d_{\min} \leq d \leq d_{\max})$  and  $(\theta_{\min} \leq \theta \leq \theta_{\max})$        ▷ validity of current conditions
4:   if not  $v$  then  $\text{better} \leftarrow \text{false}$ 
5:   else if  $d < d'$  then  $\text{better} \leftarrow \text{true}$ 
6:   else if  $d > d'$  then  $\text{better} \leftarrow \text{false}$ 
7:   else if  $\theta < \theta'$  then  $\text{better} \leftarrow \text{true}$ 
8:   else if  $\theta > \theta'$  then  $\text{better} \leftarrow \text{false}$ 
9:   else  $\text{better} \leftarrow \text{true}$ 
10:  if not  $v'$  or  $(v' \text{ and } \text{better})$  return true
11:  else return false

```

Notes:

- ▷ $[d_{\min}, d_{\max}]$: Range given by the thermal camera's DOF or manually configured by the user
 - ▷ θ_{\min} : Emission angle below which the temperature measurement might be affected by reflections of the operator. Manually configured.
 - ▷ θ_{\max} : Emission angle above which the temperature measurement error is larger than some value. Manually configured.
-

Basically, the information associated with a voxel gets updated when there is no valid information previously associated with it or when the new measurement conditions are valid and they are the same or better than the previous conditions. In this regard, the distance d in (6.2a) is valid if it is within the thermal camera's DOF or some narrower range defined by the user that guarantees the desired spatial resolution. The condition improves when the new measurement distance between the surface and the camera is shorter. Likewise, the new emission angle θ is valid if it is smaller than some predefined threshold. According to Figure 2.5a, this threshold is somewhere between 40° and 60°, depending on how much error can be tolerated. The condition improves when the new emission angle is smaller. At the same time, to avoid reflections from the operator and the thermographic camera itself whose temperatures are typically above the room temperature, the emission angle should be greater than approximately 15°C.

To minimize memory requirements and because there is no need to estimate the measurement conditions with great resolution, the values are discretized according to (6.3a) and (6.3b). The subscript 'q' denotes the discrete version of the variables. The subscript 'max' indicates a threshold value above which the measurement condition is deemed invalid. The subscript 'min' indicates a threshold value below which the measurement condition is deemed invalid. The subscript 'th' denotes a threshold value below which the observation condition is considered good. Values between the corresponding 'th' and 'max' subscripts are considered acceptable. These threshold values are manually set by the user and allow him or her to distinguish in real time surfaces that should be scanned under stricter measurement conditions (see section 6.4).

$$d_q = \begin{cases} 0 & , \text{ if } d < d_{\min} \text{ or } d > d_{\max} \\ 1 & , \text{ if } d_{\min} \leq d \leq d_{\text{th}} \\ 2 & , \text{ if } d_{\text{th}} < d \leq d_{\max} \end{cases} \quad (6.3a)$$

$$\theta_q = \begin{cases} 0 & , \text{ if } \theta < \theta_{\min} \text{ or } \theta > \theta_{\max} \\ 1 & , \text{ if } \theta_{\min} < \theta \leq \theta_{\text{th}} \\ 2 & , \text{ if } \theta_{\text{th}} < \theta \leq \theta_{\max} \end{cases} \quad (6.3b)$$

Figure 6.9 shows how data is organized within a voxel of the thermal volume. Voxels are 4-byte aligned for code execution efficiency. Some bits are reserved for future use, for example, for recording other measurement conditions.




Figure 6.9: Data structure within a voxel of the thermal volume

6.3.4 Experimental results

Two experiments were conducted to illustrate the importance of validating the thermal data before it is integrated with the 3D model and to demonstrate the effectiveness of the proposed validation algorithm.

The first experiment shows the effect that the measurement distance has on the spatial resolution of the thermal texture during the scan. For this purpose, a resolution test

target was inserted in front of the calibrator as shown in Figure 6.10a. Although the operating range of the depth sensor is 0.35 m to 1.4 m, to keep the fidelity of the 3D model as high as possible, only points on the object's surface at a distance smaller than 0.9 m were used to generate the model, since the accuracy of the depth sensor decreases quadratically with increasing distance to the surface [38]. The voxel size of the TSDF volume was set to 1.3 mm and the temperature of the calibration disc to 80 °C. The generation of the 3D thermogram started by scanning the calibrator from a distance between 0.4 m and 0.6 m. Figure 6.10a and Figure 6.10b show the thermal image and the thermal rendering of the calibrator from a distance of 40 cm. At this distance, the spatial resolution of the thermal image is 0.9 mm, slightly better than the spatial resolution of the 3D model. Then, while facing the calibration disc, the system was moved away from the calibrator up to a distance of 0.9 m. Figure 6.10c and Figure 6.10d show the resulting 3D thermogram when no validation algorithm is used and when Algorithm 6 is used, respectively. It can be observed, that the use of Algorithm 6 preserves the thermal resolution and the accuracy of the temperature measurements obtained at closer distances. In the experiment, the camera's DOF was approximately in the range 35 cm to 45 cm and d_{\min} , d_{th} and d_{\max} were set to 35 cm, 40 cm and 45 cm, respectively.

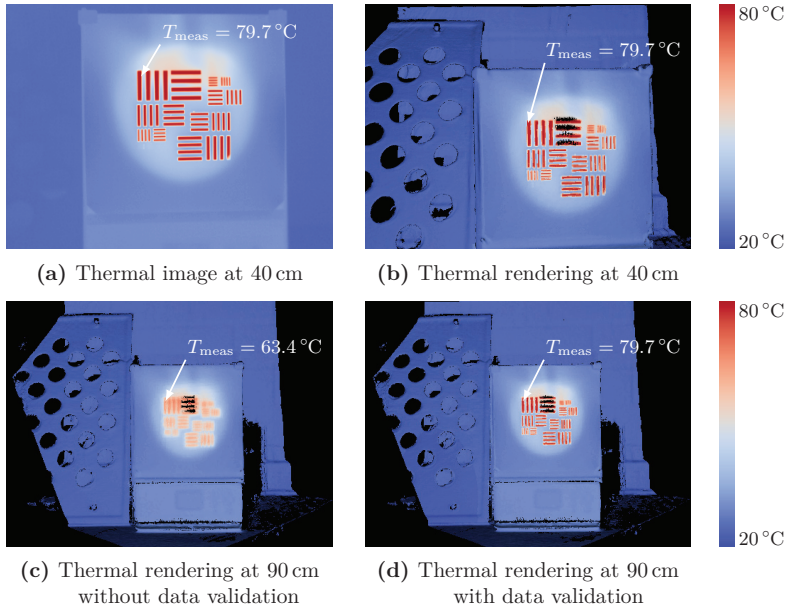


Figure 6.10: Effect of validating the distance to the surface

The second experiment shows the effect that the emission angle has on the accuracy of thermal texture during the scan. The voxel size of the TSDF volume was kept at 1.3 mm and the temperature of the calibration disc was set to 100 °C. In this experiment,

the generation of the 3D thermogram started with the cameras positioned directly in front of the calibration disc. The cameras were moved around the calibrator with a clockwise motion maintaining a distance d of approximately 50 cm until the angle α formed by the thermal camera's optical axis and the normal to the calibration disc was approximately 75° . After this, the algorithm that updates the thermal volume was manually disabled and the cameras were brought back to the position $d \approx 50$ cm and $\alpha \approx 40^\circ$, in order to render the 3D thermogram from a perspective different from that from which it was last updated. Figures 6.11a and 6.11b show the thermal image and thermal rendering at the beginning of the scan, i.e., when $\alpha \approx 0^\circ$. The temperature measurements of the calibrator disc are slightly biased towards lower values because of the lower temperature of the surroundings surfaces (see section 3.3.1). Figure 6.11c shows a thermal rendering of the 3D thermogram at the end of the experiment. The 3D thermogram was generated without validating the thermal data, that is, the temperature measurement assigned to the model vertices were continuously updated regardless of the measurement conditions. It can be observed that, even though the temperature of the calibration disc is uniform, the temperature values assigned to the corresponding vertices in the 3D thermogram vary greatly as a consequence of the directional emissivity of all materials. The effect is emphasized by the fact that the left side of the calibration disc is gradually occluded by its surrounding rim as the thermal camera moves around the calibrator and the temperature value of occluded vertices in the 3D thermogram is not updated. Figure 6.11d shows a thermal rendering of the 3D thermogram generated with the same depth and thermal images as in the previous case, but using Algorithm 6 to validate the temperature measurements with θ_{\min} , θ_{th} and θ_{\max} set to 15° , 30° and 45° , respectively. In this case, the vertices of the calibration disc maintain the temperature values measured under optimal conditions.

6.3.5 Discussion

The experiments show that validating the thermal data before they are fused with the 3D model is essential for generating high-fidelity 3D thermograms. Moreover, because the measurement conditions for each point on the object's surface vary continuously as the system is moved around the object, the validation must be carried out in real time.

The proposed validation algorithm is more effective than an averaging algorithm like the one proposed in [84] because it retains the temperature values taken under the best measurement conditions instead of averaging them with values captured under less favorable conditions. Even if a weighted averaging algorithm is used, as soon as a less accurate temperature value due to the less favorable conditions is included in the calculation of the average, the result becomes suboptimal. The proposed validation

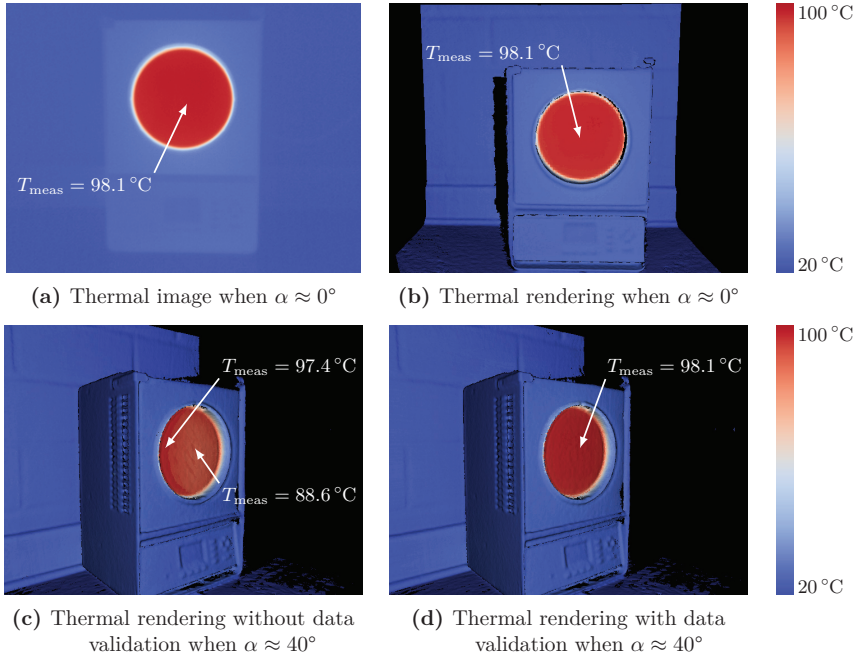


Figure 6.11: Effect of validating the emission angle

algorithm is also more efficient because it is computationally less expensive to execute and does not require figuring out optimal weight values.

At present, the proposed system validates only the distance to the surface and the angle of emission but other factors, like the linear and angular velocity of the sensors, could also be taken into account. For example, to avoid the problem depicted in Figure 3.9b, the thermal texture would only be updated if the linear and angular velocity of the thermal camera are smaller than some threshold. This feature has not been implemented in the current system because the maximum velocity with which an object can be scanned is constrained by the frame rate of the depth sensor and not by the thermographic camera. Because the depth sensor operates at 30 fps and the thermographic camera can operate up to a frame rate of 80 fps, swift movements would cause distortions of the 3D model before they cause artifacts in the thermal texture.

Another factor to be considered is the reproducibility of a temperature measurement of a point on the object's surface taken from different perspectives but under good measurement conditions. This factor would not be used to determine whether to update the thermal texture or not, but rather to provide a measure of uncertainty for each temperature value. This additional information could be used, for example, to identify reflective surfaces. The temperature value of a reflective surface reported by the thermographic

camera is likely to vary with the position of the camera, even under good measurement conditions, when the temperature distribution of the background is nonuniform.

6.4 Data representation and visualization

6.4.1 Overview

A significant advantage of infrared thermography over other temperature measurement techniques is that it allows to visualize the temperature distribution at the surface of an object. The data produced by a thermographic camera can be presented in form of a grayscale or false color image to reveal the temperature distribution of the object's surface. In the former case, the brightness of the pixels is proportional to the estimated temperature value. In the latter, the color of a pixel is obtained from an indexed color palette and the index is calculated from the estimated temperature value. The choice of the color palette affects the ability of the user to detect small temperature gradients within specific temperature ranges (see Figure 3.11). In both cases, the minimum and maximum temperatures associated with the first and last element of the gray or color scale can be manually set or they can be automatically calculated from the temperature values in the thermal image according to some criterion (min/max, 1σ , 3σ , etc.). The choice of these limits also affects the ability of the user to recognize faint details in the thermal image.

Assigning colors to the image pixels based solely on the observed temperature value provides a practical way of visualizing thermal data but it is not the most effective approach. It is often difficult for a person to distinguish and recognize objects in a thermal image when there is no additional information about how the objects look like in the visible spectrum. The human brain is able to visually recognize objects when the visual data conveys information about the orientation of the surfaces relative to each other. The colors of a conventional thermal image convey little information in this sense. The problem is accentuated when all objects within the camera's FOV are at roughly the same temperature (see Figure 3.11). For this reason, the documentation of a thermal inspection carried out with a thermographic camera is often accompanied by color images of the inspected scene. Some thermographic cameras, like the Optris PI 200 [64], even include a built-in color camera for this purpose. 3D thermography has the potential to address this issue, but the majority of 3D thermographic systems proposed so far directly assign the false colors in the thermal image to the vertices of the generated 3D model.

Another shortcoming of current thermographic systems is that it is not possible to distinguish in a thermal image temperature measurements made within the margin of error specified by the thermographic camera from those made with greater margin of error due to the directional nature of the emissivity of materials. In addition, it is sometimes difficult to assess whether the temperature measurements are made with approximately the same spatial resolution or even whether they were made within the camera's DOF. Some thermographic cameras, like the Fluke Ti400 [20], incorporate a laser based autofocus mechanism to guarantee that at least the point of interest is within the camera's DOF. Information about the measuring conditions would be useful to improve the confidence in the temperature measurements, especially if this is available during the inspection, so that the system's operator can rescan parts of the scene that were captured under unfavorable conditions. This issue can also be addressed by 3D thermographic systems, but so far none of the proposed 3D thermographic systems make this information available to the user.

Conventional infrared thermography also has potential for improvement in two aspects: the visualization of the temperature distribution of large objects with complex geometries and the estimation of heat losses or heat gains from the data gathered during the inspection. Currently, one alternative to visualize thermal information of large objects is to capture a thermal image from a distance long enough to ensure that the whole object fits within the thermographic camera's FOV. The problem here is that the spatial resolution of the thermal image decreases with increasing distance and thermographic cameras do not have great image resolution to begin with (typically 0.3 megapixels or less). Moreover, not all surfaces may be captured from an appropriate perspective. It may not even be possible to capture all areas of interest in a single thermal image, for example, when inspecting the interior of machines or buildings. The other alternative is to capture multiple thermal images of the object at close range and from different perspectives. The problem here is that the documentation and assessment of a large number of thermal images (and their supporting color images) is cumbersome and not very efficient. In either case, the estimation of heat losses or heat gains is not very accurate because 2D images contain no information about the geometry of the inspected object. 3D thermography can also address both of these issues since the 3D thermogram of the entire scene can be stored as a single file and be viewed and analyzed with a suitable software package.

6.4.2 Related work

As mentioned earlier, most 3D thermographic systems presented in the literature directly assign the false colors in the thermal image to the vertices of the generated 3D model.

Although the geometry of the 3D thermograms generated with this approach can be appreciated using point cloud editing program like Meshlab [17, 18] or CloudCompare [22], 2D projections of these 3D thermograms seen in a display during the scan or in a paper report after the inspection do not differ much from a conventional thermal image.

The system presented in [85] attempts to facilitate the identification of the inspected object by combining color and thermal data into a single image or texture; information from text and labels that identify a piece of equipment is often lost in a thermal-only representation. One problem with this approach is that it attempts to map 4 channels of information, three for color and one for temperature, onto a three dimensional color space. It is not always possible to tell if the color of a pixel represents temperature information or the actual color of a surface; color gradients in the color image can be mistaken for temperature gradients in the hybrid image. The approach might perform better if one of the color channels contains little information and can be exploited to convey thermal information, but this is rarely the case. Another problem with this approach is that it cannot be used to produce an unambiguous texture for a 3D model. The color of a surface, as recorded by a color camera, varies with the illumination of the scene and the position of the camera with respect to the surface. Which color should be combined with the color assigned to the temperature measurement if the same point is observed from multiple perspectives? Finally, this approach is restricted to cases where the scene is well illuminated.

In [78], a head mounted display is used to visualize thermal information directly mapped onto a 3D model generated in real time. This approach presents an interesting alternative to visualizing the thermal information of 3D objects but by itself does not help improve the fidelity of the 3D thermograms nor does it solve the issues mentioned at the beginning of this section.

6.4.3 Proposed approach

A person is able to distinguish the geometric shape of an object whose surface is uniformly painted with the same color because the illumination of the object by a external light source causes shading on its surface. The resulting gradients of shade provide the human brain the information it needs to infer the geometry of the object. Shading is a technique commonly used in 3D computer graphics to help create the illusion of shape and depth in a two dimensional image. In this work, shading is used to improve the discernibility of geometric shapes in a thermal image. This approach has the advantage that it can be used to convey thermal and geometric information both during the scan and on paper reports. There are several techniques to produce shading effects with

varying degree of computational complexity. In the current implementation, a computationally simple approach is preferred in order to achieve real-time operation. As in 2D thermography, the color assigned to a vertex is obtained from an indexed color palette. The minimum and maximum temperatures T_{\min} and T_{\max} corresponding to the first and last element of the color palette can be manually set by the user or they can be automatically calculated based on the temperature values associated with the current vertex map according to some criterion (min/max, 1σ , 3σ , etc.). The color index of a particular vertex is interpolated based on the temperature value associated with it. In addition, the brightness of the color is calculated as the cosine of the angle between the normal vector associated with the vertex and the vector that goes from the vertex to the depth camera. The procedure is briefly described in Algorithm 7. The choice of color palette should be restricted to color palettes that do not use brightness as the main component to convey temperature information. In this work, all thermal renderings were false colored using Matplotlib's² coolwarm colormap [33]. This colormap has little range of gray scale and the colors at the ends of the scale, as the name suggests, are commonly associated with cool and warm surfaces.

Algorithm 7 Render thermal

```

1: procedure renderThermal
2:    $T_{\min}, T_{\max} \leftarrow \text{getScaleLimits}(\mathbf{M}_v, \mathbf{V}_{\text{thermal}}, C)$   $\triangleright$  get scale limits based on criterion  $C$ 
3:   for each vertex in  $\mathbf{M}_v$  with indices  $u$  and  $v$  do
4:      $\mathbf{v} \leftarrow \mathbf{M}_v[u, v]$   $\triangleright$  vertex
5:      $\mathbf{n} \leftarrow \mathbf{M}_n[u, v]$   $\triangleright$  normal associated with vertex  $\mathbf{v}$ 
6:      $i, j, k \leftarrow \text{getVoxelIndices}(\mathbf{v})$   $\triangleright$  get voxel indices corresponding to vertex  $\mathbf{v}$ 
7:      $T \leftarrow \text{readVoxel}(\mathbf{V}_{\text{thermal}}[i, j, k])$   $\triangleright$  read temperature value from thermal volume
8:      $i \leftarrow (T - T_{\min}) / (T_{\max} - T_{\min})$   $\triangleright$  index based on temperature  $t$ 
9:      $(R, G, B) \leftarrow \mathbf{P}[i]$   $\triangleright$  color from color palette
10:     $\mathbf{r} \leftarrow \mathbf{v} - \mathbf{t}_{\text{depth}}$   $\triangleright$  vector from depth camera to vertex
11:     $k \leftarrow \text{abs}((\mathbf{r} \cdot \mathbf{n}^T) / (\|\mathbf{r}\|_2 \cdot \|\mathbf{n}\|_2))$   $\triangleright$  brightness factor
12:     $\mathbf{I}_{\text{render}, \text{Th}}[u, v] \leftarrow k(R, G, B)$   $\triangleright$  shaded false color

```

Notes:

- $\triangleright \mathbf{M}_v[u, v]$: Element $[u, v]$ of the unfiltered vertex map.
 - $\triangleright \mathbf{M}_n[u, v]$: Element $[u, v]$ of the unfiltered normal map.
 - $\triangleright \mathbf{V}_{\text{thermal}}[i, j, k]$: Voxel $[i, j, k]$ in thermal volume.
 - $\triangleright T_{\max}$: Temperature corresponding to the first color in the color palette.
 - $\triangleright T_{\min}$: Temperature corresponding to the last color in the color palette.
 - $\triangleright \mathbf{P}[i]$: i^{th} color in the color palette.
 - $\triangleright \mathbf{t}_{\text{depth}}$: Current position of the depth camera.
 - $\triangleright \mathbf{I}_{\text{render}, \text{Th}}[u, v]$: Element $[u, v]$ of the rendered thermal image
-

The parameters used to validate new thermal data can be used to generate alternative textures for the 3D model that inform the user of the proposed system about the conditions under which each temperature measurement was made. Each texture describes one measurement condition. This way, the user can easily identify which measurement conditions were not met. These additional textures are available to the user in real time

² <http://matplotlib.org/users/colormaps.html>

during the inspection so he or she has the opportunity to rescan parts of the object that were not properly scanned. The textures are generated using only four colors. Green, yellow and red are used to indicate good, acceptable and bad measurement conditions, respectively while black is used as background color and for vertices that have not been assigned a temperature value yet (for example, for object surfaces within the FOV of the depth camera but not within the FOV of the thermal camera). In other words, green, yellow and red are assigned to vertices with discrete variables d_q and θ_q equal to 1, 2 and 0, respectively, according to (6.2). The procedures to render these textures are briefly described in Algorithms 8 and 9. Here again, shading is used to improve the discernibility of the objects in the image.

Algorithm 8 Render distance to surface

```

1: procedure renderDistanceToSurface
2:   for each vertexin  $M_v$  with indices  $u$  and  $v$  do
3:      $v \leftarrow M_v[u, v]$  ▷ vertex
4:      $n \leftarrow M_n[u, v]$  ▷ normal associated with vertex  $v$ 
5:      $i, j, k \leftarrow \text{getVoxelIndices}(v)$  ▷ get voxel indices corresponding to vertex  $v$ 
6:      $d_q \leftarrow \text{readVoxel}(V_{\text{thermal}}[i, j, k])$  ▷ read distance to surface from thermal volume
7:     if  $d_q = 1$  then
8:        $(R, G, B) \leftarrow (0, 1, 0)$  ▷ color green
9:     else if  $d_q = 2$  then
10:       $(R, G, B) \leftarrow (1, 1, 0)$  ▷ color yellow
11:     else
12:       $(R, G, B) \leftarrow (1, 0, 0)$  ▷ color red
13:      $r \leftarrow v - t_{\text{depth}}$  ▷ vector from depth camera to vertex
14:      $k \leftarrow \text{abs}((r \cdot n^T) / (||r||_2 \cdot ||n||_2))$  ▷ brightness factor
15:      $I_{\text{render,DT}}[u, v] \leftarrow k(R, G, B)$  ▷ shaded false color

```

Notes:

- ▷ $M_v[u, v]$: Element $[u, v]$ of the unfiltered vertex map.
 - ▷ $M_n[u, v]$: Element $[u, v]$ of the unfiltered normal map.
 - ▷ $V_{\text{thermal}}[i, j, k]$: Voxel $[i, j, k]$ in thermal volume.
 - ▷ t_{depth} : Current position of the depth camera.
 - ▷ $I_{\text{render,DT}}[u, v]$: Element $[u, v]$ of the rendered distance to surface
-

The issues of visualizing large objects with complex geometry and of estimating heat losses and/or heat gains can be solved using a point cloud representation of the 3D thermogram. The open source implementation of the Kinect Fusion algorithm provides a procedure to generate a point cloud from the data in the TSDF volume. In the proposed system, this point cloud has been complemented with channels that contain the temperature values and the discretized measurement conditions associated with the points. The false colors can be calculated in real time during the scan or after the scan by means of a false coloring script. The second approach has the advantage that the color palette and scale limits can be modified at any time. The procedure to generate the extended point cloud is briefly described in Algorithm 10.

Algorithm 9 Render emission angle

```

1: procedure renderEmissionAngle
2:   for each vertex in  $M_v$  with indices  $u$  and  $v$  do
3:      $v \leftarrow M_v[u, v]$  ▷ vertex
4:      $n \leftarrow M_n[u, v]$  ▷ normal associated with vertex  $v$ 
5:      $i, j, k \leftarrow \text{getVoxelIndices}(v)$  ▷ get voxel indices corresponding to vertex  $v$ 
6:      $\theta_q \leftarrow \text{readVoxel}(V_{\text{thermal}}[i, j, k])$  ▷ read emission angle from thermal volume
7:     if  $\theta_q = 1$  then
8:        $(R, G, B) \leftarrow (0, 1, 0)$  ▷ color green
9:     else if  $\theta_q = 2$  then
10:       $(R, G, B) \leftarrow (1, 1, 0)$  ▷ color yellow
11:     else
12:       $(R, G, B) \leftarrow (1, 0, 0)$  ▷ color red
13:      $r \leftarrow v - t_{\text{depth}}$  ▷ vector from depth camera to vertex
14:      $k \leftarrow \text{abs}((r \cdot n^T) / (\|r\|_2 \cdot \|n\|_2))$  ▷ brightness factor
15:      $I_{\text{render,EA}}[u, v] \leftarrow k(R, G, B)$  ▷ shaded false color

```

Notes:

- ▷ $M_v[u, v]$: Element $[u, v]$ of the unfiltered vertex map.
 - ▷ $M_n[u, v]$: Element $[u, v]$ of the unfiltered normal map.
 - ▷ $V_{\text{thermal}}[i, j, k]$: Voxel $[i, j, k]$ in thermal volume.
 - ▷ t_{depth} : Current position of the depth camera.
 - ▷ $I_{\text{render,EA}}[u, v]$: Element $[u, v]$ of the rendered emission angle
-

Algorithm 10 Generate extended point cloud

```

1: procedure generatePointCloud
2:    $C_{\text{in}} \leftarrow \text{generateBasePointCloud}(V_{\text{tsdf}})$  ▷ generate base point cloud
3:   for each point in  $C_{\text{in}}$  with index  $s$  do
4:      $i, j, k \leftarrow \text{getVoxelIndices}(C_{\text{in}}[s])$  ▷ calculate voxel indices
5:      $T, d_q, \theta_q \leftarrow \text{readVoxel}(V_{\text{thermal}}[i, j, k])$  ▷ read voxel info from thermal volume
6:      $C_{\text{out}}[s].\text{point} \leftarrow \text{getCentroid}(i, j, k)$  ▷ assign vertex coordinates
7:      $C_{\text{out}}[s].\text{temperature} \leftarrow T$  ▷ assign temperature information
8:      $C_{\text{out}}[s].\text{distance} \leftarrow d_q$  ▷ assign information regarding distance to target
9:      $C_{\text{out}}[s].\text{angle} \leftarrow \theta_q$  ▷ assign information regarding emission angle

```

Notes:

- ▷ V_{tsdf} : TSDF volume.
 - ▷ $V_{\text{thermal}}[i, j, k]$: Voxel $[i, j, k]$ in thermal volume.
 - ▷ $C_{\text{in}}[s]$: Element s of the base point cloud
 - ▷ $C_{\text{out}}[s]$: Element s of the extended point cloud
-

6.4.4 Experimental results

Figure 6.12 illustrates how shading significantly improves the discernibility of the constituent parts of the infrared calibrator for the same choice of color palette and temperature range. This is demonstrated for two different scenarios, when the temperature range is small (top row) and when the temperature range is large (bottom row). The left column shows the conventional thermal image and the center and right columns show renderings of the 3D model without and with shading, respectively. Aside from a few black points at the edges of the objects, which are not properly modeled by the

TSDF volume and unintentionally expose the edges of some shapes, there is no significant difference between the conventional thermal image and the rendering of the 3D model without shading. Due to the larger temperature scale, the thermal image and the rendering without shading in the bottom row reveal slightly more information about the geometry of the scene than their counterpart in the top row; for example the control panel or the ventilation slots. Even so, much of the geometric information is still difficult to discern. The rendering with shading reveals much more information about the geometry of the scene in both scenarios.

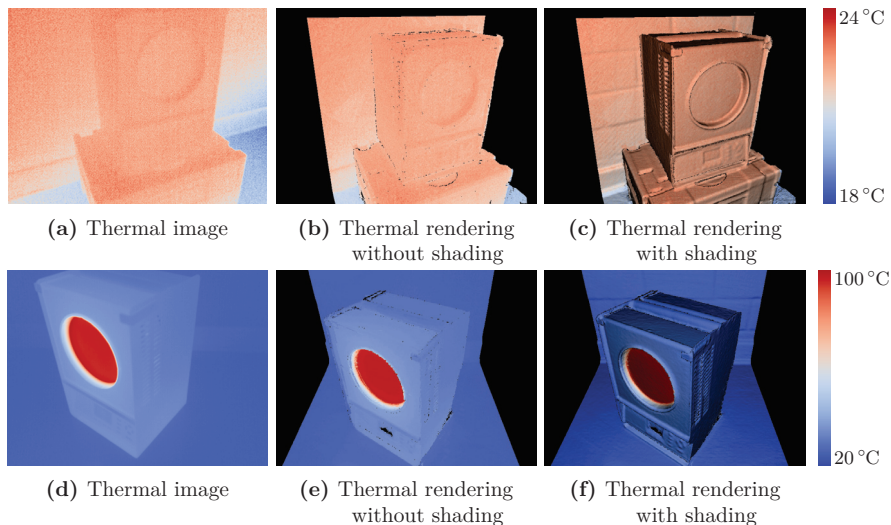


Figure 6.12: Effect of shading for small and large temperature ranges

Figure 6.13 shows two renderings of the infrared calibrator's 3D model, one texturized with temperature values according to Algorithm 7 and the other with distance values according to Algorithm 8. The temperature of the calibration disc was set to 50°C and a resolution pattern was placed in front of the disc to highlight the effects of scanning different parts of a scene from different distances. The green and yellow colored surfaces in Figure 6.13b show that at the time of the snapshot only a small part of the calibrator's surface was scanned from an appropriate measurement distance. In Figure 6.13a the thermal texture corresponding to this part of the calibrator's surface looks sharp and the temperature value it represents is within the accuracy range specified by the camera. The rest of the surfaces are depicted with a low resolution thermal texture. This is not a problem where the surface temperature is constant but it inaccurately depicts large temperature gradients. While in this experiment it is evident from the thermal rendering alone that different parts of the infrared calibrator have been scanned from different distances, this might not be the case when the areas of interest are not contiguous. In

any case, the spatial resolution of the thermal texture is limited by the spatial resolution of the geometric model.

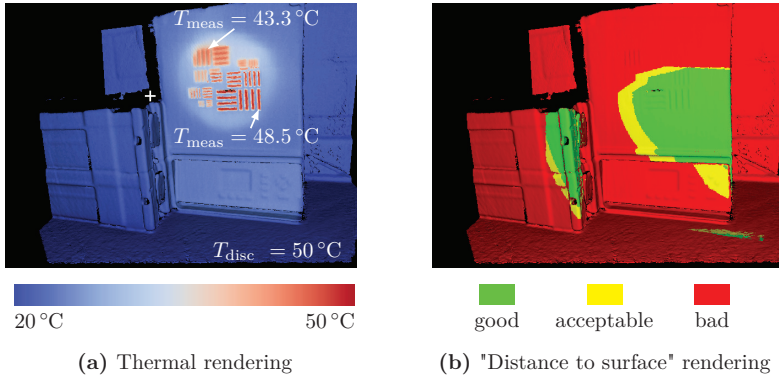


Figure 6.13: Visualizing the measurement distance for each temperature value

Similarly, Figure 6.14 shows two renderings of the infrared calibrator's 3D model, one texturized with temperature values according to Algorithm 7 and the other with emission angle values according to Algorithm 9. The green and yellow colored surfaces in Figure 6.14b indicate temperature measurements that have an error due to the directional nature of the emissivity smaller than 2 %. Red colored surfaces inform the system's operator that the measurement error is greater than 2 % and suggest that the surfaces should be rescanned from a different angle. The temperature assigned to the calibration disc in the 3D thermogram is 92.5°C because the angle formed by its normal and the infrared ray that reach the thermographic camera is large. The true temperature of the calibration disc is 100°C .

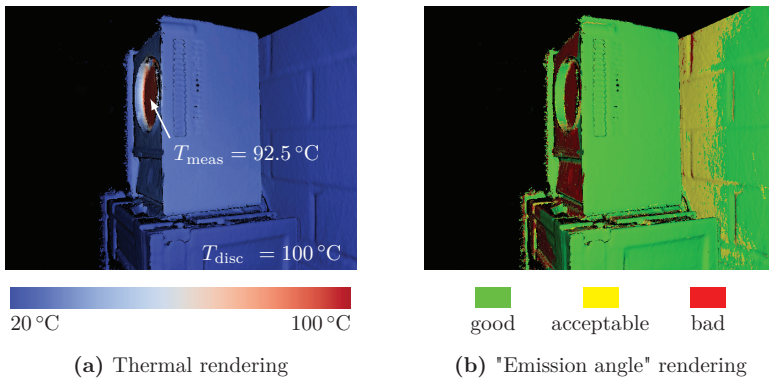


Figure 6.14: Visualizing the emission angle for each temperature value

6.4.5 Discussion

The experiments show that adding shading to a thermal image greatly facilitates the differentiation of geometric shapes in the observed scene. Because shading provides the human brain with information about the orientation of the different surfaces relative to each other, this holds true not only for the infrared calibrator but for any object with complex geometry. However, adding shading to a thermal image implies changing the brightness of its pixels. As a result, some of the colors in the thermal rendering will no longer be found in the temperature scale that accompanies the image; for example, the color of the infrared calibrator's right side in Figure 6.12c does not correspond to any color in the temperature scale. This is not really a problem since a thermal image or a thermal rendering is not meant to provide an exact temperature value based on the color of its pixels but rather to reveal the distribution of temperature gradients at the object's surface. It is important, however, to choose a color map that does not strongly rely on changes in brightness to convey temperature differences; shading a grayscale image, for example, could introduce fictitious temperature gradients on curved surfaces with uniform temperature. Moreover, the user can always minimize the effects of shading of a particular surface of interest by positioning the camera in such a way that its optical axis is perpendicular to the surface. In addition, if the user wants to know the temperature value measured for a particular point on object's surface, a cursor can be enabled and the temperature of the pixel under the cursor will be shown in the display, as it is done in conventional 2D thermography.

In some cases, it may be necessary to inspect a series of identical objects that only differ from each other by a label or a printed number or text and this information may be lost in a thermal-only representation. In such cases, the visual data that identify the object can overlap or directly replace the thermal data in the thermal rendering or the texture. This is basically the same technique used in bi-spectral thermal cameras.

The experiments have also shown that the use of alternative textures to visualize the measurement conditions for each temperature value can help obtain more reliable measurements of an object's surface temperature. The algorithm that validates and updates the thermal volume guarantees that only the temperature measurements made under the best conditions are used in the generation of the 3D thermogram, but it is the operator of system who enables these conditions. The alternative textures provide the visual feedback the operator requires to achieve favorable conditions. Additional measurement conditions can be visualized in a similar way by means of new texture layers. The user can easily switch between the different layers during the scan, for example, by pressing a button.

It is worth noticing that the sharpness of the thermal rendering and the ability to discern shapes in it strongly depends on the spatial resolution of the 3D model as well as on the spatial resolution of the thermal texture. These two factors depend in turn on the spatial resolution of the depth and thermal images, respectively, and on the voxel size of the [TSDF](#) volume. The spatial resolution of the depth and thermal images improves with decreasing distance within the operating range of the depth sensor and the [DOF](#) of the thermal camera.

6.5 Summary

This chapter elaborates on the key aspects concerning the generation, representation and visualization of high-fidelity 3D thermograms.

Section [6.1](#), 3D model generation, explains the principle of operation of the open source version of the Kinect Fusion algorithm. It also presents two minor but helpful modifications made to it. The first modification allows the user to know whether the object is within the scanning range of the depth sensor and whether this range intersects the cubic volume defined at the beginning of the scan. The second modification reduces the accumulation of errors in the 3D model by preventing the integration of new depth data when the sensor is abruptly moved or when the pose change estimate is abnormally large due to noisy data. This section also evaluates the effect that the voxel size has on the quality of the 3D model. It shows that, for the 3D thermal imaging system used in this work, higher resolution does not necessarily mean higher fidelity. The voxel size must be larger than the precision of the depth measurements so that the Kinect Fusion algorithm can reduce the measurement errors through averaging within a voxel.

Section [6.2](#), data fusion, describes how the temperature values measured with the thermographic camera are assigned to the 3D model. This section addresses two main issues: the estimation of the thermographic camera pose with respect to the 3D model and the parallax effect. The pose of the thermographic camera is calculated via interpolation from two consecutive poses of the depth camera and using the known fixed geometric transformation between the cameras. An analysis of this approach shows that, for the proposed operating conditions, the error in assuming a linear instead of a curved trajectory is negligible. In addition, an extended analysis presented in [appendix F](#) calculates the magnitude of the expected registration errors caused by data synchronization errors and by the [ICP](#) algorithm when it is assumed that its accuracy is 10 %. These errors are comparable to the size of the [IFOV](#) of the thermographic camera and are therefore tolerable. The errors caused by the parallax effect are minimized by filtering the 3D model

vertices generated from the perspective of the depth camera before assigning temperature values to them. The proposed filtering algorithm keeps only the vertex closest to the thermographic camera from each group of vertices that project onto the same pixel in the thermal image. The algorithm significantly improves the fidelity and appearance of the generated 3D thermogram. Vertices at the edges of objects can be problematic, however, because they are difficult to reliably detect with the depth camera.

Section 6.3, data validation, introduces a new approach to improve the fidelity with which 3D thermograms depict the surface temperature of an object. The approach prevents previous temperature measurements from being overwritten by new temperature measurements when the latter are taken under worse measurement conditions. It exploits the available geometric information to validate the thermal data before they are integrated into the existing 3D thermogram. The experiments show that 3D thermograms generated using data validation depict the surface temperature of the inspected object with much higher fidelity than without data validation. The current implementation only validates thermal data in terms of the emission angle and the distance between the object's surface and the thermographic camera; however, the validation of other influencing factors, like linear and angular scan velocity, is also possible.

Finally, section 6.4, data representation and visualization, proposes the combined use of shading and a color palette with little range of gray scale to improve the discernibility of an object's features in renderings of the 3D thermogram. It also proposes the use of additional textures to convey visual information regarding the conditions under which the temperature measurements were taken. These additional textures provide the operator with real-time feedback about the quality of the measurements and allows him/her to immediately rescan surfaces that were not scanned under the desired conditions. The acquisition of good quality data is essential for the generation of high-fidelity 3D thermograms.

Chapter 7

3D Thermal imaging: case study

Conventional thermographic cameras are used in a large number of applications. They are used, for instance, to inspect the thermal insulation of industrial furnaces in order to evaluate their energy efficiency. Industrial furnaces are typically large objects with complex geometries that exhibit high thermal inertia and wide temperature ranges. Because of these characteristics they pose an interesting case study for evaluating the performance of a 3D thermal imaging system. This chapter presents the results of scanning an operating industrial furnace with the 3D thermal imaging system developed in this work and discusses the problems associated with the generation of 3D thermograms of large objects with complex geometries.

7.1 Scanning an industrial furnace

The 3D thermal imaging system developed in this work was used to scan the access door of an operating industrial furnace. The total scanned volume was approximately $1.6\text{ m} \times 1.7\text{ m} \times 0.5\text{ m}$ ($H \times W \times D$). The system was configured to work with a voxel count of $480 \times 480 \times 480$ voxels to be able to generate the 3D thermograms in real time. The volume size, on the other hand, was set to $0.96\text{ m} \times 0.96\text{ m} \times 0.96\text{ m}$ to work with a cubic voxel of edge length 2 mm. This voxel size offers a good compromise between the amount of reconstruction noise caused by the depth measurements errors, scan volume size and the spatial resolution of the 3D model.

The access door was scanned in parts and from multiple perspectives. Scan distances ranged from 0.3 m to 1 m. To facilitate the process of merging the 3D thermograms generated from each scan, the scanned surfaces overlapped. Raw time-stamped depth and thermal images were saved to disk for post-processing. The scan duration was limited

to 20 s. The thermal camera was allowed to warm up for at least 30 min before the first scan. A total of 35 scans were carried out.

The 3D thermograms were generated from the scans using the validation parameters: $d_{\min} = 35$ cm, $d_{\text{th}} = 65$ cm, $d_{\max} = 95$ cm, $\theta_{\min} = 15^\circ$, $\theta_{\text{th}} = 30^\circ$ and $\theta_{\max} = 50^\circ$ and the color scale limits: $T_{\min} = 30^\circ\text{C}$ and $T_{\max} = 270^\circ\text{C}$. The 3D thermograms were manually cleaned, merged and meshed using the point cloud editing software Meshlab. The complete 3D thermogram is shown in Figure 7.1 from multiple perspectives.




Figure 7.1: 3D thermogram of the access door of an operating industrial furnace

Due to the presence of different types of errors, the manual composition of a large 3D thermogram from smaller individual 3D thermograms is a difficult and time-consuming task. One of the main problems is volume closure. Perfectly aligning all point clouds is not possible because all point clouds are slightly distorted. Alignment errors are in general smaller than two millimeters between adjacent and overlapping point clouds but can add up to 1 cm across multiple point clouds. Figures 7.2a and 7.2b show two point clouds generated from two independent scans. Because these two point clouds do not overlap much they are aligned with respect to nearby point clouds with which they share large overlapping sections. Figures 7.2c shows a clear misalignment error of about 1 cm between the upper beams.

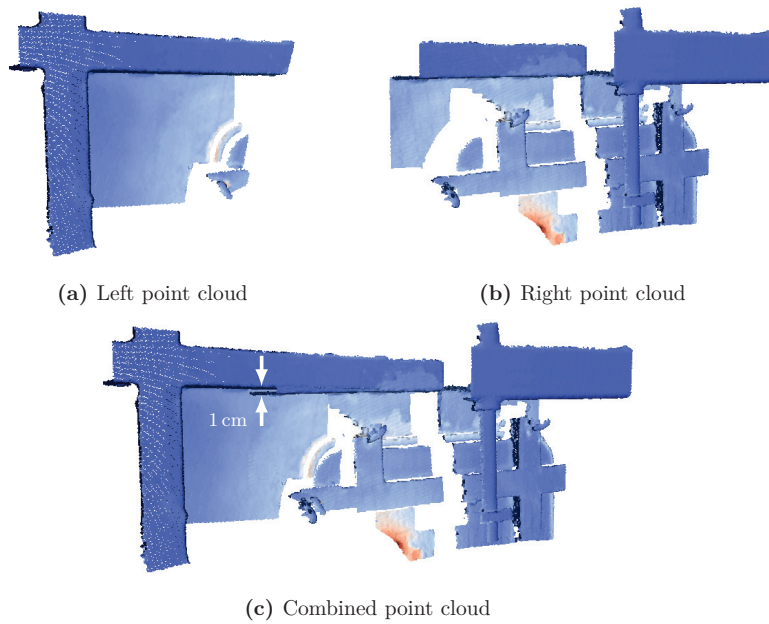


Figure 7.2: Registration errors due to depth measurement errors

Another main problem is parallax. The parallax filter proposed in section 6.2 works well as long as the depth camera is able to detect the points closest to the thermal camera that are involved in the parallax effect. Figure 7.3 illustrates one scenario where this condition is not met. The figure presents three sets of images captured at different points in time during the generation of the 3D thermogram. The first two were generated at the beginning of the scan and the third one towards the end. Each set consists of a thermal image, a rendering of the 3D model and a rendering of the 3D thermogram. Until $t = 1.2$ s, the parallax filter is able to prevent most vertices behind (and above) the plate from being assigned the temperature value measured for the plate. (A few vertices near the edges are assigned the wrong temperature value but this occurs because edges cannot

be determined and represented very accurately in the current system implementation and because the area represented by a single thermal pixel at this distance is larger than the spatial resolution of the 3D model.) At $t = 1.6$ s, part of the plate has left the FOV of the depth camera but not the FOV of the thermographic camera and registration errors occur. At the end of the scan, the accumulated registration errors can be clearly appreciated.




Figure 7.3: Registration errors due to parallax when an object in the foreground leaves the FOV of the depth camera

Another scenario where the parallax filter fails to work is illustrated in Figure 7.4. Here, the base of the shaft that holds the access door is too close to the depth camera for it to be detected correctly and cannot be integrated into the 3D model. As a result, its temperature value, as measured by the thermographic camera, is assigned to the vertices that describe the surface in the background.

The third main problem is caused by reflections. Figure 7.5 shows the surface behind the shaft that holds the access door from four different perspectives. In Figures 7.5a and 7.5c the temperature measured by the thermographic camera is approximately 80°C whereas in Figures 7.5b and 7.5d it is about 130°C . This substantial difference in temperature




Figure 7.4: Registration errors due to parallax when an object has not yet entered the operating range of the depth camera

readings is caused by the reflection of the hot pipe onto the door's surface. In Figure 7.5b the reflection is oriented towards the upper right corner and in Figure 7.5d towards the lower right corner.




Figure 7.5: Registration errors due to reflections

7.2 Discussion

In principle, the entire volume of the access door could have been scanned using the same voxel length of 2 mm and a voxel count of $800 \times 850 \times 250$ voxels which is smaller than the maximum voxel count of $640 \times 640 \times 640$ voxels that the system's GPU is able to manage in terms of memory. Nevertheless, it was preferred to scan the furnace access door in parts for several reasons. First, the scanned object contains many parallel flat surfaces with few or no features. ICP is likely to estimate the system's pose with less accuracy or even lose track of the object if the transition from one perspective to another includes one of these regions. Second, the current implementation of the system skips depth and thermal data frames when working in real time with voxel counts larger than $480 \times 480 \times 480$ voxels. This can also cause the loss of tracking of the object and the corruption of the model. Third, the larger the volume, the larger the accumulation of errors during the generation of the 3D model and the harder it is to align a surface with itself when the trajectory of the system forms a closed loop. Fourth, scanning a large object in one pass would result in very large files containing depth and thermal data since, in the current implementation, these are saved in raw format. The first issue could be solved by incorporating other mechanisms to determine the system's trajectory, for example, visual odometry or inertial measurements. The second issue can be solved by optimizing the code and by using a more powerful GPU. The third issue could be addressed using loop closure methods like the one proposed in [89]. The fourth issue can be addressed using video compression algorithms.

The errors due to parallax depicted in Figure 7.3 can be further reduced by generating a separate vertex map from the perspective of the thermal camera. This requires optimizing the code and using a more powerful GPU because ray casting is a computational expensive process but it also requires a new way of describing the 3D model that allows the generation of vertex maps from any perspective and not just from the perspective of the depth camera as is the case with the TSDF volume. The incorrect temperature registration behind the edges of objects could also be prevented by using dilation operations on the vertex map before the application of the parallax filter. Increasing the resolution of both the depth and thermal cameras would also reduce this type of error. The errors due to parallax depicted in Figure 7.4 can be easily avoided by keeping a minimum distance between the scanned object and the 3D thermal imaging system.

Unfortunately, the inclusion of false temperature values caused by reflections into the thermogram cannot always be avoided. In the scenario depicted in Figure 7.5, the reflections do not come from objects in the background but from the inspected object itself. This problem is more conspicuous in 3D thermograms because the same surface may be scanned multiple times from different perspectives. Depending on the strategy used to

combine the measurements, the resulting thermal texture may end up resembling a collage of temperature measurements. Nevertheless, the incongruence of the temperature measurements can be exploited to identify reflective surfaces of an object. This information could be depicted, for example, as an additional texture of the 3D model in which the color assigned to a particular vertex depends on the variance of the temperature measurements associated with that vertex. A strategy to improve the visual appearance of the 3D thermogram could be to keep the lowest temperature measurement made for the same point, as this represents the measurement with the lowest component of reflected thermal radiation (at least within the range of emission angles in which the emissivity is relatively constant.)

7.3 Summary

This chapter presents the results of scanning the access door of an operating industrial furnace with the 3D thermal imaging system developed in this work and discusses the main problems encountered when trying to generate a large complex 3D thermogram from smaller overlapping 3D thermograms. It is argued that most of the problems can be solved using better hardware, optimizing the code or including new methods developed by other researchers and that the problem caused by reflections cannot be avoided but can at least be made known to the user.

Chapter 8

Conclusions and outlook

8.1 Conclusions

The main objective of this work was to develop a system, methods and strategies that take advantage of new technological advancements to generate 3D thermograms that depict the surface temperature of objects with higher fidelity than 2D or 3D thermograms generated with other thermographic systems. To this end, in-depth analyses of the different theoretical and practical aspects that affect the fidelity and visual representation of thermograms were carried out and new methods and strategies to improve the current state of the art were developed.

Several hardware configurations based on different combinations of sensors were considered for the implementation of a 3D thermographic system prototype. For close range 3D thermography, a solution based on a thermographic and a depth camera, both light and compact, was preferred. The small size of the sensors allows them to be mounted very close to each other. This characteristic helps maximize the **FOV** common to both sensors and minimize registration errors due to the parallax effect. Their little weight, on the other hand, makes the whole 3D thermographic system easy to maneuver. This property greatly facilitates the acquisition of data from favorable perspectives. Furthermore, the depth camera generates a dense point cloud, which is crucial for reconstructing accurate 3D models and for accurately estimating the pose of the 3D thermographic system with respect to the inspected object.

It was argued that a real-time high-fidelity 3D thermographic system requires at least three independent calibration procedures: radiometric, geometric and temporal calibration.

With regard to the radiometric calibration of the thermographic camera, a series of original experiments presented in section 5.1 showed that the radiometric calibration carried out at the factory does not take into account that the signal produced by an individual detector element in the thermographic camera's FPA is influenced by the thermal radiation that impinges on other detector elements in the array. This behavior is not particular to the Optris PI 450 but was also observed with the high-end thermographic camera VarioCAM hr head 600 and is very briefly described in the literature. The analysis of the experimental data suggests that the measurement deviations are caused by thermal conduction between the detector elements and the array's substrate. Because the total amount of thermal radiation that enters the thermographic camera changes constantly as the 3D thermal imaging system is moved around the inspected object, this behavior should be taken into account when generating high-fidelity 3D thermograms. The relevance of the measurement errors will depend on the degree of accuracy that is needed. It could play an important role, for example, in medical applications where small errors in the temperature measurements can lead to false diagnostics or in scenarios where high temperature contrasts are present. In theory, it could be possible to model the energy exchange between the elements of the FPA and use this model to compensate for the measurement errors but this would require detailed knowledge about the physical makeup of the FPA.

To facilitate the fusion of depth and thermal data, the sensors must be geometrically calibrated. In section 5.2, an intrinsic and extrinsic calibration of the system was carried out with an improved calibration target and following a procedure that is more rigorous than other procedures found in the literature. Although the most influential parameters were estimated with high precision and the MRE was very low, manual fine-tuning of some of the calibration parameters was still necessary in order to properly map the thermal texture onto the 3D model. An analysis of this issue revealed that the margin of error in the estimation of the camera orientation is considerable but, because these errors are compensated by corresponding errors in the estimation of the intrinsic parameters c_u and c_v , the MRE remains very low. This means that, with the current approach, the pose of one camera relative to the other cannot be estimated with enough precision.

Because the sensors are not synchronized and do not allow for external hardware synchronization, they had to be synchronized via software. Section 5.3 provides a rigorous analysis of the proposed synchronization approach which is based on a median filter. In this approach, the clock of the host computer is used to synchronize the timestamps generated by the sensors' internal clocks. The filter eliminates sporadic large differences between the timestamps. These are caused by random delays associated with the transfer of new data from the sensors to the computer and can cause noticeable registration errors. It was shown that the timing errors after filtering cause registration errors that

are comparable to the size of thermographic camera's **IFOV** for the specified operating range and can therefore be tolerated.

It was also argued that the generation of the 3D thermogram should be carried out in real time in order to acquire the best possible data. The main steps involved in generation of 3D thermograms are the reconstruction of the 3D model, the determination of the correspondence between the vertices of the 3D model and the temperature values measured by the thermographic camera and the validation of the new thermal data.

The 3D model was generated using the Kinect Fusion algorithm. The principle of operation of the open source implementation of the algorithm is described in section 6.1. The algorithm was modified to integrate depth data only when the change in position or orientation between consecutive depth frames is within predefined limits. This small modification helped reduce the accumulation of registration errors over time. It was also observed that there is an optimal voxel size with which objects are best represented. The use of smaller voxel results in noisy and distorted 3D models while the use of larger voxel results in excessively smoothed 3D models and loss of detail. Because new truncated signed distances are averaged with previous values before they are assigned to their corresponding voxels, it helps to contain most of the measurement uncertainty within the same voxel so it is reduced through averaging. In other words, the voxel edge length should be greater than the measurement uncertainty of the depth sensor.

In section 6.2, it was argued that the most effective way to figure out the correspondence between the vertices of the 3D model and the temperature values measured by the thermographic camera is to project onto the thermal image plane only the vertices seen by the thermographic camera at the time a new thermal image is generated. One challenge in this approach is to estimate the pose of the thermographic camera with respect to the 3D model. Because it cannot be estimated reliably and with enough accuracy directly from the thermal images, it must be derived from the pose of the depth camera. It was shown that for reasonable scan velocities the error in estimating the pose of the thermographic camera using linear interpolation instead of some curved trajectory is negligible. Furthermore, an in-depth analysis shows how time synchronization errors and pose change estimation errors introduced by the **ICP** algorithm affect the mapping of the thermal data onto the 3D model. It was calculated that time synchronization errors smaller than 6 ms and pose change estimation errors smaller than 10% cause tolerable registration errors that are comparable to the size of thermographic camera's **IFOV** within the specified operating range. The other challenge is that in the current implementation of the open source Kinect Fusion algorithm the vertices of the 3D model can only be generated from the perspective of the depth camera when a new depth image is captured. Because depth and thermal images are captured at different times

and from different positions in space, they do not depict the surfaces of objects from the same perspective. Consequently, registration errors due to the parallax effect are introduced. A filter to eliminate these errors was proposed and implemented. From each set of vertices that project onto the same pixel in the thermal image, the filter keeps only those vertices closest to the thermographic camera. The filter significantly improves the fidelity and appearance of the generated 3D thermogram but only works as long as the depth camera is able to detect the vertices closest to the thermal camera that are involved in the parallax effect. This means that small registration errors due to the parallax effect can still be observed, for example, near the edges of the objects.

The accuracy of the temperature measurements made with a thermographic camera depends on the position of the camera relative to the observed surface. This means that temperature values with varying degrees of accuracy will be captured for the same point on the object's surface as the 3D thermographic system is moved around. To make sure that the thermal texture of the 3D thermogram is generated with the most accurate temperature measurements a novel algorithm that verifies the conditions under which new measurements are taken was proposed in section 6.3. The algorithm exploits the available geometrical information to decide when to update the thermal texture of the 3D model. The tests showed that more accurate thermal textures are obtained with the use of this algorithm.

Finally, the ability of the 3D thermographic system to generate high-fidelity 3D thermograms depends to a large extent on the quality of the data that is supplied to it. The user of the system plays an important role in this regard. To facilitate the use of the system, various types of visual feedback are provided to the user. On the one hand, the user is able to see at any time during the scan the position and orientation of the 3D thermographic system with respect to the scan volume as well as the effective **FOV** of the depth sensor (see Figure 6.2). On the other hand, additional textures of the 3D model inform the user in real time which surfaces of the object have been scanned under the desired measurement conditions and which parts need to be rescanned from a more favorable perspective (see Figure 6.13 and Figure 6.14). In addition, it was shown, that the use of shading and a color palette with little range of gray scale improves the discernibility of object's features in renderings of the 3D thermogram.

8.2 Outlook

Even in the research and development stage, the 3D thermal imaging system developed in this work was able to generate 3D thermograms with very good fidelity. Nonetheless,

there are still many aspects related to the generation of 3D thermograms that have potential for improvement.

Non-uniformity correction: In the current implementation, the thermographic camera's [NUC](#) procedure is triggered at the beginning of a scan and remains disabled during the scan. It is assumed that the camera has had enough time to warm up and that the camera housing temperature remains stable during the short scan duration. A more robust approach would allow [NUC](#) procedures during the scan and disable the update of the thermal volume while they are executed. The system would nonetheless continue updating the 3D model to avoid losing track of the object. Depending on the capabilities of the thermographic camera, this could be achieved, for example, by letting the camera decide when to carry out the procedure and monitoring a flag that indicates when the thermal images are invalid or by disabling the camera's automatic [NUC](#) procedure and letting the system trigger the procedure when the temperature change of the camera housing has exceeded some predefined threshold. In this case, the duration of the procedure must be known to the system.

Radiometric calibration: During the development and test of the 3D thermal imaging system, it was observed that the temperature value measured by the thermographic camera at a particular point on an object's surface is affected by the radiation emitted by the surrounding surfaces. In [section 5.1](#), a series of experiments were carried out in order to understand this behavior and quantify the magnitude of the measurement error. It was argued that this effect is caused by the conduction of heat between neighboring elements in the detector array. A possible next step would be to develop a mathematical model that describes this heat exchange. This would require knowing the geometrical structure and the thermal properties of the detector elements in the [FPA](#). This model could be used to correct the measurement errors caused by thermal conduction via software and thus improve the accuracy of the temperature measurements over a broader range of measurement conditions.

Geometric calibration: Despite the seemingly accurate camera calibration results with the current approach, it is still necessary to carry out a manual fine-tuning of some of the intrinsic and extrinsic parameters in order to correctly register depth and thermal data. It would be interesting to see whether other target features and feature detectors can improve the accuracy of the estimation of the camera orientation and the intrinsic parameters c_u and c_v .

Self-localization: Currently, the integration of a new point cloud generated by the depth sensor with the existing 3D model depends on the accurate estimation of the new depth sensor's pose. The sensor's pose estimation, however, relies on the accuracy of the 3D model generated so far. A consequence of this mutual dependence is that errors

accumulate over time. One option is to use the thermographic camera to detect and track features on the scene and auto locate itself by solving the Perspective-n-Point problem. This approach has been proposed in [83]. One drawback with this solution is that it only works in scenarios that are rich with well-defined thermal features. Another drawback is that the pose estimate is far less accurate than that provided by the Kinect Fusion algorithm due to the low resolution of thermal images. This option can nonetheless be useful as a fail-safe mechanism in cases where the ICP algorithm momentarily loses track of the scene. During this period of time, the thermographic camera can roughly track the position of the system with respect to the scene without integrating new data into the existing 3D thermogram until the ICP algorithm is able to get back on track. It would be interesting to see alternative options that actually improve the accuracy of the self-localization under a broad range of measuring conditions.

3D model reconstruction: In the current implementation of the open source Kinect Fusion algorithm, the scan volume and the spatial resolution of the 3D model are limited by the memory and processing power of the computer's graphic card. This limits the applicability of the current system to the generation of 3D thermograms of small objects. In order to scan large machines or buildings, the algorithm has to be modified to move data corresponding to parts of the scene that are not currently within the FOV of the system to CPU memory and to the hard drive. Work in this direction has already been demonstrated in [88–90] but not for a 3D thermal imaging system. On the other hand, in the current implementation, the description of the vertices that define the 3D model varies with the position of the depth sensor with respect to the model. Ideally, the description of these vertices should depend only on the initial definition of the 3D model's reference frame and not vary with the position of the depth sensor thereafter. It would be interesting to see the quality of 3D models that can be achieved this way.

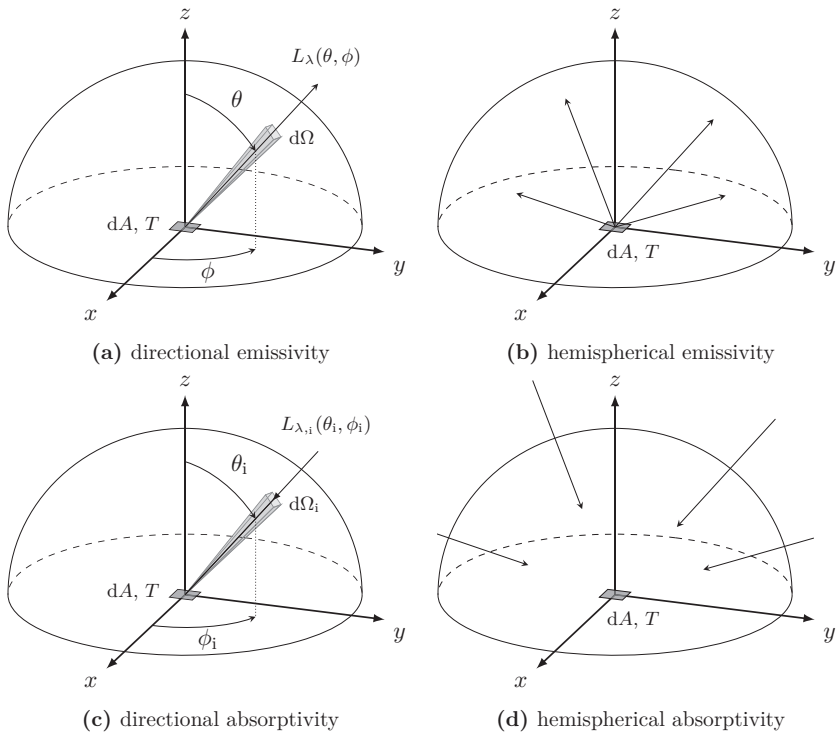
Data validation and visualization: Currently, thermal data are validated based on the angle of emission and the distance to the target before they are integrated into the 3D thermogram. This information is available to the user in form of additional textures of the 3D model. In this way, the user can find out in real time whether a surface was scanned under the desired measurement conditions and rescan if not. Another texture could be added to the system to inform the user whether the exposure time of a particular surface to the thermographic camera satisfies the minimum requirement. This information can be derived from the linear and angular velocity with which the system is moved. Similarly, another texture could be used to highlight surfaces of an object whose temperature values vary significantly as they are scanned from different positions even when the measurements are taken under good measurement conditions. Such texture would reveal reflective surfaces whose temperature values should not be trusted.

Automation: In the current implementation, the user decides how to move and rotate the 3D thermographic system in order to generate the desired 3D thermogram. It can use the additional textures of the 3D model as feedback to guarantee that all measurements are made under the appropriate measurement conditions. However, when scanning very large structures like buildings at close range, this monotonous activity can become tedious. Such a task is an ideal candidate for automation, as it would allow the use of human resources in a more effective way. It would be interesting to develop inspection planning algorithms based on the feedback provided the proposed 3D thermographic system to autonomously carry out the thermographic inspection. Together with the 3D thermal imaging system, these algorithms could be used for example by a drone to create 3D thermogram of buildings without human intervention.

Appendix A

Radiative properties

Figure A.1 illustrates the definitions of radiative properties presented in section 2.2.5.



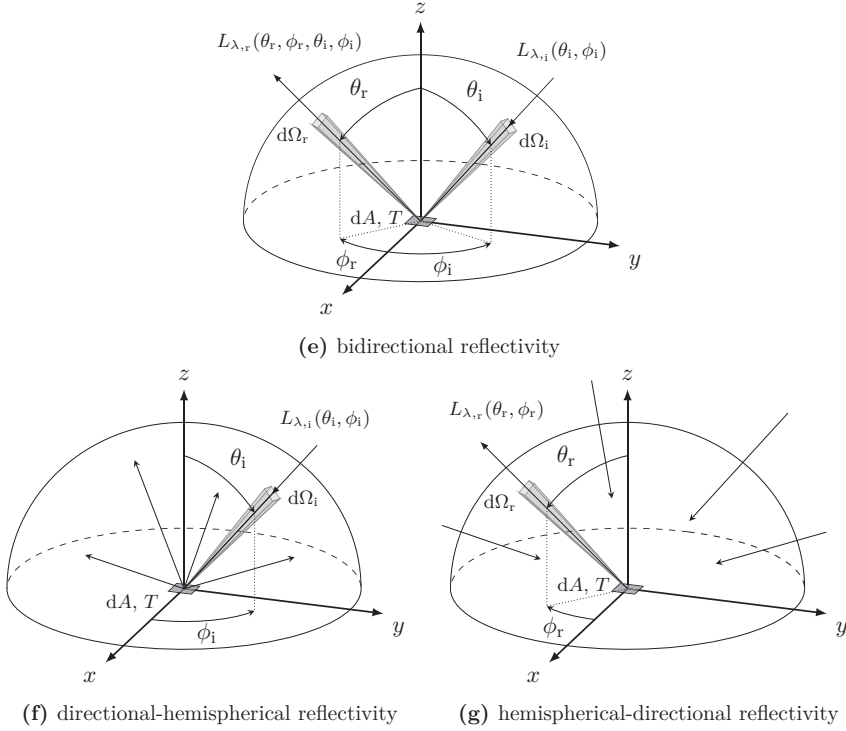


Figure A.1: Radiative properties

Appendix B

Reflection and refraction at an interface

Snell's law and the Fresnel equations describe the behavior of electromagnetic waves at the boundary between two media with different refractive indices. For convenience, this appendix presents the formula that relates the angles of incidence and refraction as well as formulas to calculate the reflectivity of a material with complex refractive index at its interface with air or vacuum. A more detailed explanation of these and other related equations can be found in [29, 39].

B.1 Snell's law

Snell's law describes the change of direction that an electromagnetic wave experiences as it crosses the boundary between two media with different refractive indices. This scenario is illustrated in Figure B.1 and is described mathematically by (B.1).




Figure B.1: Refraction of light at an interface

$$n_1 \sin \theta_1 = n_2 \sin \theta_2 \quad (\text{B.1})$$

n_1, n_2 refractive indices of the two media

θ_1 angle of incidence (measured with respect to the normal of the boundary)

θ_2 angle of refraction (measured with respect to the normal of the boundary)

B.2 Fresnel equations

The following equations describe the fraction of incident radiative power that gets reflected at the boundary between air or vacuum and a medium with complex refractive index \bar{n} .




Figure B.2: Reflection and refraction of light at an interface

$$\theta = \theta_i = \theta_r \quad (\text{B.2a})$$

$$\bar{n} = n + i\kappa \quad (\text{B.2b})$$

$$z = n^2 - \kappa^2 - \sin^2 \theta \quad (\text{B.2c})$$

$$a = \sqrt{z^2 + 4n^2\kappa^2} \quad (\text{B.2d})$$

$$\rho_{\perp}(\theta) = \frac{(\sqrt{a+z} - \sqrt{2}\cos\theta)^2 + a-z}{(\sqrt{a+z} + \sqrt{2}\cos\theta)^2 + a-z} \quad (\text{B.2e})$$

$$\rho_{\parallel}(\theta) = \frac{(\sqrt{a+z} - \sqrt{2}\sin\theta\tan\theta)^2 + a-z}{(\sqrt{a+z} + \sqrt{2}\sin\theta\tan\theta)^2 + a-z} \rho_{\perp} \quad (\text{B.2f})$$

$$\rho(\theta) = \frac{\rho_{\perp} + \rho_{\parallel}}{2} \quad (\text{B.2g})$$

\bar{n} complex refractive index

n refractive index

κ extinction coefficient

$\rho_{\perp}(\theta)$ component of reflectivity for perpendicular polarized light

$\rho_{\parallel}(\theta)$ component of reflectivity for parallel polarized light

$\rho(\theta)$ reflectivity for unpolarized light

Appendix C

Device specifications

Tables C.1 summarizes the main characteristics of the infrared calibrator used in this work.

Table C.1: Fluke 4180 precision infrared calibrator

Technical specifications	Fluke 4180
Temperature range ($T_{amb} = 23\text{ }^{\circ}\text{C}$, $\varepsilon = 0.95$)	$-15\text{ }^{\circ}\text{C}$ to $120\text{ }^{\circ}\text{C}$
Nominal emissivity	0.95
Thermometer emissivity compensation	0.9 to 1.0
Display accuracy	$\pm 0.40\text{ }^{\circ}\text{C}$ at $0\text{ }^{\circ}\text{C}$
	$\pm 0.50\text{ }^{\circ}\text{C}$ at $100\text{ }^{\circ}\text{C}$
	$\pm 0.55\text{ }^{\circ}\text{C}$ at $120\text{ }^{\circ}\text{C}$
Uniformity (12.7 cm dia of center of target)	$\pm 0.10\text{ }^{\circ}\text{C}$ at $0\text{ }^{\circ}\text{C}$
	$\pm 0.25\text{ }^{\circ}\text{C}$ at $120\text{ }^{\circ}\text{C}$
Target diameter	15.2 cm

Tables C.2 summarizes the main characteristics of the depth sensor used in this work.

Table C.2: PrimeSense Carmine 1.09 depth sensor

Technical specifications	Primesense Carmine 1.09
Depth image size	640×480 pixels
Operation range	0.35 m to 1.4 m
FOV	$57.5^{\circ} \times 45^{\circ}$
Spatial x/y resolution (2-sigma values @ 0.5 m)	0.9 mm
Depth resolution (2-sigma values @ 0.5 m)	1.0 mm
Frame rate	30 Hz
Interface	USB 2.0/3.0

Tables C.3 summarizes the main characteristics of the thermographic cameras used in this work.

Table C.3: Optris PI 450 and Variocam hr head 600 thermographic cameras

Technical specifications	Optris PI 450	Variocam hr head 600
Detector	uncooled FPA (25 μm \times 25 μm)	uncooled FPA (25 μm \times 25 μm)
Image resolution	382 \times 288 pixels	640 \times 480 pixels
Spectral range	7.5 μm to 13 μm	7.5 μm to 14 μm
Temperature ranges	−20 $^{\circ}\text{C}$ to 100 $^{\circ}\text{C}$ 0 $^{\circ}\text{C}$ to 250 $^{\circ}\text{C}$ 150 $^{\circ}\text{C}$ to 900 $^{\circ}\text{C}$	−40 $^{\circ}\text{C}$ to 1200 $^{\circ}\text{C}$
Frame rate	80 Hz	50 Hz
FOV	62 $^{\circ}$ \times 49 $^{\circ}$	30 $^{\circ}$ \times 23 $^{\circ}$
focal length	8 mm	30 mm
f-number	0.8	1.0
Thermal sensitivity (NETD)	40 mK	40 mK
Accuracy	± 2 $^{\circ}\text{C}$ or ± 2 %	± 2 $^{\circ}\text{C}$ or ± 2 %
Interface	USB 2.0	Gigabit Ethernet

Appendix D

Radiometric calibration - supplemental material

This appendix contains the results of the second experiment described in section 5.1 when using the high-end thermographic camera VarioCAM hr head 600 instead of the Optris PI450. Figure D.1 shows thermal images of the infrared calibrator and the patterns taken at four different distances.

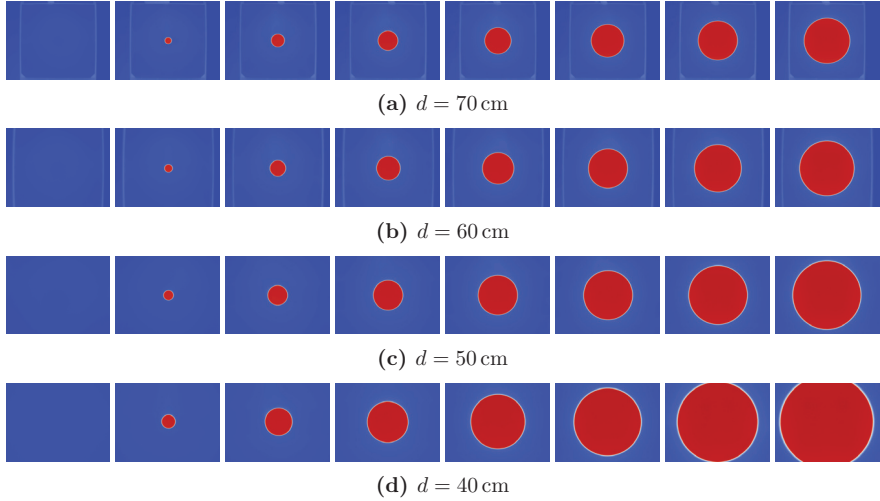


Figure D.1: Thermal images of the patterns right after placing them in front of the infrared calibrator taken at different distances d from the calibrator. $T_{\text{calibrator}} = 120^\circ\text{C}$.

Figure D.2 shows the temperature measurement indicated by the four pixel that see the center of the calibration disc in each scenario. The shape of the curve is similar to that obtained with the Optris PI450 but in this case the magnitude of the error is smaller.




Figure D.2: Temperature measured at the center of the calibrator for patterns with different hole sizes at different distances d from the calibrator. $T_{\text{calibrator}} = 120^\circ\text{C}$.

Appendix E

Geometric calibration - error margins

According to (2.22b), (2.22c) and (2.23a) the image coordinates u of a point with camera coordinates (X_c, Y_c, Z_c) can be described by (E.1).

$$u = f_u \cdot \frac{X_c}{Z_c} + c_u \quad (\text{E.1})$$

This equation can be rewritten as (E.2), where θ_u denotes the angle of incidence in the plane XZ.

$$u = f_u \cdot \tan(\theta_u) + c_u \quad (\text{E.2})$$

Let θ_{u_0} and $\theta_{u_{w-1}}$ indicate the angles of the rays that correspond to the leftmost and rightmost pixel in the image, respectively. Then, the distance between the centroids of these two pixels is given by (E.3), where w represents the image width.

$$w - 1 = f_u \cdot [\tan(\theta_{u_{w-1}}) - \tan(\theta_{u_0})] \quad (\text{E.3})$$

If the camera is rotated by an angle $\Delta\theta$ with respect to its Y axis according to the camera orientation depicted in Figure 2.9, the distance between the points where the two rays intersect the initial image plane would change by a value ε as indicated by (E.4).

$$w - 1 + \varepsilon = f_u \cdot [\tan(\theta_{u_{w-1}} + \Delta\theta) - \tan(\theta_{u_0} + \Delta\theta)] \quad (\text{E.4})$$

Solving for $\Delta\theta$ yields (E.5).

$$d \cdot \tan^2(\Delta\theta) + e \cdot \tan(\Delta\theta) + f = 0 \quad (\text{E.5})$$

Here, $a = \tan(\theta_{u_{w-1}})$, $b = \tan(\theta_{u_0})$, $c = (w - 1 + \varepsilon)/f_u$, $d = a - b - abc$, $e = ac + bc$ and $f = a - b - c$.

The solution of the quadratic equation returns two values. These values determine the range of angles by which the camera can be rotated that would cause the distance between the points where the rays θ_{u_0} and $\theta_{u_{w-1}}$ intersect the initial image plane to change by less than ε . They also determine the range of values within which c_u varies to compensate the offset caused by the rotation. The same logic applies to f_v and c_v in the YZ plane.

Example 1

Let $w = 382$ pixels and $f_u = 371.4$. If the uncertainty associated with the detection of the centroids of the circles in the image is ± 0.05 pixels (i.e. $\varepsilon = 0.1$) and the result of the geometric calibration gives $c_u = 205.5$, then:

$$\begin{aligned} a &= \frac{w - 1 - c_u}{f_u} = 4.725 \times 10^{-1} \\ b &= -\frac{c_u}{f_u} = -5.533 \times 10^{-1} \\ c &= \frac{w - 1 + \varepsilon}{f_u} = 1.026 \\ d &= a - b - abc = 1.294 \\ e &= ac + bc = -8.289 \times 10^{-2} \\ f &= a - b - c = -2.693 \times 10^{-4} \end{aligned}$$

and:

$$\begin{aligned} \Delta\theta &\in [-0.18, 3.84]^\circ \\ c_u &\in [174.1, 207.0] \text{ pixels} \end{aligned}$$

This means that any value of c_u within the range $[174.1, 207.0]$ pixels accompanied by a corresponding adjustment $\Delta\theta$ within the range $[-0.18, 3.84]^\circ$ would change the distance between the leftmost and the rightmost pixels by no more than 0.1 pixels. This change would go unnoticed due to the uncertainty in the detection of the feature coordinates.

Example 2

Let $w = 640$ pixels and $f_u = 567.7$. If the uncertainty associated with the detection of the centroids of the circles in the image is ± 0.05 pixels (i.e. $\varepsilon = 0.1$) and the result of

the geometric calibration gives $c_u = 321.3$, then:

$$a = \frac{w - 1 - c_u}{f_u} = 5.596 \times 10^{-1}$$

$$b = -\frac{c_u}{f_u} = -5.660 \times 10^{-1}$$

$$c = \frac{w - 1 + \varepsilon}{f_u} = 1.126$$

$$d = a - b - abc = 1.482$$

$$e = ac + bc = -7.139 \times 10^{-3}$$

$$f = a - b - c = -1.761 \times 10^{-4}$$

and:

$$\Delta\theta \in [-0.50, 0.78]^\circ$$

$$c_u \in [311.2, 327.9] \text{ pixels}$$

In this case, both the range of $\Delta\theta$ and the range of c_u are smaller than in the previous example due to the higher image resolution of the camera.

Appendix F

Data fusion errors

F.1 Registration error due to timestamp errors

The interpolation factor α and the interpolated position $\mathbf{p}_{\text{depth}}(t[k])$ described by (6.1a) and (6.1b) in Section 6.2 are shown here again for convenience.

$$\alpha = \frac{t[k] - t[i - 1]}{t[i] - t[i - 1]} \quad (\text{F.1a})$$

$$\mathbf{p}_{\text{depth}}(t[k]) = \alpha \mathbf{p}_{\text{depth}}(t[i]) + (1 - \alpha) \mathbf{p}_{\text{depth}}(t[i - 1]) \quad (\text{F.1b})$$

The variables $t[i]$ and $t[i - 1]$ indicate the time when depth images i and $i - 1$ are captured. The variable $t[k]$ indicates the time when thermal image k is captured, such that $t[i - 1] \leq t[k] \leq t[i]$. The vector $\mathbf{p}_{\text{depth}}(t)$ denotes the position of the depth camera with respect to the 3D model at time t .

Let $\Delta t_{\text{D}}[i]$ represent the time difference between the timestamp assigned to the depth image frame i and the corresponding value in the linear regression calculated with all timestamps. Similarly, let $\Delta t_{\text{T}}[k]$ represent the time difference between the timestamp assigned to the thermal image frame k and the corresponding value in the linear regression calculated with all timestamps. The linear regressions represent the ideal timestamps that should be assigned to the corresponding depth and thermal image frames given that the camera frame rates are constant. Assuming that the average offset between the depth and thermal timestamps is null (i.e. it has already been accounted for in the assigned timestamps), the interpolation factor including the time differences $\Delta t_{\text{D}}[i]$ and

$\Delta t_T[k]$ can be rewritten as:

$$\alpha' = \frac{(t[k] + \Delta t_T[k]) - (t[i-1] + \Delta t_D[i-1])}{(t[i] + \Delta t_D[i]) - (t[i-1] + \Delta t_D[i-1])} \quad (\text{F.2a})$$

$$= \frac{(t[k] - t[i-1]) + (\Delta t_T[k] - \Delta t_D[i-1])}{(t[i] - t[i-1]) + (\Delta t_D[i] - \Delta t_D[i-1])} \quad (\text{F.2b})$$

Typical values of $\Delta t_D[i]$ and $\Delta t_T[k]$ are presented in section 5.3 and are shown again in Figure F.1 for convenience.




Figure F.1: Variability of image timestamps after filtering

Since $t[i] - t[i-1] \gg \Delta t_D[i] - \Delta t_D[i-1]$, the interpolation factor can be rewritten as:

$$\alpha' \approx \frac{(t[k] - t[i-1]) + (\Delta t_T[k] - \Delta t_D[i-1])}{t[i] - t[i-1]} \quad (\text{F.3a})$$

$$\approx \frac{t[k] - t[i-1]}{t[i] - t[i-1]} + \frac{\Delta t_T[k] - \Delta t_D[i-1]}{t[i] - t[i-1]} \quad (\text{F.3b})$$

$$\approx \alpha + \frac{\Delta t_T[k] - \Delta t_D[i-1]}{t[i] - t[i-1]} \quad (\text{F.3c})$$

At a frame rate of 30 fps, $t[i] - t[i-1] = 33.3$ ms. According to Figure F.1, in the worst case scenario $\Delta t_T[k] - \Delta t_D[i-1] \approx 6$ ms. Therefore, in the worst case scenario $\alpha' \approx \alpha + 0.18$. Since $0 \leq \alpha \leq 1$, the maximum interpolation factor error according to the data is 18%.

In the case of pure translation, if the linear scan velocity is 0.3 m/s, the error in the estimation of the position of the depth camera is:

$$\Delta \mathbf{p}_{\text{depth}}(t[k]) = \mathbf{p}'_{\text{depth}}(t[k]) - \mathbf{p}_{\text{depth}}(t[k]) \quad (\text{F.4a})$$

$$\begin{aligned} &= \alpha' \mathbf{p}_{\text{depth}}(t[i]) + (1 - \alpha') \mathbf{p}_{\text{depth}}(t[i - 1]) \\ &\quad - \alpha \mathbf{p}_{\text{depth}}(t[i]) - (1 - \alpha) \mathbf{p}_{\text{depth}}(t[i - 1]) \end{aligned} \quad (\text{F.4b})$$

$$= 0.18 \cdot (\mathbf{p}_{\text{depth}}(t[i]) - \mathbf{p}_{\text{depth}}(t[i - 1])) \quad (\text{F.4c})$$

$$= 0.18 \cdot 10 \text{ mm} = 1.8 \text{ mm} \quad (\text{F.4d})$$

This is the same error in the estimation of the position of the thermal camera, since the two sensors are rigidly attached.

In the case of pure rotation centered on the optical center of the thermographic camera, if the angular scan velocity is $18^\circ/\text{s}$, the registration error of a point 1 m away from the thermal camera is:

$$18^\circ \text{ s}^{-1} \cdot \frac{1}{30} \text{ s} \cdot \frac{\pi}{180} \cdot 1 \text{ m} \cdot 0.18 = 1.9 \text{ mm} \quad (\text{F.5})$$

In practice, both types of error occur simultaneously. If the errors are aligned the maximum combined error due to inaccuracies in the timestamps is 3.7 mm.

F.2 Registration error due to depth camera pose errors

Let $\Delta \mathbf{p}_{\text{depth}}(t[i])$ and $\Delta \mathbf{p}_{\text{depth}}(t[i - 1])$ be the errors in the estimation of the position of the depth camera for frames i and $i - 1$, respectively. The interpolated position $\mathbf{p}_{\text{depth}}(t[k])$ including the errors $\Delta \mathbf{p}_{\text{depth}}(t[i])$ and $\Delta \mathbf{p}_{\text{depth}}(t[i - 1])$ is then given by:

$$\begin{aligned} \mathbf{p}'_{\text{depth}}(t[k]) &= \alpha(\mathbf{p}_{\text{depth}}(t[i]) + \Delta \mathbf{p}_{\text{depth}}(t[i])) \\ &\quad + (1 - \alpha)(\mathbf{p}_{\text{depth}}(t[i - 1]) + \Delta \mathbf{p}_{\text{depth}}(t[i - 1])) \end{aligned} \quad (\text{F.6a})$$

Hence, the error in the calculation of the interpolated position is:

$$\Delta \mathbf{p}_{\text{depth}}(t[k]) = \mathbf{p}'_{\text{depth}}(t[k]) - \mathbf{p}_{\text{depth}}(t[k]) \quad (\text{F.7a})$$

$$= \alpha \Delta \mathbf{p}_{\text{depth}}(t[i]) + (1 - \alpha) \Delta \mathbf{p}_{\text{depth}}(t[i - 1]) \quad (\text{F.7b})$$

In other words, the error of the interpolated position is equal to the interpolation of the individual errors. Thus, without loss of generality, the maximum error $\Delta \mathbf{p}_{\text{depth}}(t[k])$ is equal to maximum value that $\Delta \mathbf{p}_{\text{depth}}(t[i])$ can take. The value of $\Delta \mathbf{p}_{\text{depth}}(t[i])$ depends

on the accuracy of the depth sensor and the ICP algorithm. The latter in turn depends on the number of vertices used in the calculations, the geometry of the point clouds and the magnitude of the relative pose between them. If the error of the ICP algorithm is assumed to be less than 10% of the displacement between consecutive depth frames, then, for a linear scan velocity of 0.3 m/s and a frame rate of 30 fps, the error in the estimation of the position of the depth camera is:

$$\Delta p_{\text{depth}}(t[k]) = 0.10 \cdot (p_{\text{depth}}(t[i]) - p_{\text{depth}}(t[i-1])) = 1 \text{ mm} \quad (\text{F.8})$$

This is the same error for the position of the thermographic camera, since the two sensors are rigidly attached.

A similar analysis can be carried out for the case of pure rotation. In Figure F.2, O indicates the optical center of the thermographic camera, A represents a point seen by some detector at time $t[i-1]$ and B represents another point seen by the same detector at time $t[i]$ after the camera is rotated. The arc \widehat{AB} represents points seen by the same detector at time $t[k]$ for different values of α . The spherical caps centered at A and B denote the uncertainty in the estimation of the correct orientation of the camera with respect to the 3D model. Arcs drawn between points on the spherical caps centered at A and B represent the possible rotational trajectories that the camera can follow in the time interval $t[i-1] \rightarrow t[i]$.




Figure F.2: Rotational uncertainty

From the figure it is clear that the trajectory farthest away from the arc \widehat{AB} must be tangential to the circles that define the spherical caps and that the maximum error occurs at the beginning or at the end of the trajectory. If the error of the ICP algorithm is assumed to be less than 10% of the rotation between consecutive depth frames, then, for an angular scan velocity of $18^\circ/\text{s}$, a frame rate of 30 fps and a distance to the object of 1 m, the maximum registration error due to rotation would be equal to:

$$18^\circ \text{ s}^{-1} \cdot \frac{1}{30} \text{ s} \cdot \frac{\pi}{180} \cdot 1 \text{ m} \cdot 0.10 = 1 \text{ mm} \quad (\text{F.9})$$

Publications

Peer reviewed publications

- **A. Ordonez Müller**, A. Kroll, "Generating High Fidelity 3D Thermograms with a Handheld Real-time Thermal Imaging System", *IEEE Sensors Journal*, vol. 17, no. 3, p. 774-783, February 2017.
- **A. Ordonez Müller**, A. Kroll, "On the Temperature Assignment Problem and the Use of Confidence Textures in the Creation of 3D Thermograms", in Proceedings of the 9th International Conference on Sensing Technology (ICST 2015), Auckland, New Zealand, December 2015, pp. 223-228

Other publications

- **A. Ordonez Müller**, A. Kroll, "On the use of Cooperative Autonomous Mobile Robots and Optical Remote Sensing in Inspection Robotics", *Automation*, Baden-Baden, Germany, July 2014.
- **A. Ordonez Müller**, A. Kroll, "Effects of beam divergence in hand-held TD-LAS sensors on long distance gas concentration measurements", in Proceedings of the International Workshop on Advanced Infrared Technology and Applications (AITA), vol. 12, Turin, Italy , September 2013, pp. 9-13
- **A. Ordonez Müller**, A. Kroll, "On Range Extension of Tunable Diode Laser Absorption Spectroscopy (TDLAS) Based Devices in Remote Gas Sensing Applications", in Proceedings of the International Symposium on Olfaction and Electronic Nose (ISOEN), EXCO, Daegu, Korea , July 2013
- J. Hegenberg, R. Herrmann, D. Ziegner, L. Schmidt, T. Günther, **A. Ordoñez Müller**, A. Kroll, T. Barz, D. Schulz, "Forschungsprojekt Robotair: Praxisaugliches Boden-Luft-Servicerobotersystem für die Inspektion industrieller Druckluftversorgung und die Verbesserung der Arbeitsumgebungsfaktoren", *Technische Sicherheit*, pp. 16–22, Mai 2015

Bibliography

- [1] L. Acampora, F. D. Filippis, A. Martucci, and L. Sorgi, “3D reconstruction of thermal images,” in *Proceedings of the 26th Aerospace Testing Seminar*, Los Angeles, California, USA, March 2011, pp. 263–276.
- [2] M. A. Akhloufi and B. Verney, “Multimodal registration and fusion for 3D thermal imaging,” *Mathematical Problems in Engineering*, 2015.
- [3] P. Aksenov, I. Clark, D. Grant, A. Inman, L. Vartikovski, and J.-C. Nebel, “3D thermography for quantification of heat generation resulting from inflammation,” in *Proceedings of the 8th 3D Modelling Symposium*, Paris, France, 2003.
- [4] M. I. Alba, L. Barazzetti, M. Scaioni, E. Rosina, and M. Previtali, “Mapping Infrared Data on Terrestrial Laser Scanning 3D Models of Buildings,” *Remote Sensing*, vol. 3, no. 9, pp. 1847–1870, August 2011.
- [5] A. Appel, “Some techniques for shading machine renderings of solids,” in *Proceedings of the AFIPS Spring Joint Computing Conference*, Atlantic City, New Jersey, USA, April-May 1968, pp. 37–45.
- [6] Atomic and Molecular Physics Division, Harvard-Smithsonian Center for Astrophysics, “The HITRAN database,” <http://hitran.org/>, Accessed on May 04, 2017.
- [7] S. Barone, A. Paoli, and A. V. Razionale, “A biomedical application combining visible and thermal 3D imaging,” in *Proceedings of the XVIII Congreso Internacional de Ingeniería Gráfica (INGEGRAF)*, Barcelona, Spain, May-June 2006, pp. 1–9.
- [8] P. J. Besl and N. D. McKay, “A method for registration of 3D shapes,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 239–256, February 1992.
- [9] D. Borrmann, A. Nüchter, M. Đakulović, I. Maurović, I. Petrović, D. Osmanković, and J. Velagić, “The Project ThermalMapper - Thermal 3D Mapping of Indoor Environments for Saving Energy,” in *Proceedings of the 10th International IFAC Symposium on Robot Control (SYROCO)*, Dubrovnik, Croatia, September 2012, pp. 31–38.

- [10] J. Y. Bouguet, “MATLAB calibration Toolbox,” http://www.vision.caltech.edu/bouguetj/calib_doc/, Accessed on Apr 13, 2017.
- [11] H. Budzier and G. Gerlach, *Thermal Infrared Sensors*. Wiley, 2011.
- [12] H. Budzier and G. Gerlach, “Calibration of uncooled thermal infrared cameras,” *Journal of Sensors and Sensor Systems*, vol. 4, no. 1, pp. 187–197, 2015.
- [13] M. Cabrelles, S. Galcerá, S. Navarro, J. L. Lerma, T. Akasheh, and N. Haddad, “Integration of 3D laser scanning photogrammetry and thermography to record architectural monuments,” in *Proceedings of the 22nd Symposium of the International Committee for Documentation of Cultural Heritage (CIPA)*, Kyoto, Japan, October 2009.
- [14] G. Cardone, A. Ianiro, G. dello Ioio, and A. Passaro, “Temperature maps measurements on 3D surfaces with infrared thermography,” *Experiments in Fluids*, vol. 52, no. 2, pp. 375–385, February 2012.
- [15] G. Champelboux, S. Lavalée, R. Szeliski, and L. Brunie, “From accurate range imaging sensor calibration to accurate model-based 3D object localization,” in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, Champaign, Illinois, USA, June 1992, pp. 83–89.
- [16] Y. Chen and G. Medioni, “Object modeling by registration of multiple range images,” in *Proceedings of the IEEE International Conference on Robotics and Automation*, Sacramento, California, USA, April 1991, pp. 2724–2729.
- [17] P. Cignoni, M. Callieri, M. Corsini, M. Dellepiane, F. Ganovelli, and G. Ranzuglia, “MeshLab: an open-source mesh processing tool,” in *Proceedings of the Eurographics Italian Chapter Conference*, V. Scarano, R. D. Chiara, and U. Erra, Eds. The Eurographics Association, 2008.
- [18] Cignoni, Paolo and Callieri, Marco and Corsini, Massimiliano and Dellepiane, Matteo and Ganovelli, Fabio and Ranzuglia, Guido, “MeshLab,” <http://www.meshlab.net/>, Accessed on May 04, 2017.
- [19] G. G. Demisse, D. Borrmann, and A. Nüchter, “Interpreting thermal 3D models of indoor environments for energy efficiency,” *Journal of Intelligent & Robotic Systems*, vol. 77, no. 1, pp. 55–72, January 2015.
- [20] Fluke Corporation, “Professional Series Infrared Cameras,” <http://en-us.fluke.com/products/infrared-cameras/professional-series-infrared-cameras.html>, Accessed on May 04, 2017.

- [21] R. Gade and T. B. Moeslund, "Thermal cameras and applications: a survey," *Machine Vision and Applications*, vol. 25, no. 1, pp. 245–262, January 2014.
- [22] D. Girardeau-Montaut, "CloudCompare," <http://www.danielgm.net/cc/>, Accessed on May 04, 2017.
- [23] D. González-Aguilera, P. Rodríguez-Gonzálvez, J. Armesto, and S. Lagüela, "Novel approach to 3D thermography and energy efficiency evaluation," *Energy and Buildings*, vol. 54, pp. 436–443, November 2012.
- [24] H. González-Jorge, B. Riveiro, E. Vazquez-Fernandez, J. Martínez-Sánchez, and P. Arias, "Metrological evaluation of Microsoft Kinect and Asus Xtion sensors," *Measurement*, vol. 46, no. 6, pp. 1800–1806, 2013.
- [25] I. Grubišić, L. Gjenero, T. Lipić, I. Sović, and T. Skala, "Active 3D scanning based 3D thermography system and medical applications," in *Proceedings of the 34th International Convention MIPRO*, Opatija, Croatia, May 2011, pp. 269–273.
- [26] I. Grubišić, L. Gjenero, T. Lipić, I. Sović, and T. Skala, "Medical 3D thermography system," *Periodicum Biologorum*, vol. 113, pp. 401–406, November 2011.
- [27] K. Hajebi and J. S. Zelek, "Structure from infrared stereo images," in *Proceedings of the Canadian Conference on Computer and Robot Vision*, May 2008, pp. 105–112.
- [28] Y. Ham and M. Golparvar-Fard, "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, August 2013.
- [29] E. Hecht, *Optics*, 5th ed. Pearson Education Limited, 2016.
- [30] V. Hilsenstein, "Surface reconstruction of water waves using thermographic stereo imaging," in *Proceedings of the Conference on Image and Vision Computing New Zealand (IVCNZ)*, Dunedin, New Zealand, November 2005, pp. 102–107.
- [31] L. Hoegner and U. Stilla, "Thermal leakage detection on building façades using infrared textures generated by mobile mapping," in *Proceedings of the Urban Remote Sensing Joint Event*, Shanghai, China, May 2009, pp. 1–6.
- [32] J. R. Howell, R. Siegel, and M. P. Mengüç, *Thermal Radiation Heat Transfer*, 5th ed. CRC Press, 2011.
- [33] J. D. Hunter, "Matplotlib: A 2D graphics environment," *Computing In Science & Engineering*, vol. 9, no. 3, pp. 90–95, 2007.

- [34] D. Iwaszczuk, L. Hoegner, and U. Stilla, “Matching of 3D building models with IR images for texture extraction,” in *Proceedings of the Joint Urban Remote Sensing Event*, Munich, Germany, April 2011, pp. 25–28.
- [35] S. Izadi, R. Newcombe, D. Kim, O. Hilliges, D. Molyneaux, S. Hodges, P. Kohli, A. Davison, and A. Fitzgibbon, “Kinectfusion: Real-time dynamic 3D surface reconstruction and interaction,” in *Proceedings of the ACM SIGGRAPH Talks*, Vancouver, British Columbia, Canada, August 2011.
- [36] X. Ju, J.-C. Nebel, and J. P. Siebert, “3D thermography imaging standardization technique for inflammation diagnosis,” in *Proceedings of SPIE - Infrared Components and Their Applications*, vol. 5640, Beijing, China, January 2005, pp. 266–273.
- [37] J. Kammerl, “OpenNI 2 Camera Driver,” http://wiki.ros.org/openni2_camera, Accessed on April 14, 2017.
- [38] K. Khoshelham and S. O. Elberink, “Accuracy and resolution of Kinect depth data for indoor mapping applications,” *Sensors*, vol. 12, no. 9, pp. 1437–1454, February 2012.
- [39] M. Kriss, *Handbook of Digital Imaging*. Wiley, 2015, vol. 3.
- [40] A. Kylili, P. A. Fokaides, P. Christou, and S. A. Kalogirou, “Infrared thermography (IRT) applications for building diagnostics: a review,” *Applied Energy*, vol. 134, pp. 531–549, December 2014.
- [41] S. Lagüela, L. Díaz-Vilariño, J. Armesto, and P. Arias, “Thermographic 3D models as the foundation for building information models,” in *Proceedings of the 11th International Conference on Quantitative Infrared Thermography (QIRT)*, Naples, Italy, June 2012.
- [42] S. Lagüela, H. González-Jorge, J. Armesto, and P. Arias, “Calibration and verification of thermographic cameras for geometric measurements,” *Infrared Physics & Technology*, vol. 54, no. 2, pp. 92–99, 2011.
- [43] S. Lagüela, H. González-Jorge, J. Armesto, and J. Herráez, “High performance grid for the metric calibration of thermographic cameras,” *Measurement Science and Technology*, vol. 23, no. 1, 2012.
- [44] S. Lagüela, J. Martínez, J. Armesto, and P. Arias, “Energy efficiency studies through 3D laser scanning and thermographic technologies,” *Energy and Buildings*, vol. 43, no. 6, pp. 1216–1221, June 2011.

- [45] B. B. Lahiri, S. Bagavathiappan, T. Jayakumar, and J. Philip, “Medical applications of infrared thermography: A review,” *Infrared Physics & Technology*, vol. 55, no. 4, pp. 221–235, 2012.
- [46] B. Lane and E. P. Whittenton, “Calibration and measurement procedures for a high magnification thermal camera,” National Institute of Standards and Technology, Tech. Rep., December 2015.
- [47] F. Liebmann and T. Kolat, “Standards for radiation thermometry,” in *Proceedings of the National Conference of Standards Laboratories (NCSL) International Workshop and Symposium*, Sacramento, California, USA, July-August 2012.
- [48] T. Luhmann, J. Ohm, J. Piechel, and T. Roelfs, “Geometric calibration of thermographic cameras,” *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 38, no. 5, pp. 411–416, June 2010.
- [49] G. Machin, R. Simpson, and M. Broussely, “Calibration and validation of thermal imagers,” in *Proceedings of the 9th International Conference on Quantitative Infrared Thermography (QIRT)*, Krakow, Poland, July 2008.
- [50] F. L. Markley, Y. Cheng, J. L. Crassidis, and Y. Oshman, “Averaging Quaternions,” *Journal of Guidance, Control, and Dynamics*, vol. 30, no. 4, pp. 1193–1197, July-August 2007.
- [51] S. May, “Optris Camera Driver,” http://wiki.ros.org/optris_drivers, Accessed on August 02, 2017.
- [52] G. Medioni and S. B. Kang, *Emerging Topics in Computer Vision*. Prentice Hall Computer, 2005.
- [53] S. Meister, S. Izadi, P. Kohli, M. Hämmerle, C. Rother, and D. Kondermann, “When can we use kinectfusion for ground truth acquisition?” in *Proceedings of the Workshop on Color-Depth Camera Fusion in Robotics, IEEE International Conference on Intelligent Robots and Systems*, 2012.
- [54] Mikhail N. Polyanskiy, “Refractive index database,” <https://refractiveindex.info>, Accessed on June 08, 2017.
- [55] D. L. Mills, *Computer Network Time Synchronization*, 2nd ed. CRC Press, 2011.
- [56] P. Moghadamand and S. Vidas, “Heatwave: The next generation of thermography devices,” in *Proceedings of SPIE - Thermosense: Thermal Infrared Applications XXXVI*, vol. 9105, Baltimore, Maryland, USA, May 2014.

- [57] P. J. Mohr, D. B. Newell, and B. N. Taylor, “Codata recommended values of the fundamental physical constants: 2014,” *Reviews of Modern Physics*, vol. 58, no. 3, July-September 2016.
- [58] Š. Mráček, J. Váňa, R. Dvořák, M. Draňanský, and S. Yanushkevich, *New Trends and Developments in Biometrics*. InTech, 2012, ch. 2.
- [59] K. Nagatani, K. Otake, and K. Yoshida, *Three-Dimensional Thermography Mapping for Mobile Rescue Robots*. Springer, 2014, pp. 49–63.
- [60] Y.-M. H. Ng and R. Du, “Acquisition of 3D surface temperature distribution of a car body,” in *Proceedings of the IEEE International Conference on Information Acquisition (ICIA)*, Hong Kong and Macau, China, June-July 2005, pp. 16–20.
- [61] C. V. Nguyen, S. Izadi, and D. Lovell, “Modeling Kinect sensor noise for improved 3D reconstruction and tracking,” in *Proceedings of the 2nd Joint Conference 3DIM/3DPVT: 3D Imaging, Modeling, Processing, Visualization and Transmission*, 2012, pp. 524–530.
- [62] OpenCV, “Camera calibration and 3D reconstruction,” http://docs.opencv.org/2.4/modules/calib3d/doc/camera_calibration_and_3d_reconstruction.html, Accessed on April 13, 2017.
- [63] Optris GmbH, “Optris Calculator,” <http://www.optris.com/optics-calculator>, Accessed on May 06, 2016.
- [64] Optris GmbH, “Optris PI 200,” <http://www.optris.de/infrarotkamera-pi200>, Accessed on May 04, 2017.
- [65] J. J. Orteu, Y. Rotrou, T. Sentenac, and L. Robert, “An innovative method for 3D shape, strain and temperature full-field measurement using a single type of camera: Principle and preliminary results,” *Experimental Mechanics*, vol. 48, no. 2, pp. 163–179, April 2008.
- [66] D. Pagliari, F. Menna, R. Roncella, F. Remondino, and L. Pinto, “Kinect Fusion improvement using depth camera calibration,” *The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 40, no. 5, pp. 479–485, June 2014.
- [67] A. Pelagotti, A. D. Mastio, F. Ucheddu, and F. Remondino, “Automated multispectral texture mapping of 3D models,” in *Proceedings of the 17th European Signal Processing Conference (EUSIPCO)*, Glasgow, Scotland, August 2009, pp. 1215–1219.

- [68] S. Prakash, P. Y. Lee, and T. Caelli, “3D mapping of surface temperature using thermal stereo,” in *Proceedings of the 9th International Conference on Control, Automation, Robotics and Vision (ICARCV)*, Singapore, December 2006, pp. 1–4.
- [69] H. Preston-Thomas, “The international temperature scale of 1990 (ITS-90),” *Metrologia*, vol. 27, no. 1, pp. 3–10, 1990.
- [70] M. Quigley, K. Conley, B. P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, “ROS: an open-source robot operating system,” in *Proceedings of the International Conference on Robotics and Automation (ICRA) - Workshop on Open Source Software*, Kobe, Japan, May 2009.
- [71] J. Rangel González, S. Soldan, and A. Kroll, “3D thermal imaging: Fusion of thermography and depth cameras,” in *Proceedings of the 12th International Conference on Quantitative Infrared Thermography (QIRT)*, Bordeaux, France, July 2014.
- [72] S. F. Ray, *Applied Photographic Optics*, 3rd ed. Focal Press, 2002.
- [73] B. Ribbens, J. Peeters, B. Bogaerts, G. Steenackers, S. Sels, R. Penne, and G. V. Barel, “4D active and passive thermography measurement system using a KUKA KR16 robot and time-of-flight imaging,” in *Proceedings of the 13th International Conference on Quantitative Infrared Thermography (QIRT)*, Gdańsk, Poland, July 2016, pp. 670–677.
- [74] A. Rogalski, “Recent progress in infrared detector technologies,” *Infrared Physics & Technology*, vol. 54, no. 3, pp. 136–154, 2011.
- [75] D. Rzeszutarski and B. Więcek, “Calibration for 3D reconstruction of thermal images,” in *Proceedings of the 9th International Conference on Quantitative Infrared Thermography (QIRT)*, Krakow, Poland, July 2008.
- [76] W. Satzger, G. Zenzinger, and V. Carl, “Quantitative 3D Thermography,” in *Proceedings of the 8th Conference on Quantitative Infrared Thermography (QIRT)*, Padova, Italy, June 2006.
- [77] P. Sawicki, R. Stein, and B. Więcek, “Directional emissivity correction by photogrammetric 3D object reconstruction,” in *Proceedings of the 4th Conference on Quantitative Infrared Thermography (QIRT’98)*, Łódź, Poland, September 1998, pp. 327–332.
- [78] C. Schoenauer, E. Vonach, G. Gerstweiler, and H. Kaufmann, “3D building reconstruction and thermal mapping in fire brigade operations,” in *Proceedings of the 4th Augmented Human International Conference*, March 2013, pp. 548–555.

- [79] R. C. Simpson, H. C. McEvoy, G. Machin, K. Howell, M. Naeem, P. Plassmann, F. Ring, P. Campbell, C. Song, J. Tavener, and I. Ridley, "In-field-of-view thermal image calibration system for medical thermography applications," *International Journal of Thermophysics*, vol. 29, no. 3, pp. 1123–1130, June 2008.
- [80] K. Skala, T. Lipić, I. Sović, L. Gjenero, and I. Grubišić, "4D thermal imaging system for medical applications," *Periodicum Biologorum*, vol. 113, pp. 407–416, November 2011.
- [81] The Point Cloud Library, "Open source implementation of the Kinect Fusion algorithm," <http://pointclouds.org/>, Accessed on May 06, 2016.
- [82] S. Vidas, R. Lakemond, S. Denman, C. Fookes, S. Sridharan, and T. Wark, "A mask-based approach for the geometric calibration of thermal-infrared cameras," *IEEE Transactions on Instrumentation and Measurement*, vol. 61, no. 6, pp. 1625–1635, June 2012.
- [83] S. Vidas, P. Moghadam, and S. Sridharan, "Real-time mobile 3D temperature mapping," *IEEE Sensors Journal*, vol. 15, no. 2, p. 1145, February 2015.
- [84] S. Vidas and P. Moghadaman, "Heatwave: A handheld 3D thermography system for energy auditing," *Energy and Buildings*, vol. 66, pp. 445–460, November 2013.
- [85] S. Vidas, P. Moghadamand, 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 (ICRA)*, Karlsruhe, Germany, May 2013, pp. 2311–2318.
- [86] M. Vollmer and K.-P. Möllmann, *Infrared Thermal Imaging*. Wiley-VCH, 2010.
- [87] J. Wardlaw, M. Gryka, F. Wanner, G. Brostow, and J. Kautz, "A New Approach to Thermal Imaging Visualisation – Thermal Imaging in 3D," University College London, Tech. Rep., July 2010.
- [88] T. Whelan, M. Kaess, M. F. Fallon, H. Johannsson, J. J. Leonard, and J. B. McDonald, "Kintinuous: Spatially extended Kinect Fusion," in *Proceedings of the RSS Workshop on RGB-D: Advanced Reasoning with Depth Cameras*, Sydney, Australia, July 2012.
- [89] T. Whelan, M. Kaess, J. Leonard, and J. B. McDonald, "Deformation-based loop closure for large scale dense RGB-D SLAM," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Tokyo, Japan, November 2013, pp. 548–555.

- [90] T. Whelan, M. Kaess, H. Johannsson, M. Fallon, J. J. Leonard, and J. McDonald, “Real-time large-scale dense RGB-D SLAM with volumetric fusion,” *International Journal of Robotics Research*, vol. 34, no. 4-5, pp. 598–626, April 2015.
- [91] T. L. Williams, *Thermal Imaging Cameras - Characteristics and Performance*. CRC Press, 2009.
- [92] R. Yang and Y. Chen, “Design of a 3D infrared imaging system using structured light,” *IEEE Transactions on Instrumentation and Measurement*, vol. 60, pp. 608–617, February 2011.
- [93] R. Yang, Y. Chen, and X. Wu, “Geometric calibration of IR camera using trinocular vision,” *Journal of Lightwave Technology*, vol. 29, no. 24, pp. 3797–3803, December 2011.
- [94] S. Zhang, *Handbook of 3D Machine Vision: Optical Metrology and Imaging*. CRC Press, Taylor & Francis Group, 2013.
- [95] Z. Zhang, “A flexible new technique for camera calibration,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 11, pp. 1330–1334, November 2000.

Schriftenreihe Mess- und Regelungstechnik der Universität Kassel

Herausgegeben von / Edited by

Univ.-Prof. Dr.-Ing. Andreas Kroll, Universität Kassel

- Band 1:** Klassifikationsgestützte on-line Adaption eines robusten beobachter-basierten Fehlerdiagnoseansatzes für nichtlineare Systeme, Kassel 2011
Patrick Gerland
- Band 2:** Zur Identifikation mechatronischer Stellglieder mit Reibung bei Kraftfahrzeugen, Kassel 2012
Zhenxing Ren
- Band 3:** Sensordatenfusionsansätze in der Thermografie zur Verbesserung der Messergebnisse, Kassel 2014
Samuel Soldan
- Band 4:** Multi-Robot Task Allocation for Inspection Problems with Cooperative Tasks Using Hybrid Genetic Algorithms, Kassel 2014
Chun Liu
- Band 5:** Gasleckortungsmethode für autonome mobile Inspektionsroboter mit optischer Gasfernmesstechnik in industrieller Umgebung, Kassel 2015
Gero Bonow
- Band 6:** Dynamische Analyse großer verkoppelter Systeme mit Methoden der Komplexen Netzwerke am Beispiel des Inverse-Response-Verhaltens, Kassel 2015
Andreas Geiger

Infrared thermography enables the non-contact measurement of an object's surface temperature and presents the results in form of thermal images. The analysis of these images provides valuable information about an object's thermal state. However, the fidelity of the thermal images strongly depends on the pose of the thermographic camera with respect to the surface. 3D thermography offers the possibility to overcome this and other limitations that affect conventional 2D thermography but most 3D thermographic systems developed so far generate 3D thermograms from a single perspective or from few noncontiguous points of view and do not operate in real time. As a result, the 3D thermograms they generate do not offer much advantage over conventional thermal images. However, recent technological advances have unlocked the possibility of implementing affordable handheld 3D thermal imaging systems that can be easily maneuvered around an object and that can generate high-fidelity 3D thermograms in real time. This thesis explores various aspects involved in the real-time generation of high-fidelity 3D thermograms at close range using a handheld 3D thermal imaging system, presents the results of scanning an operating industrial furnace and discusses the problems associated with the generation of 3D thermograms of large objects with complex geometries.

ISBN 978-3-7376-0624-0



9 783737 606240 >