



Institución Universitaria

Development of a system to improve resolution and accuracy in temperature measurement in thermal images through machine learning techniques and noise reduction

Paula Alejandra Ortiz Santana

Instituto Tecnológico Metropolitano
Faculty of Engineering
Medellín, Colombia
2025

Development of a system to improve resolution and accuracy in temperature measurement in thermal images through machine learning techniques and noise reduction

Paula Alejandra Ortiz Santana

Thesis submitted as a partial requirement to obtain the degree of:

Master in Automation and Industrial Control

Thesis Advisors:

Ph.D. Jorge Alexis Herrera Ramírez

Ph.D. Juan Sebastian Botero Valencia

Research Lines:

Computer Vision and Photonics

Control Systems and Robotics

Research Group:

Automation, Electronics and Computational Sciences

Control Systems and Robotics

Instituto Tecnológico Metropolitano

Faculty of Engineering

Medellín, Colombia

2025

Acknowledgments

First, I want to thank my parents, Henry Ortiz and Malluly Santana, for their unwavering support, for believing in me throughout this process, and for instilling in me the values of perseverance and discipline. I am also deeply grateful to my partner, Juan David, for his patience, unconditional support, and for being a constant source of stability along this journey. And I sincerely extend this gratitude to my entire family for their unconditional support, even despite the distance, and especially to my cousin Mariana for her optimism, joy, and companionship during this time.

I sincerely thank my advisors, Professors Jorge Herrera and Juan Sebastián Botero, for their academic guidance, rigor, and commitment, which were essential to the development of this work. I also thank Professor Erick Reyes for his continued support, technical contributions, and constant availability throughout the project.

I dedicate this work, in a very special way, to my uncle, Milton César Téllez Ortiz. Much of my personal and professional development is due to him. His guidance, his example, and his love of knowledge profoundly shaped my path. Although he is no longer with us, his presence remains in every step I take and every decision I make. This work exists thanks to what he instilled in me, and therefore it is for him and thanks to him.

Finally, I would like to acknowledge and sincerely thank the financial support that made this research possible. This work was funded through the research and development project “Microwave applicator for the treatment of melanoma through hyperthermia”, supported by the Call for Strengthening Science, Technology, and Innovation in Higher Education Institutions (Call 890-2020 of the Ministry of Science, Technology, and Innovation of Colombia), project code 2021-1086.

Resumen

Las cámaras térmicas se utilizan en aplicaciones industriales, médicas y científicas, gracias a su capacidad para registrar distribuciones espaciales de temperatura a partir de radiación infrarroja. Sin embargo, a pesar de su amplia aplicabilidad, estos sensores suelen tener un costo elevado y las referencias más económicas presentan limitaciones en cuanto a resolución espacial y precisión en las mediciones, lo que puede dificultar la detección de bordes y la identificación precisa de zonas de interés. En la literatura, los enfoques que han intentado mejorar estas limitaciones, utilizando técnicas como la interpolación, el filtrado o los modelos de aprendizaje automático, han mostrado avances. Sin embargo, muchos de estos métodos dependen de bases de datos no públicas o carecen de protocolos de adquisición, omitiendo las particularidades de los sensores, dejando un vacío en la confiabilidad práctica de los datos.

Este documento propone que la combinación de adquisición multimodal controlada, técnicas de reducción de ruido y modelos de aprendizaje automático entrenados con datos propios puede mejorar tanto la resolución espacial como la precisión de las mediciones de temperatura. Para ello, se desarrolló un sistema que integra cámaras RGB y térmicas, un protocolo de adquisición en condiciones controladas y una base de datos etiquetada. Sobre esta base, se implementaron y evaluaron diferentes arquitecturas, como redes generativas adversarias, VAE y U-Net, evaluadas mediante métricas cuantitativas como PSNR y error medio de temperatura.

Los resultados mostraron incrementos en el PSNR desde 17.43 dB hasta 20.45 dB en el modelo MR-CNN, y desde 18.93 dB hasta 28.13 dB en el modelo VAE. Adicionalmente, el error medio de temperatura se redujo de 1.380 °C a 0.781 °C, mientras que el error cuadrático medio (RMSE) disminuyó de 1.45 °C a 0.82 °C bajo condiciones controladas de adquisición, lo que confirma una mejora cuantitativa en la calidad estructural y en la precisión térmica de

las imágenes.

Este trabajo se enmarca en el aprendizaje automático aplicado al procesamiento de imágenes dentro de sistemas multisensor de termografía infrarroja, con énfasis en la mejora de la resolución espacial bajo un protocolo de adquisición controlado.

Palabras clave: Aprendizaje automático, Procesamiento de imágenes, Sistemas multisensor, Termografía infrarroja, Resolución espacial.

Abstract

Thermal cameras are used in industrial, medical, and scientific applications due to their ability to record spatial temperature distributions from infrared radiation. However, despite their wide applicability, these sensors often have a high cost, and lower-cost alternatives present limitations in terms of spatial resolution and measurement accuracy, which may hinder edge detection and precise identification of regions of interest. In the literature, approaches that have attempted to overcome these limitations using techniques such as interpolation, filtering, or machine learning models have shown progress. However, many of these methods rely on non-public datasets or lack well-defined acquisition protocols, overlooking sensor-specific characteristics and leaving a gap in the practical reliability of the data.

This document proposes that the combination of controlled multimodal acquisition, noise reduction techniques, and machine learning models trained on proprietary data can improve both spatial resolution and temperature measurement accuracy. To this end, a system integrating RGB and thermal cameras, a controlled acquisition protocol, and a labeled database was developed. Based on this framework, different architectures were implemented and evaluated, such as generative adversarial networks, VAE and U-Net, assessed using quantitative metrics such as PSNR and mean temperature error.

The results showed increases in PSNR from 17.43 dB to 20.45 dB for the MR-CNN model, and from 18.93 dB to 28.13 dB for the VAE model. Additionally, the mean temperature error was reduced from 1.380 °C to 0.781 °C, while the root mean square error (RMSE) decreased from 1.45 °C to 0.82 °C under controlled acquisition conditions, confirming a quantitative improvement in both structural quality and thermal measurement accuracy.

This work falls within the framework of machine learning applied to image processing in

multisensor infrared thermography systems, with an emphasis on improving spatial resolution under a controlled acquisition protocol.

Keywords: Machine learning, Image processing, Multisensor systems, Infrared thermography, Spatial resolution.

Contents

Nomenclature	10
1 Introduction	11
1.1 Context and justification	11
1.2 Motivation	14
1.3 Objectives	15
1.3.1 General Objective	15
1.3.2 Specific Objectives	15
1.4 Thesis Structure	15
1.5 Thesis Structure	15
1.5.1 Research Design	17
2 Fundamentals of Thermal Imaging, Complementary Sensors, and Hyperthermia	26
2.1 Fundamentals of Thermal Imaging and Complementary Sensors	26
2.1.1 Principles of Infrared Detection	26
2.1.2 Blackbody Radiation	27
2.1.3 Fundamental Laws of Blackbody Radiation	28
2.1.4 Emissivity and Reflected Radiation	30
2.1.5 Radiometric Measurement Model	32
2.1.6 Infrared Spectral Bands	33
2.2 Technical Limitations of Active Thermography	34
2.3 Optical and Thermal Imaging Sensors	36
2.3.1 Physical Basis and Operating Principles of Thermal Sensors	36

2.3.2	Infrared Sensors: Operating Principles and Technical Specifications . . .	38
2.3.3	Imaging Sensors Employed in the Experiments	39
2.4	Fundamentals of Hyperthermia	41
2.4.1	Physical Basis of Microwave-Tissue Interaction	41
2.4.2	Thermal Behavior of Biological Tissues	43
2.4.3	Hyperthermia in Melanoma Treatment	44
3	Thermal and Electromagnetic Characterization of Tissue-Mimicking Phantoms	46
3.1	Phantom Design and Fabrication	46
3.2	Chemical Characterization and Electromagnetic Analysis	48
3.2.1	Chemical Characterization	48
3.2.2	Thermal analysis	50
3.2.3	Optical Response	53
3.3	Experimental Setup and Reproducibility	56
3.4	Thermal Characterization under Microwave Excitation	57
3.5	THz Electromagnetic Characterization	59
4	Construction and Labeling of the Database	64
4.1	Radiometric Validation	64
4.1.1	Inter-Sensor Comparative Analysis	74
4.2	Controlled Acquisition Protocol	80
4.2.1	Ethical Considerations	84
4.3	Structure and Validation of the Database	85
5	Thermal Image Processing	93
5.1	Noise Reduction Techniques	93
5.2	Multimodal Registration and Alignment	94
5.2.1	Multimodal Fusion of Thermal and RGB Images	96
6	Integration of the Thermal Acquisition System with Closed-Loop Control	98
6.1	System Integration and Experimental Setup	98

6.2	Closed–Loop Thermal Control Evaluation	101
6.2.1	Control strategies and experimental methodology	101
6.2.2	Results and comparative analysis	105
7	Machine Learning Models for Thermal Super-Resolution	115
7.1	Implemented Architectures	116
7.1.1	Multi-Resolution Convolutional Neural Network	116
7.1.2	Multimodal Generative Adversarial Network	117
7.1.3	Progressive Variational Autoencoder	118
7.2	Training, Validation, and Testing	119
7.2.1	Multi-resolution convolutional neural network training	120
7.2.2	Multimodal generative adversarial network training	122
7.2.3	Progressive multimodal variational autoencoder training	125
7.3	Comparative Performance Analysis	127
8	Scientific Outputs and Derived Publications	131
9	Conclusions and Future Work	133
9.1	Conclusions	133
9.2	Recommendations and Future Research Lines	135

Nomenclature

ΔT Temperature difference

HR High Resolution

iPID Incremental Proportional–Integral–Derivative controller

IR Infrared

LPIPS Learned Perceptual Image Patch Similarity

LR Low Resolution

ML Machine Learning

MR Medium Resolution

MS-SSIM Multi-Scale Structural Similarity Index

PID Proportional–Integral–Derivative controller

RF Radio Frequency

RGB Red, Green, Blue

SSIM Structural Similarity Index

U-Net Convolutional encoder–decoder network with skip connections

Chapter 1

Introduction

1.1 Context and justification

Thermal cameras are used in industrial inspections, medical applications, and scientific research to obtain non-contact measurements of the spatial distribution of surface temperature. In biomedical contexts, thermal imaging plays a relevant role in applications such as inflammation monitoring [1, 2], vascular [3, 4], and hyperthermia supervision, where both spatial and temporal temperature variations are critical. However, some infrared thermal acquisition sensors, especially low-cost ones, often present limitations in terms of spatial resolution and even thermal accuracy, which can restrict processes or applications in contexts where detailed thermal analysis is required, especially in surface temperature variations.

These limitations go beyond visual quality and directly affect quantitative tasks such as thermal edge detection, segmentation of regions of interest, estimation of local temperature gradients, and reliable measurement of small temperature differences [2]. In applications where variations below one degree Celsius are relevant, spatial undersampling, sensor noise, thermal drift, or nonlinear sensor responses can introduce systematic errors that compromise the interpretation of thermal data and subsequent decision-making processes [2]. In therapeutic hyperthermia applications, the clinical target temperature range typically lies between 41°C and 45°. Temperatures below this interval may not produce the desired therapeutic effect, whereas exceeding 45°C may lead to uncontrolled tissue damage. Consequently, temperature

monitoring systems must ensure measurement uncertainties below approximately $\pm 1^\circ\text{C}$, with tighter tolerances preferred in clinical environments. These safety margins directly condition the design requirements of the proposed acquisition and processing system, particularly in terms of radiometric calibration, noise reduction, and quantitative thermal validation. As observed in Figure 1.1, the recovery of structural details in thermal images is strongly dependent on spatial resolution, evidencing that low-resolution acquisitions limit edge definition and local contrast, while higher resolutions or super-resolution post-processing allow a clearer delineation of thermal patterns.




Figure 1.1: Example of thermal images with low, medium, and high spatial resolution. The improvement in structural details is associated with acquisition or post-processing techniques such as super-resolution. Image taken from [5]

Several approaches have been proposed in the literature to mitigate these limitations [6]. Some studies focus on classical interpolation and filtering techniques to improve the apparent resolution of thermal images. Although computationally efficient, these methods are fundamentally constrained by the absence of new spatial information and often lead to over-smoothing or artificial artefacts that degrade thermal edges and local contrast.

Other studies aimed to supplement the gathered information by employing RGB or depth sensors. Using the higher spatial resolution of visible-spectrum sensors, these approaches aim to guide spatial reconstruction, alignment, and structural enhancement of thermal images. While effective in improving geometric coherence, these methods remain sensitive to registration errors, illumination changes, and scene-dependent variability.

Authors have recently suggested machine learning models trained on specific databases, demonstrating promising results in super-resolution and noise reduction [5, 7]. Architectures

such as U-Net, denoising autoencoders, and generative adversarial networks have demonstrated the ability to learn nonlinear mappings between low-resolution and high-resolution thermal representations, often incorporating multimodal inputs [8, 9]. These models offer the potential to surpass classical approaches by simultaneously learning spatial structure and sensor-specific noise characteristics.

Despite these advances, significant methodological limitations remain. Many studies rely on non-public datasets or disregard low-cost devices that exhibit nonlinear behavior and variable thermal noise [10, 11]. Furthermore, many studies do not consider controlled acquisition conditions or validation using thermally referenced quantitative metrics. This reveals a fundamental methodological gap: the lack of an integrated framework that integrates controlled thermal capture, multimodal sensing, quantitative thermal validation, and machine learning-based augmentation, with an explicit focus on accessible and low-cost thermal imaging equipment.

In response to this gap, this research proposes the development of an integrated thermal image acquisition and processing system that integrates RGB and thermal cameras, noise reduction techniques, and ML models trained with experimentally acquired data. The system emphasizes reproducibility, traceable temperature measurement, and objective evaluation of both spatial resolution and thermal accuracy. The performance is assessed using quantitative metrics such as signal-to-noise ratio and structural similarity index together with learned perceptual image patch similarity, multi-scale structural similarity, and thermal-domain indicators, including temperature distribution accuracy, thermal gradient preservation, and hot spot detection accuracy [6, 12], enabling the analysis of processing effects on both structural fidelity and thermal measurement reliability.

This thesis presents the design, implementation, and validation processes of the system. The results obtained allow for verification of the fulfillment of the objectives set forth and evaluation of its applicability in environments where it is not possible to use high-end thermal sensors for technical or budgetary reasons.

1.2 Motivation

This research is motivated by the growing requirement to obtain accurate and high-quality thermal information from low-cost infrared sensors. Low-resolution thermal cameras are increasingly used in laboratories, portable diagnostic systems, and embedded platforms; however, their measurements frequently exhibit spatial distortions, thermal noise, and temporal instability that do not correspond to actual temperature changes, but rather to the sensing device's inherent limitations. Experimental observations demonstrate inconsistencies in pixel-level temperature estimates as well as spatial distortion of thermal patterns, which have a direct impact on thermal scene interpretation and limit their applicability in quantitative analysis. These issues inspire the creation of a systematic acquisition and processing framework capable of reducing sensor-induced artifacts in controlled environments.

On the other hand, the integration of complementary sensors, such as RGB cameras, provides additional spatial and structural information that can compensate for the geometric limitations of thermal sensors. Multimodal acquisition facilitates accurate image registration, scene understanding, and structural guidance for thermal image enhancement, enabling more robust super-resolution and denoising strategies. Likewise, incorporating ML models trained with data acquired under controlled and reproducible conditions further enables the learning of complex relationships between spatial structure, thermal distribution, sensor behavior, and environmental factors. Such models can adapt to the nonlinear characteristics of low-cost thermal cameras, provided that the training data are thermally referenced and experimentally validated.

Finally, the motivation and interest in building a proprietary database originate from the scarcity of public datasets that meet the necessary conditions of alignment, thermal labeling, controlled acquisition conditions, and reproducibility. This necessitates a complete pipeline for data acquisition, labeling, and organization that supports the training, validation, and testing of ML models in a scientifically rigorous manner.

1.3 Objectives

1.3.1 General Objective

To develop an integrated system for the acquisition and processing of thermal images that combines noise reduction techniques, machine learning, and complementary sensors to improve spatial resolution and accuracy in temperature measurement.

1.3.2 Specific Objectives

- To develop a method for acquiring thermal images using complementary sensors and image processing techniques, quantitatively evaluating their impact on improving thermal accuracy and spatial resolution.
- To build and label a thermal image database that enables the training and validation of machine learning models aimed at enhancing spatial resolution and temperature measurement accuracy.
- To implement and validate machine learning architectures for improving thermal image resolution and temperature measurement accuracy supported by complementary sensors.

1.4 Thesis Structure

1.5 Thesis Structure

This thesis is organized into nine chapters, each addressing a specific aspect of the development, validation, and application of a thermal imaging system aimed at improving spatial resolution and temperature measurement accuracy.

Chapter 1 introduces the general context of the research, presenting the justification and motivation for the study. It defines the general and specific objectives, describes the research design, and outlines the structure of the thesis.

Chapter 2 provides the theoretical and technical foundations required for this work. It reviews the principles of infrared thermography, blackbody radiation, emissivity, and infrared spectral bands, as well as the operating principles and limitations of thermal and complementary imaging sensors. In addition, the chapter introduces the fundamentals of hyperthermia, including microwave–tissue interaction mechanisms, thermal behavior of biological tissues, and its application in melanoma treatment.

Chapter 3 focuses on the thermal and electromagnetic characterization of tissue-mimicking phantoms. It describes the phantom design and fabrication process, its thermal response under microwave excitation, its electromagnetic characterization in the terahertz frequency range, and the experimental setups employed to ensure measurement reproducibility.

Chapter 4 presents the construction and labeling of the thermal image database used in this research. This chapter details the radiometric validation process, the controlled acquisition protocol, and the structure and validation procedures applied to ensure data consistency.

Chapter 5 addresses thermal image processing techniques. It covers noise reduction methods and multimodal registration and alignment strategies, including the fusion of thermal and RGB images, which are required for subsequent model development.

Chapter 6 describes the integration of the thermal acquisition system with a closed-loop control framework. It presents the system integration and experimental setup, followed by the evaluation of different closed-loop thermal control strategies.

Chapter 7 focuses on the machine learning models developed for thermal image super-resolution. It describes the implemented architectures and the training, validation, and testing procedures, and presents the performance evaluation based on quantitative metrics.

Chapter 8 presents the scientific outputs derived from this research, including publications and related academic contributions.

Finally, Chapter 9 summarizes the main conclusions of the research and discusses recommendations and future research directions derived from the results obtained in this work.

1.5.1 Research Design

This research follows an experimental and quantitative design. The experimental approach is based on controlled laboratory acquisition conditions and the implementation of different thermal image super-resolution architectures. The quantitative nature of the study is reflected in the use of numerical performance metrics, including Peak Signal-to-Noise Ratio (PSNR), Root Mean Square Error (RMSE), and temperature error values expressed in degrees Celsius, to evaluate spatial resolution enhancement and thermal measurement accuracy.

List of Figures

1.1	Example of thermal images with low, medium, and high spatial resolution. The improvement in structural details is associated with acquisition or post-processing techniques such as super-resolution. Image taken from [5]	12
2.1	Representation emitted, and reflected infrared radiation on a surface with partial emissivity. Image generated by AI using Google AI Studio	31
2.2	Infrared spectral bands from near-infrared (NIR) to far-infrared (FIR), showing typical wavelength ranges and applications. Image generated by AI using Google AI Studio.	33
2.3	Schematic representation of the RGB sensing process, illustrating how incident visible light is filtered through a Bayer pattern and converted into digital color information. Image generated by AI using Google AI Studio.	38
2.4	Schematic representation of the microwave–tissue interaction process. Microwaves induce dielectric heating by exciting polar molecules and ions, leading to temperature rise within biological structures. The resulting thermal diffusion and perfusion determine the spatial distribution and biological effects of heat. Image generated by AI using Google AI Studio.	43
3.1	Schematic representation of the tissue-mimicking phantom, including geometry and material composition used to emulate the thermal and dielectric properties of biological soft tissues.	46
3.2	EDS elemental distribution map of the tissue-mimicking phantom, showing the spatial distribution of carbon, sodium, and chloride. Adapted from [13].	49

3.3	Thermogravimetric analysis (TGA) of the tissue-mimicking phantom.	51
3.4	Thermal properties of the phantom evaluated at different temperatures. (a) Thermal conductivity (b) Specific heat capacity.	52
3.5	Experimental setup used for VIS–NIR diffuse reflectance measurements of the tissue-mimicking phantom.	54
3.6	VIS–NIR optical characterization of the tissue-mimicking phantom. (a) Mean normalized reflectance spectrum with standard deviation. (b) Apparent absorbance spectrum derived from reflectance. (c) Wavelength-dependent coefficient of variation of absorbance.	54
3.7	Principal component analysis of the VIS–NIR apparent absorbance spectra measured at multiple positions on the phantom surface.	56
3.8	Experimental setup for microwave heating and thermal acquisition.	57
3.9	Representative infrared images of the tissue-mimicking phantom acquired at different exposure times under microwave excitation.	58
3.10	Temporal evolution of phantom surface temperature under 20 dBm microwave excitation.	58
3.11	Temporal evolution of phantom surface temperature under 10 dBm microwave excitation.	59
3.12	Schematic of the SHR reflection configuration, showing the THz beam path from the Tx antenna to the sample and back to the Rx antenna.	60
3.13	THz-TDS of the phantom (0.5 cm thickness). (a) Time-domain traces: reference (mirror) and 15 sequential acquisitions of the phantom, with a zoomed view of the window of interest. (b) Corresponding frequency-domain spectra, with a magnified view in 0.05–2 THz. (c) Normalized amplitude metrics (peak-to-peak and RMS) computed within the same time window. (d) Drift expressed as an equivalent optical path-length change estimated from the pulse shift and linear regression.	62

3.14	THz-TDS of the phantom (0.3 cm thickness). (a)–(b) Same as Fig. 3.13 for 15 sequential acquisitions. (c)–(d) Amplitude metrics and drift computed over the initial subset of measurements for which the pulse remains within the fixed temporal window used in the analysis.	63
4.1	Experimental setup used for the calibration procedure: (A) Ametek Jofra CTC–650A dry–block temperature calibrator, (B) thermal cameras under evaluation, and (C) host computer for data acquisition and control.	65
4.2	Mean temperature measured by the Optris camera and the Thermal Camera HAT as a function of the reference temperature.	67
4.3	Boxplots of the thermal measurements obtained with the Optris camera and the Thermal Camera HAT for each reference temperature.	67
4.4	Thermal image mosaic acquired with the Optris camera at different reference temperatures and working distances.	71
4.5	Thermal image mosaic acquired with the FLIR camera at different reference temperatures and working distances.	71
4.6	Cross-comparison between Optris and Thermal Camera HAT measurements. Markers indicate mean values across five repeated measurements per setpoint, while error bars represent the corresponding standard deviation for each sensor. The solid line corresponds to the linear regression fit, and the dashed line indicates the ideal identity relation.	75
4.7	Cross-comparison between Optris and FLIR temperature measurements acquired at a working distance of 10 cm. Each point represents the mean temperature obtained from five independent measurement sweeps for a given setpoint. Horizontal and vertical error bars indicate the corresponding standard deviations of the Optris and FLIR measurements, respectively. The solid line represents the linear regression between both sensors.	77
4.8	Normalized radar chart comparing the statistical performance of the Waveshare, Optris, and FLIR thermal cameras based on R^2 , RMSE, bias, coefficient of variation, and standard deviation.	79

4.9	Schematic representation of the RF heating and thermal imaging system. The configuration includes: (A) RF signal generator, (B) microwave amplifier, (C) circulator, (D) matched load, (E) directional antenna, and (F) infrared cameras positioned at multiple viewing angles. Adapted from [13].	81
4.10	Spatial configuration used during data acquisition. Cameras (A1, A2) correspond to the HR and MR thermal systems mounted on tripods (E1, E2). The antenna (C) was placed above the sample (B) at a constant 5 mm distance, with a black curtain (D) to reduce reflections. Adapted from [13].	83
4.11	Hierarchical structure of the thermal image database organized by spatial resolution (HR, MR, LR). Each directory contains original and augmented subfolders. Adapted from [13].	86
4.12	High-resolution (HR) thermal images after minimal augmentations preserving structural integrity. Adapted from [13].	88
4.13	Medium-resolution (MR) thermal images with moderate augmentations, including rotation and noise injection. Adapted from [13].	89
4.14	Low-resolution (LR) thermal images with aggressive augmentations simulating noisy or degraded conditions. Adapted from [13].	90
4.15	Distribution of thermal images by material type across all resolution levels. Adapted from [13].	91
4.16	Example of thermal images recorded with the three infrared cameras under controlled conditions. (A) Thermal Camera HAT (80×60), (B) Optris PI 450i (382×288), and (C) FLIR A655sc (640×480). The sample corresponds to pork skin heated up to 45°C at a 30° viewing angle. Adapted from [13].	92
5.1	Example of the multimodal registration process. From left to right: RGB image, unaligned thermal image, and the result after applying the geometric registration model.	94
5.2	Multimodal fusion workflow. The process starts from the initial thermal and RGB images, applies geometric registration, and subsequently performs the selected fusion strategy.	96

5.3 Comparison of the implemented multimodal fusion methods. The image shows the results obtained using wavelet fusion, multiscale pyramid fusion, and blending, all applied to the same registered image pair. 97

6.1 Schematic representation of the integrated closed-loop thermal system. The architecture includes: (A) USRP - 2922 SDR acting as RF generator, (B) RF linear power amplifier, (C) host computer executing the control algorithms, (D) 2.45 GHz applicator antenna, (E) thermal camera for monitoring, and (F) biological or phantom sample. 99

6.2 Photograph of the experimental mounting system. The structure holds the RF applicator antenna (c), the Thermal Camera HAT used for closed-loop control(b), and the FLIR and Optris cameras(a) employed as reference instruments for validation. 100

6.3 Open-loop thermal response of the system obtained under maximum RF excitation, after low-pass filtering. 101

6.4 Flowchart of the hybrid PID and hybrid incremental PID control architectures. When the temperature error exceeds the regulation band, a constant RF actuation is applied. Inside the 8 °C region below the setpoint, the controller switches to PID or iPID regulation. 103

6.5 Closed-loop thermal regulation scheme used in all experiments, including RF excitation, thermal sensing, control algorithm and actuation stages. 105

6.6 Closed-loop thermal response of the conventional PID controller using the phantom sample. From top to bottom: maximum temperature evolution, temperature error, and RF gain applied to the actuator. 107

6.7 Closed-loop response using the hybrid PID controller with the phantom sample. From top to bottom: thermal response, temperature error and controller gain. . 109

6.8 Closed-loop response using the hybrid PID controller with the pork meat sample. From top to bottom: thermal response, temperature error and controller gain. 110

6.9	Closed-loop thermal response using the hybrid incremental PID controller with the phantom sample. From top to bottom: temperature evolution, temperature error and controller gain.	112
6.10	Closed-loop thermal response using the hybrid incremental PID controller with the pork meat sample. From top to bottom: temperature evolution, temperature error and controller gain.	113
7.1	Flow diagram of the multi-resolution convolutional neural network. LR and MR thermal images are processed through independent convolutional branches, fused at the feature level, and reconstructed into an HR thermal output.	117
7.2	Flow diagram of the multimodal generative adversarial network.	118
7.3	Flow diagram of the progressive variational autoencoder for multi-resolution thermal image modeling.	119
7.4	Representative qualitative samples showing low-resolution input (LR), super-resolved output (SR), and high-resolution reference (HR) thermal images obtained during training.	121
7.5	Evolution of training metrics for the MR-CNN, including total loss, PSNR, and temperature error (Temp-L1) over 50 epochs.	122
7.6	Evolution of training metrics for the multimodal GAN, including total loss, loss components, PSNR, and temperature error (Temp-L1) over the training epochs.	123
7.7	Representative qualitative samples showing low-resolution input (LR), super-resolved output (SR), and high-resolution reference (HR) thermal images generated by the multimodal GAN during training.	124
7.8	Evolution of reconstruction-related metrics during progressive VAE training, including PSNR across epochs.	126
7.9	Representative reconstruction samples obtained during progressive VAE training at epoch 5, showing LR inputs and corresponding reconstructed thermal images.	127

List of Tables

2.1	Main specifications of the infrared cameras employed in the experiments.	40
2.2	Main specifications of the RGB imaging sensors used in the experimental setups.	41
2.3	Representative thermal properties of biological tissues at normothermic conditions, based on data from [14–16].	44
3.1	Weight percentage (wt.%) of the elements identified in the phantom sample by EDS analysis. Adapted from [13].	50
3.2	Integrated apparent absorbance of the tissue-mimicking phantom over selected VIS–NIR spectral bands. Values are reported as mean \pm standard deviation across measurement positions.	55
4.1	Inter–sensor measurements obtained with the Optris and the Thermal Camera HAT (Waveshare) for each reference temperature. Each value corresponds to the ROI–averaged temperature extracted from five independent measurement sweeps.	66
4.2	Descriptive statistics of the inter–sensor measurements for the Optris and the Thermal Camera HAT (Waveshare).	68
4.3	Thermal measurements obtained at different acquisition distances using the Optris and FLIR cameras. Each value corresponds to the average temperature extracted from the ROI for each measurement round.	70
4.4	Descriptive statistics of the thermal measurements obtained with the Optris and FLIR cameras at different acquisition distances.	72
4.5	Materials included in the database and their respective labels. Adapted from [13].	85

4.6	Initial and final surface temperatures for each material at 5.5 dBm RF power. Ambient temperature ranged between 16–19°C. Adapted from [13].	89
7.1	Final quantitative performance over the independent testing subset (mean ± standard deviation).	128

Chapter 2

Fundamentals of Thermal Imaging, Complementary Sensors, and Hyperthermia

2.1 Fundamentals of Thermal Imaging and Complementary Sensors

2.1.1 Principles of Infrared Detection

Infrared thermography is a non-contact imaging technique to acquire and analyze spatial temperature distributions by detecting the thermal radiation emitted by objects as a direct consequence of their temperature. This technique enables the measurement of energy within the infrared (IR) region of the electromagnetic spectrum, generally ranging from $0.7 \mu\text{m}$ to $14 \mu\text{m}$. Although this radiation is invisible to the human eye, it provides valuable information for estimating surface temperature distributions and thermal gradients on the surface of objects and systems.

From a physical standpoint, thermal cameras do not directly measure temperature; instead, they measure radiative energy, which is subsequently converted into temperature values through radiometric models. This distinction is fundamental, as the detected signal depends not only on the object's temperature but also on its radiative properties and the surrounding environment.

Any object with a temperature above absolute zero (0 K, $-273.15\text{ }^{\circ}\text{C}$) emits electromagnetic radiation due to the thermal motion of its particles, and well-established physical laws of thermal radiation govern both the intensity and spectral distribution of this radiation. However, the conversion of detected IR radiation into accurate temperature estimates requires consideration of additional factors beyond temperature alone. Surface emissivity, reflected ambient radiation, atmospheric transmittance, sensor spectral response, and background conditions all contribute to the radiometric signal received by the thermal detector. Neglecting these factors can lead to systematic measurement errors, particularly when using low-cost thermal sensors that exhibit higher noise levels and nonlinear response behavior.

Based on these considerations, several physical principles are required to explain the processes of emission, propagation, and detection of IR radiation. These principles form the foundation for understanding the limitations of thermal imaging systems, the sources of measurement uncertainty, and the motivation for subsequent processing techniques such as noise reduction, multimodal fusion, and machine learning-based enhancement. The most relevant of these principles are introduced below.

2.1.2 Blackbody Radiation

A blackbody is defined as an object that absorbs all incident electromagnetic radiation, regardless of wavelength or direction, without reflection or transmission and re-emits energy with a spectral distribution determined only by its absolute temperature [17]. As such, a blackbody represents the maximum possible radiative emission for a given temperature and serves as a universal reference in thermal radiation theory.

In practical terms, a blackbody can be approximated by a cavity with a small opening and diffusely reflective internal walls. Radiation entering the aperture undergoes multiple reflections and is almost completely absorbed, so the radiation that leaves the cavity depends only on its internal temperature. This principle is widely exploited in radiometric calibration sources and reference standards for thermal imaging systems [18].

In IR thermography, the blackbody model is fundamental for defining the emissivity of real surfaces, which expresses how efficiently a material emits energy compared to a perfect blackbody under the same thermal conditions. Since most real materials exhibit emissivity values lower than unity and may vary with wavelength, surface finish, and temperature, accurate temperature estimation requires either prior knowledge or experimental calibration of emissivity. As highlighted by Usamentiaga *et al.*, emissivity is one of the most critical factor in IR temperature measurement because it directly affects the relationship between detected radiance and actual surface temperature [19]. Errors in emissivity estimation can propagate nonlinearly into temperature measurements, particularly in low-emissivity surfaces or when operating under uncontrolled environmental conditions.

2.1.3 Fundamental Laws of Blackbody Radiation

The quantitative description of thermal radiation emitted by an ideal blackbody is governed by three fundamental physical laws that establish the relationship between temperature, emitted energy, and spectral distribution. These laws not only laid the groundwork for quantum theory but also remain essential for modern radiometry, IR thermography, and thermal sensor calibration. In the context of thermal imaging, these laws provide the physical basis for understanding how temperature information is encoded in IR radiation and how this information is subsequently detected, processed, and interpreted by thermal cameras.

Planck's Law describes the spectral radiance $L_\lambda(\lambda, T)$ of a blackbody at an absolute temperature T as a function of wavelength λ :

$$L_\lambda(\lambda, T) = \frac{2\pi hc^2}{\lambda^5} \cdot \frac{1}{\exp\left(\frac{hc}{\lambda k_B T}\right) - 1} \quad (2.1)$$

where h is Planck's constant, c is the speed of light in vacuum, and k_B is Boltzmann's constant. This formulation solved the so-called ultraviolet catastrophe predicted by classical physics and shows that the emitted energy depends solely on the temperature and wavelength of the radiation [19–21]. Importantly, Planck's law shows that the radiative signal detected by a thermal camera depends simultaneously on wavelength and temperature, which directly

motivates the use of specific IR spectral bands in thermographic systems.

From a thermographic perspective, Planck's law explains why IR sensors measure radiance rather than temperature directly, and why accurate temperature estimation requires radiometric calibration and compensation. In practical systems, detectors operate over limited wavelength ranges and exhibit non-uniform spectral response, meaning that the measured radiance represents only a fraction of the total emitted energy.

From Planck's distribution emerges Wien's Displacement Law, which defines the wavelength λ_{\max} at which the spectral radiance reaches its maximum value:

$$\lambda_{\max} T = b \quad (2.2)$$

where $b \approx 2.897 \times 10^{-3}$ m K. This law indicates that as temperature increases, the peak of the emitted spectrum shifts toward shorter wavelengths. For objects at ambient and physiological temperatures (approximately 300–320 K), this peak lies within the long-wave infrared (LWIR) region, which explains the widespread use of LWIR cameras in biomedical and industrial thermography.

Although Wien's law is generally valid, several studies have discussed its deviations depending on the representation domain and specific conditions [21–23]. These discussions highlight that the peak wavelength is a theoretical construct and that real thermal cameras capture radiation over a finite spectral band, making detector bandwidth and spectral sensitivity key factors in practical temperature measurement

Finally, the Stefan–Boltzmann Law expresses the total radiant power M emitted per unit area by a blackbody, integrating the spectral emission over all wavelengths and directions:

$$M = \sigma T^4, \quad (2.3)$$

where $\sigma \approx 5.67 \times 10^{-8}$ W m⁻² K⁻⁴ is the Stefan–Boltzmann constant. This law is derived from the integration of Planck's distribution and establishes that the emitted power increases with the fourth power of temperature [17, 18, 24]. This strong nonlinear relationship implies that small temperature variations can produce measurable changes in radiative power, but it also

highlights the increasing sensitivity of thermal measurements at higher temperatures.

From a practical standpoint, the Stefan–Boltzmann law is essential for understanding the dynamic range and sensitivity limits of infrared detectors, as well as the challenges associated with measuring small temperature differences at low contrast levels. In low-cost thermal cameras, detector noise, quantization, and calibration uncertainty can significantly obscure the ideal T^4 dependence predicted by theory.

Together, these three laws provide the theoretical foundation for describing radiative transfer, defining emissivity in real materials, and performing calibration procedures in infrared thermographic systems [19]. They also define the fundamental physical limits that post-processing and ML-based enhancement methods cannot surpass, even if apparent spatial resolution or perceptual image quality is improved.

2.1.4 Emissivity and Reflected Radiation

When IR radiation reaches a real surface, it can be absorbed, reflected, or transmitted through the material. In the case of opaque materials, transmission is negligible, so only absorption and reflection are relevant. The relationship between absorptivity α , reflectivity ρ , and emissivity ε can then be expressed as $\varepsilon = \alpha = 1 - \rho$, following Kirchhoff’s law, which states that under thermal equilibrium, a material’s emissivity is equal to its absorptivity. Emissivity therefore, quantifies the efficiency with which a surface emits thermal radiation relative to an ideal blackbody at the same temperature and wavelength. Since most real materials have emissivity values lower than unity, the radiative signal detected by a thermal camera is not purely emitted radiation but a combination of emitted and reflected components.

Figure 2.1 illustrates these interactions conceptually. Part of the incoming radiation (represented by red streamlines) is absorbed and re-emitted by the surface as thermal radiation, while another portion is reflected depending on the surface’s reflectivity. The balance between emitted and reflected energy depends on several factors, including temperature, wavelength, and material characteristics such as roughness, oxidation, moisture content, and coatings.

In infrared thermography, emissivity and reflected radiation are among the most critical sources of measurement uncertainty. The signal recorded by a thermal camera represents a




Figure 2.1: Representation emitted, and reflected infrared radiation on a surface with partial emissivity. Image generated by AI using Google AI Studio

superposition of emitted radiation from the object and reflected radiation from the environment, meaning that incorrect emissivity settings or neglecting reflections can lead to substantial temperature estimation errors.

Previous studies have shown that small inaccuracies in emissivity values can produce temperature deviations of several degrees Celsius, particularly when environmental reflections are ignored. Barreira et al. [25] demonstrated that moisture-induced emissivity variations can exceed 10%, significantly affecting temperature readings. Similarly, Ogasawara et al. [26] proposed a method to measure emissivity using polarized reflection near the Brewster angle, showing that surface roughness and polarization effects play an important role in accurately estimating emissivity for opaque materials.

In practical thermographic work, reflected radiation is often characterized by estimating the apparent reflected temperature ($T_{\text{reflected}}$) of the surroundings. For materials with low emissivity, such as metals, the reflected component can dominate the detected signal, leading to over- or underestimation of surface temperature if not corrected. To reduce these errors, it is common to use high-emissivity reference patches or measure the reflected temperature with reflective foils

placed in the camera's field of view, which helps calibrate the system and improve measurement accuracy [19, 25, 26].

2.1.5 Radiometric Measurement Model

Infrared temperature measurement is governed by the radiometric balance between the radiation emitted by the object, the radiation reflected from the surroundings, and the contribution of the propagation medium. According to classical radiative transfer theory and infrared thermography principles [27–29], the spectral radiance detected by an infrared sensor can be expressed, in simplified form, as

$$L_{\text{det}} = \tau \left[\varepsilon L_{bb}(T_{\text{obj}}) + (1 - \varepsilon) L_{bb}(T_{\text{ref}}) \right] + L_{\text{atm}},$$

where L_{det} is the radiance reaching the detector, ε is the surface emissivity, $L_{bb}(T_{\text{obj}})$ is the blackbody radiance at the object temperature, $L_{bb}(T_{\text{ref}})$ represents the reflected ambient radiance, τ is the atmospheric transmittance along the optical path, and L_{atm} accounts for atmospheric self-emission.

This formulation shows that the temperature estimated by an infrared camera is determined not only by the object temperature, but also by surface emissivity, reflected ambient radiation, atmospheric transmittance (which depends on the propagation distance), and acquisition geometry, including viewing angle and field of view. Consequently, variations in these parameters may generate systematic offsets or increased measurement dispersion if the experimental conditions are not adequately controlled.

From a methodological standpoint, this radiometric framework provides the physical reference required to evaluate the necessity of an explicit calibration procedure. Instead of presupposing the need for corrective transformations, the experimental stages described in Chapter 4 were designed to quantify the overall measurement deviation under controlled conditions of distance, alignment, and ambient stability. By experimentally assessing the magnitude of dispersion and bias within the operational temperature range, it becomes possible to determine whether the combined influence of emissivity assumptions, reflected radiation, and geometric

configuration justifies the implementation of an additional calibration step.

2.1.6 Infrared Spectral Bands

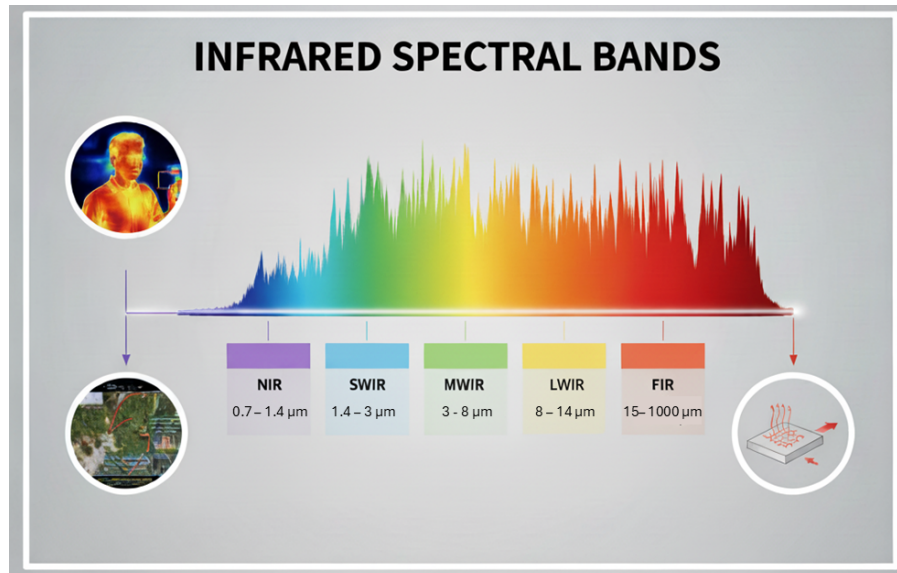


Figure 2.2: Infrared spectral bands from near-infrared (NIR) to far-infrared (FIR), showing typical wavelength ranges and applications. Image generated by AI using Google AI Studio.

Infrared (IR) radiation occupies the region of the electromagnetic spectrum between the visible and microwave ranges, extending approximately from 0.7 μm to 1000 μm in wavelength. This spectral region is generally subdivided into near-infrared (NIR, 0.7–1.4 μm), short-wave infrared (SWIR, 1.4–3 μm), mid-wave infrared (MWIR, 3–8 μm), long-wave infrared (LWIR, 8–14 μm), and far-infrared (FIR, 15–1000 μm), as illustrated in Figure 2.2. Each band corresponds to distinct physical emission and interaction mechanisms, determining their suitability for applications in remote sensing, thermography, and spectroscopy.

The atmospheric transmission of infrared radiation is not uniform across the spectrum; absorption bands primarily due to water vapor, carbon dioxide, and ozone define windows of high transmittance. These windows correspond to the MWIR and LWIR regions, which are therefore widely used for thermographic imaging and temperature measurement. In these bands, objects near ambient temperature ($\approx 300\text{ K}$) emit strongly according to Planck's law, making them particularly relevant for non-contact temperature estimation and material characterization.

The analysis of hyperspectral reflectance data in the range of 6 to 14.5 μm has shown that thermal infrared wavelengths can effectively differentiate between natural and human-made materials such as plastics and minerals. Studies have revealed that most of the diagnostic wavelengths fall within the main atmospheric transmission windows of the thermal infrared spectrum, highlighting the importance of the MWIR and LWIR regions for remote sensing and material identification [30]. Combining spectral information from both regions has proven useful for improving the quantitative estimation of mineral compositions. Using regression models such as PLSR, PCR, and SVR, researchers have found that integrating MWIR and LWIR data increases the accuracy of mineralogical predictions, showing that these two spectral ranges complement each other in material classification and temperature estimation [31].

From an engineering perspective, research on MWIR and LWIR imaging systems has shown that the characteristics of the detector and the choice of wavelength are decisive factors for detection, recognition, and identification performance. MWIR systems tend to perform better in low-humidity environments and when observing high-temperature targets, while LWIR configurations are more effective for passive detection of objects at or near ambient temperature, as they are particularly sensitive to blackbody radiation within the 8–14 μm range [32].

A comprehensive understanding of these spectral bands is essential for accurate thermographic interpretation. The detected radiative signal depends on parameters such as surface temperature, emissivity, and the chosen wavelength range, in addition to atmospheric effects. Consequently, precise thermographic measurements require selecting sensors adapted to the appropriate spectral window and implementing emissivity and reflectance corrections consistent with the optical behavior of the materials under observation.

2.2 Technical Limitations of Active Thermography

The performance of active thermography is mainly limited by how heat diffuses through materials. When an external energy source heats the surface, the resulting temperature field evolves according to the heat diffusion equation, which causes the thermal signal to decrease exponentially with depth. The ability of the excitation to penetrate the material is described by the

thermal diffusion length, a parameter that increases with the square root of the thermal diffusivity and decreases with the square root of the modulation frequency. As a result, the deeper the region being analyzed, the weaker its contribution to the surface temperature detected by the sensor. This behavior creates an inherent compromise between the capacity to detect deep subsurface features and the spatial resolution achievable on the surface [33, 34].

The diffusive nature of heat propagation also imposes a low-pass filtering effect on the temperature field, where high spatial frequencies corresponding to small subsurface features are progressively attenuated during the propagation to the surface. This results in a blurring of the thermal image and limits the capacity of the technique to resolve fine details. The inverse problem of reconstructing subsurface temperature distributions from surface measurements becomes ill-posed, as small variations in the data can cause large uncertainties in the estimated parameters. Regularization techniques can mitigate this effect, but they introduce model dependency and reduce quantitative accuracy [33].

Another major limitation arises from the non-ideal conditions of thermal excitation. In practice, the assumption of uniform surface heating is rarely satisfied, leading to spatial variations in the initial temperature field that affect the linearity of the thermal response. Even small deviations in excitation intensity or angle modify the boundary conditions of the heat equation, introducing systematic errors in the estimation of defect depth or thermal properties. These effects are particularly critical when short pulse durations or localized energy sources are employed, as they increase the influence of boundary and emissivity variations [35].

The properties of the material strongly affect how heat spreads during active thermography. When the material is not uniform or conducts heat differently in various directions, the temperature does not propagate evenly. This behavior is described by the thermal diffusivity tensor, which represents how heat moves along different axes. In materials made of several layers or components, such as composites, changes in thermal conductivity and heat capacity between layers cause complex temperature patterns. In these cases, the surface temperature response no longer follows simple linear models, which makes it difficult to relate the measured signal directly to the depth or size of internal defects [36].

The radiometric detection process also imposes constraints. Infrared cameras operate within

specific spectral bands, typically within the 3–5 μm (MWIR) or 8–14 μm (LWIR) regions, where atmospheric transmittance and detector sensitivity vary. The recorded signal depends not only on the temperature field but also on surface emissivity, reflectivity, and ambient radiation. These factors introduce multiplicative and additive noise components, reducing the effective signal-to-noise ratio and limiting the detection threshold for small temperature contrasts. The nonlinearity between radiance and temperature at high thermal gradients further complicates calibration and quantitative interpretation [35].

In active thermography, the ability to distinguish features at different depths depends on the thermal diffusion length associated with the chosen modulation frequency. Because heat diffuses gradually through materials, there is a limited range of usable frequencies. High modulation frequencies improve the sharpness of the detected image but restrict how deep the thermal wave can reach. In contrast, low frequencies allow the heat to penetrate deeper but reduce the level of detail on the surface image. Selecting the appropriate frequency therefore requires balancing resolution and depth according to the material's thermal properties and the expected position of internal defects [34]. This trade-off makes it physically impossible to obtain both high resolution and high depth sensitivity at the same time.

The system's temporal performance is also limited by the response speed and sensitivity of the infrared detector. In pulsed thermography, the temperature variations occur very quickly after the energy input, so the detector must record data at a high sampling rate to accurately capture these rapid changes. If the acquisition speed or dynamic range is insufficient, the recorded signal becomes distorted or loses detail, which affects the accuracy of the defect analysis.

2.3 Optical and Thermal Imaging Sensors

2.3.1 Physical Basis and Operating Principles of Thermal Sensors

Thermal cameras are based on infrared sensors designed to detect electromagnetic radiation within specific spectral bands, typically the mid-wave infrared (MWIR, 3–5 μm) and the long-wave infrared (LWIR, 7.5–14 μm) regions. These spectral windows are selected because they

coincide with atmospheric transmission bands and with the peak thermal emission of objects at ambient and physiological temperatures. The operation of thermal sensors relies on the conversion of incident infrared radiation into electrical signals, which can be achieved through two main detection mechanisms: photon-based and thermal-based detection [37–39].

Photon detectors, such as mercury cadmium telluride (HgCdTe, MCT) and indium antimonide (InSb) photodiodes, operate by directly converting incident infrared photons into electrical charge. When photons with sufficient energy are absorbed by the semiconductor material, charge carriers are generated, producing an electrical current proportional to the incident radiance. These detectors exhibit high sensitivity and fast temporal response, making them suitable for applications requiring high frame rates or low noise levels. However, their operation is strongly affected by thermally generated carriers, which introduce significant noise. As a result, photon detectors typically require cryogenic cooling to suppress thermal noise and ensure stable and accurate measurements. The choice of semiconductor material determines the detector's spectral response, as each material is sensitive to a specific range of photon energies [38, 40].

Thermal detectors, in contrast, respond to infrared radiation through a temperature increase induced by absorbed energy rather than through direct photon-electron interactions. This category includes bolometers, pyroelectric detectors, and thermopiles. In uncooled microbolometer-based cameras, which are widely used in LWIR imaging, the absorbed radiation raises the temperature of a sensing element, commonly made of vanadium oxide (VO_x) or amorphous silicon (a-Si). This temperature change modifies the electrical resistance of the material, and the resulting resistance variation is measured as a voltage signal proportional to the incident radiance. The performance of microbolometers is governed by parameters such as thermal isolation, heat capacity, responsivity, and noise behavior. A key figure of merit is the Noise Equivalent Temperature Difference (NETD), which represents the smallest temperature variation that can be distinguished from the sensor noise [37, 41].

Thermal imaging sensors are integrated into focal plane arrays (FPAs), where each element acts as an independent detector pixel. Advances in microelectromechanical systems (MEMS) technology have allowed the fabrication of compact, high-density uncooled arrays with improved uniformity, reduced noise, and enhanced response time [38, 40]. The optical system focuses

the infrared radiation onto the detector array, and the resulting analog signals are digitized, corrected, and processed to generate a thermal image in which each pixel represents the apparent temperature or radiance of a corresponding region in the scene.

The performance of a thermal camera is therefore determined by the combined characteristics of the detector material, sensor architecture, spectral response, and signal processing chain. These factors directly influence sensitivity, accuracy, and spatial resolution, and ultimately define the reliability of thermal measurements in scientific, industrial, and biomedical applications [33, 39, 41].

2.3.2 Infrared Sensors: Operating Principles and Technical Specifications

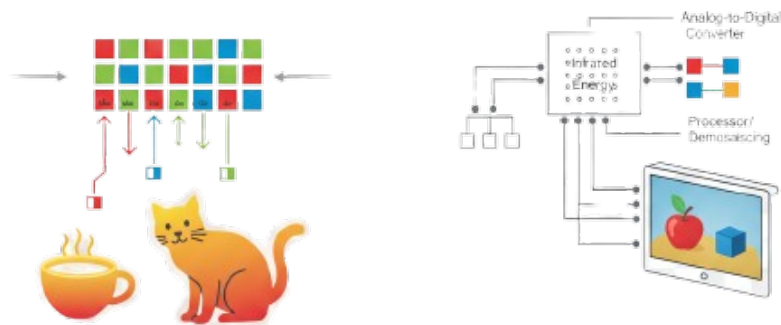


Figure 2.3: Schematic representation of the RGB sensing process, illustrating how incident visible light is filtered through a Bayer pattern and converted into digital color information. Image generated by AI using Google AI Studio.

RGB sensors operate within the visible portion of the electromagnetic spectrum, approximately from 400 to 700 nm, and acquire images by measuring the intensity of incident light on an array of photosensitive elements. To enable color perception, these sensors are typically combined with a color filter array, most commonly the Bayer mosaic, where each pixel is covered by a red, green, or blue filter. As a result, each pixel records information corresponding to a single color channel, and the full-color image is subsequently reconstructed through digital demosaicing and interpolation algorithms.

These sensors are commonly implemented using charge-coupled device (CCD) or complementary metal-oxide-semiconductor (CMOS) architectures. In both cases, incident photons

generate electrical charges within the photosensitive elements, which are then converted into voltage signals and digitized by an analog-to-digital converter. Figure 2.3 schematically illustrates this process, highlighting the conversion from incident visible light to digital color information.

RGB cameras are characterized by high spatial resolution, low cost, and widespread availability, which has led to their extensive use in scientific, industrial, and consumer imaging applications. Their high pixel density enables detailed spatial representation of scene geometry, edges, and textures, making them particularly suitable for tasks such as object detection, feature extraction, and image registration. These characteristics are especially relevant when RGB sensors are used in conjunction with lower-resolution sensing modalities. However, RGB sensors present inherent limitations. Their performance strongly depends on ambient illumination conditions, and their sensitivity degrades significantly in low-light or uncontrolled lighting environments. Moreover, RGB cameras are incapable of directly measuring temperature or capturing radiative information outside the visible spectrum [42–44]. As a result, they cannot provide thermal information or operate reliably in scenarios where visible illumination is insufficient or irrelevant. These limitations restrict their standalone use in thermographic applications but motivate their integration as complementary sensors in multimodal imaging systems, where visible structural information can support and enhance thermal image analysis [9,45]

2.3.3 Imaging Sensors Employed in the Experiments

Two imaging modalities were employed in the experimental setup: thermal infrared cameras and an RGB camera. The thermal sensors were used for the acquisition of temperature fields and for the detection of thermal contrasts induced by microwave excitation, while the RGB camera provided high-resolution visual information to support spatial reference and multimodal alignment. The main technical characteristics of the imaging systems are described below.

Thermal Imaging Sensors

Table 2.1 summarizes the main specifications of the infrared cameras used in the experiments. The selected devices differ in spatial resolution, spectral range, field of view (FoV), and thermal

sensitivity, quantified through the Noise Equivalent Temperature Difference (NETD). These parameters directly influence the capability of each sensor to resolve spatial details and detect small temperature variations under active thermographic conditions.

Table 2.1: Main specifications of the infrared cameras employed in the experiments.

Camera	Resolution	Spectral range	FoV	NETD
FLIR A655sc (HR)	640×480	7.5–14.0 μm	$15^\circ \times 11^\circ$	<30 mK
Optris PI 450i (MR)	382×288	8–14 μm	Interchangeable lenses	≤ 40 mK
Thermal Camera HAT (LR)	80×60	8–14 μm	$45^\circ \times 45^\circ$	≈ 150 mK

The spatial resolution defines the number of pixels in the focal plane array, determining the level of spatial detail in the image. The spectral range indicates the infrared wavelengths detected by the sensor, corresponding here to the long-wave infrared (LWIR) region from 7.5 to 14 μm . The field of view (FoV) specifies the angular coverage of the optics, determining the observable area. Finally, the Noise Equivalent Temperature Difference (NETD) measures the minimum resolvable temperature variation; lower NETD values indicate higher sensitivity.

The selected cameras were intentionally chosen to represent three distinct categories: high-resolution, medium-resolution, and low-resolution thermal imaging systems. This selection enables a systematic evaluation of how spatial resolution and thermal sensitivity affect defect characterization, temperature field reconstruction, and the accuracy of thermal–RGB data fusion. By considering sensors with markedly different performance levels, the experimental framework allows the analysis of limitations inherent to low-cost devices and the assessment of processing strategies aimed at mitigating these constraints.

Thermal imaging sensors exhibit intrinsic measurement noise that affects the stability of the radiometric signal. In uncooled microbolometer-based devices, this behavior is commonly quantified through the Noise Equivalent Temperature Difference (NETD), which represents the minimum temperature variation distinguishable from sensor noise [38, 41]. As indicated in Table 2.1, the low-resolution sensor presents a higher NETD compared to the high-resolution system, implying reduced thermal sensitivity and increased pixel-level dispersion.

Since the learning-based enhancement models employed in this work operate directly on the acquired image data, the inherent noise level of each sensor becomes part of the training and evaluation process. The present study does not incorporate an explicit noise model within

the learning architecture. Instead, reconstruction performance is assessed experimentally under controlled acquisition conditions, and the influence of thermal noise is interpreted through its effect on measurable image quality metrics.

RGB Imaging Sensor

RGB images acquired using two different imaging devices corresponding to the initial and final experimental setups. In the initial configuration, the RGB data used to build the dataset were captured using the rear camera of a Motorola Moto G84 smartphone. In the final experimental setup, RGB acquisition was performed using a USB camera module based on the Sony IMX323 image sensor.

Table 2.2 summarizes the main specifications of the RGB imaging sensors employed. The listed parameters include spatial resolution, sensor type, lens characteristics, and supported video formats, which define the spatial detail, field of view, and acquisition modes available during data collection.

Table 2.2: Main specifications of the RGB imaging sensors used in the experimental setups.

Device	Sensor / Camera	Resolution	Lens / FoV	Output Format
Motorola Moto G84 (Initial setup)	Smartphone rear camera	50 MP	f/1.88, wide-angle	RAW, JPEG, FHD video
ELP USB Camera (Final setup)	Sony IMX323 CMOS	1920 × 1080	100° fixed lens	H.264, UVC

The smartphone camera supports high-resolution still image acquisition and manual capture modes, including RAW output, which were used during the initial dataset generation.

The final experimental setup integrates a compact USB camera module based on the Sony IMX323 CMOS sensor, providing Full HD video acquisition at 30 fps with H.264 compression over a UVC interface. The fixed wide-angle lens ensures consistent spatial coverage across acquisitions.

2.4 Fundamentals of Hyperthermia

2.4.1 Physical Basis of Microwave-Tissue Interaction

Microwave hyperthermia is based on the interaction between electromagnetic fields and biological tissues, where a fraction of the absorbed electromagnetic energy is converted into heat

through dielectric mechanisms. When biological tissues are exposed to microwave radiation, the oscillating electric field induces polarization of dipolar molecules—predominantly water—as well as ionic conduction within the extracellular and intracellular fluids. These mechanisms produce energy dissipation at the molecular scale, generating heat through dielectric losses, a phenomenon commonly referred to as dielectric heating [14–16]. The sequence of processes linking electromagnetic exposure to temperature rise and thermal diffusion is schematically illustrated in Figure 2.4.

The propagation, absorption, and attenuation of electromagnetic energy within biological tissues are governed primarily by their dielectric properties, namely relative permittivity (ϵ) and electrical conductivity (σ). Tissues with high water content, such as muscle and skin, exhibit higher permittivity and conductivity, leading to more efficient microwave absorption and increased heat generation. In contrast, tissues such as adipose tissue and bone, which contain lower water content, generally exhibit reduced absorption and greater penetration depth, resulting in lower local temperature rise under similar exposure conditions [46,47]. Because each tissue type interacts differently with electromagnetic energy, temperature distributions within the body can become uneven [13,48]. To manage this variability and maintain both safety and treatment effectiveness, applicator design and control systems must be optimized accordingly.

The conversion of absorbed electromagnetic energy into heat and its subsequent redistribution within tissue are commonly described using bioheat transfer models. These models account for conductive heat transfer within tissue, convective heat exchange mediated by blood perfusion, and, to a lesser extent, metabolic heat generation [49]. The transient and steady-state temperature evolution depends on the tissue's thermal diffusivity, perfusion rate, and boundary conditions, which together determine the extent of thermal spread and the temporal stability of the heated region.

From a biological perspective, temperature thresholds play a critical role in defining the therapeutic and adverse effects of hyperthermia. Temperatures exceeding approximately 45 °C can cause protein denaturation and irreversible cellular damage, whereas maintaining tissue temperatures in the range of 40–43 °C is associated with therapeutic hyperthermia, promoting

increased blood perfusion, enhanced oxygenation, and improved radiosensitivity in tumor regions [?, 14, 47]. Precise temperature monitoring and control are therefore essential components of microwave hyperthermia systems.

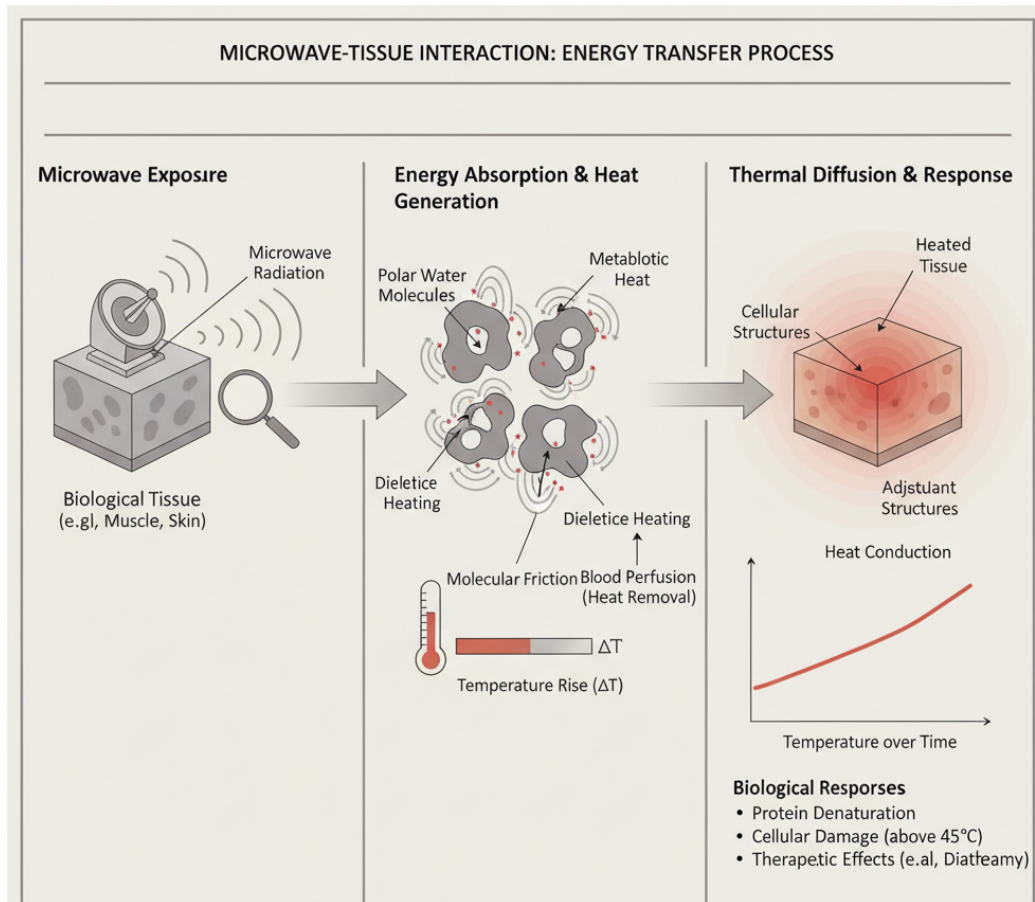


Figure 2.4: Schematic representation of the microwave–tissue interaction process. Microwaves induce dielectric heating by exciting polar molecules and ions, leading to temperature rise within biological structures. The resulting thermal diffusion and perfusion determine the spatial distribution and biological effects of heat. Image generated by AI using Google AI Studio.

2.4.2 Thermal Behavior of Biological Tissues

The thermal response of biological tissues during hyperthermia depends on their intrinsic physical and physiological parameters, such as density, specific heat, thermal conductivity, and blood perfusion rate. These properties influence both the rate of temperature rise and the heat dissipation capacity when tissues are exposed to external electromagnetic energy. Table 2.3 summarizes representative values of these parameters obtained from experimental and modeling studies [14–16]. Variations among tissues affect the spatial and temporal distribution

of temperature, highlighting the need for perfusion-dependent heat transfer models in localized hyperthermia treatments.

Table 2.3: Representative thermal properties of biological tissues at normothermic conditions, based on data from [14–16].

Tissue	Density (kg/m ³)	Specific heat (J/kg·K)	Conductivity (W/m·K)	Perfusion (kg/m ³ s)
Muscle	1090	3600	0.50	0.014
Skin	1109	3391	0.37	0.0026
Fat	911	2348	0.21	0.001
Tumor (Melanoma)	1045	3800	0.56	0.009
Blood	1050	3617	0.52	N/A

The higher heat capacity and conductivity of tissues with elevated water content, such as muscle or melanoma, facilitate rapid thermal equilibrium, whereas adipose layers, characterized by low conductivity and perfusion, act as thermal barriers, affecting the uniformity of temperature distribution in superficial hyperthermia.

2.4.3 Hyperthermia in Melanoma Treatment

The use of a hyperthermia as a therapeutic adjunct in melanoma has been explored for several years in preclinical studies seeking to understand how it interacts with radiotherapy and other cancer treatments. Early studies showed that combining radiation with hyperthermia could improve tumor response compared to radiotherapy alone [50]. Over time, different ways of applying heat were evaluated, and it was observed that hyperthermia helped to reduce the dose needed to achieve a complete response, especially when applied after radiotherapy, while maintaining safety in healthy tissues [51]. In a European multicenter study, applying hyperthermia at 43°C for 60 minutes after radiotherapy improved tumor control at two years and showed a relevant effect in statistical analyses, without increasing acute or late toxicities [52]. More extensive reviews have confirmed that hyperthermia can enhance other treatments such as radiotherapy and chemotherapy through mechanisms such as radiosensitization, changes in tumor perfusion, and temperature-dependent biological responses, which has driven the development of more precise tools for planning, controlling, and monitoring temperature during treatment [14,53]. In recent years, clinical analyses and technological advances have highlighted that localized hyperthermia, combined with radiotherapy and systemic therapies, can contribute

to the local control of melanoma and other neoplasms, emphasizing the importance of improving thermal dosimetry and moving toward more standardized protocols [47].

Chapter 3

Thermal and Electromagnetic Characterization of Tissue-Mimicking Phantoms

3.1 Phantom Design and Fabrication

The tissue-mimicking phantom developed in this work was designed to emulate the thermal and dielectric behavior of biological soft tissues under microwave and terahertz excitation. The phantom matrix was formulated using a polymer–saline solution composed of deionized water, sodium chloride, agar powder, and polyethylene powder, providing control over both thermal diffusivity and electromagnetic permittivity, as illustrated in Fig. 3.1.




Figure 3.1: Schematic representation of the tissue-mimicking phantom, including geometry and material composition used to emulate the thermal and dielectric properties of biological soft tissues.

The fabrication process consisted of dissolving sodium chloride in deionized water at 80 °C under continuous magnetic stirring. Agar powder was then gradually incorporated to achieve the desired gel viscosity, followed by the addition of polyethylene powder to adjust dielectric losses. The mixture was poured into cylindrical molds and left to cure at room temperature for 24 h to obtain homogeneous solid phantoms with stable mechanical and thermal properties.

The phantom formulation followed the procedure reported in [48], which was designed to reproduce the electrical properties of human muscle tissue at microwave frequencies. The mixture was prepared using deionized water as the base medium, with sodium chloride added to adjust ionic conductivity, agar powder acting as the gelling agent to stabilize the structure, and polyethylene powder to modify the dielectric response of the matrix.

The preparation procedure consisted of dissolving sodium chloride and sodium azide in deionized water under magnetic stirring with controlled heating. The solution temperature was increased to approximately 70°C to allow the incorporation of agar powder until complete dissolution. Subsequently, the magnetic stirrer was replaced with a mechanical mixing motor to homogeneously disperse the polyethylene powder together with the TX-151 thickening agent, which is part of the formulation reported in [48].

After mixing, the liquid preparation was poured into a mold and allowed to solidify. The resulting phantom was kept in a sealed container covered with plastic film for 24 hours to ensure structural stabilization. For storage, the sample was fully wrapped with plastic film and maintained at room temperature without direct sunlight exposure. The phantom used in this study was stored for approximately five months prior to characterization, during which gradual water loss produced minor variations in its electrical properties.

The formulation of the phantom follows the procedure reported in [48], where a polymer–saline mixture is used to approximate the thermal and dielectric behavior of biological soft tissues. In this type of formulation, deionized water acts as the base medium of the matrix, while sodium chloride is used to introduce ionic conductivity in the mixture. Agar powder provides the gelling structure that stabilizes the material once it solidifies, allowing the aqueous solution to maintain a fixed shape during experiments. Polyethylene powder is incorporated in the mixture to modify the effective dielectric response of the material. This combination of components

is commonly used in tissue-mimicking phantoms intended for microwave characterization, where both the electrical conductivity and the dielectric response of the material are adjusted to approach the behavior of hydrated biological tissues.

3.2 Chemical Characterization and Electromagnetic Analysis

3.2.1 Chemical Characterization

Chemical characterization of the tissue-mimicking phantom was performed using energy-dispersive X-ray spectroscopy (EDS) to evaluate the elemental composition and spatial distribution of the constituent materials. Figure 3.2 presents a representative elemental distribution map, revealing a predominantly carbon-based matrix with homogeneously distributed sodium and chloride.

The elemental composition obtained from the EDS analysis is summarized in Table 3.1, which reports the mean weight percentage (wt.%) and standard deviation for each detected element. Carbon constitutes the majority of the matrix, corresponding to the polymeric components of the phantom, namely agarose and polyethylene. Sodium and chloride originate from the dissolved sodium chloride added to adjust the ionic content of the material and, consequently, its dielectric conductivity in the microwave frequency range.

The observed standard deviations reflect microscale heterogeneity associated with manual mixing and gel casting processes. Localized variations in salt concentration and polymer network density contribute to spot-to-spot fluctuations in the measured elemental content. Elements with low atomic number, such as oxygen and hydrogen, are not detected by EDS; therefore, the remaining mass fraction is attributed primarily to oxygen in the agarose structure and residual moisture within the phantom.

EDS was used in this work to verify the elemental presence and spatial distribution of the main constituents of the phantom, particularly sodium and chlorine associated with the NaCl solution incorporated in the formulation. Although EDS is mainly sensitive to surface

composition and has limited sensitivity for light elements such as hydrogen and oxygen, it still provides useful qualitative information about the elemental distribution within the analyzed region.

In this context, the purpose of the analysis was to confirm the presence and distribution of the ionic components within the polymeric matrix rather than to obtain an exact volumetric quantification of the full composition. For this reason, EDS was considered adequate for the intended characterization. The analysis is therefore used as complementary information to support the consistency of the phantom formulation.




Figure 3.2: EDS elemental distribution map of the tissue-mimicking phantom, showing the spatial distribution of carbon, sodium, and chloride. Adapted from [13].

Table 3.1: Weight percentage (wt.%) of the elements identified in the phantom sample by EDS analysis. Adapted from [13].

Element	Mean (wt.%)	Deviation (wt.%)
Carbon	92.81	1.62
Sodium	1.61	0.46
Chloride	1.34	0.43

3.2.2 Thermal analysis

Thermogravimetric analysis (TGA) was conducted on a Discovery 550 Analyzer (TA Instruments) to assess the thermal stability of the muscle-equivalent phantom under hyperthermia-relevant conditions. Samples were heated in ambient air from room temperature up to 800 °C at a constant 10 °C/min, allowing us to pinpoint any mass-loss transitions and confirm that the phantom matrix remains structurally intact well above therapeutic temperature ranges [54]. In addition, to analyze the thermal behavior of the phantom at body temperature (37°C), it was heated at a rate of 1°C/min up to 37°C. Finally, an isothermal process was conducted for 70 min.

Thermogravimetric analysis (TGA) was carried out up to 800 °C to examine the thermal behavior and decomposition profile of the phantom material. Although hyperthermia treatments typically operate within the 37–45 °C range, extending the analysis to higher temperatures makes it possible to identify mass-loss events related to moisture evaporation and polymer degradation, which provides additional insight into the material composition.

Therefore, the objective of the TGA experiment was not to reproduce the therapeutic temperature range, but to obtain a broader characterization of the thermal stability of the phantom and confirm that no mass loss or structural degradation occurs within the temperature interval relevant for hyperthermia applications.

Furthermore, the phantom’s capacity to conduct and store heat was evaluated using a KD2 Pro Thermal Properties Analyzer (Decagon Devices, Inc.). The KD2 Pro features a dual-needle sensor (30 mm active length, 6 mm needle spacing) that captures both thermal conductivity and volumetric heat capacity simultaneously. To reduce ambient fluctuations, the sample was placed

inside a climatic chamber (Mettler, model HPP 110). The samples were kept at temperatures of 25 °C, 35 °C, and 45 °C, with a constant relative humidity of 65 %. The sensor was placed 3 cm into the phantom and allowed to equilibrate for 15 minutes. This time was enough to ensure that the temperature was uniform across the surrounding sample before taking a reading. In addition, each measurement was performed four times to verify accuracy and reproducibility.

The TGA and differential thermal analysis (DTG) of the phantom are shown in Fig. 3.3. The phantom was analyzed to evaluate its chemical composition (Fig.3.3a) and behavior at body temperature (Fig. 3.3b). The weight loss of (76.76%) between 25°C and 200°C is attributed to the evaporation of water from the phantom material [55]. The weight loss of 22.05% between 200°C and 800°C is associated with the combustion of polyethylene, agar, and TX-151 (Super Stuff) [56,57]. Finally, the ash content of 1.19% corresponds to the inorganic components such as sodium, which is confirmed by EDS.




Figure 3.3: Thermogravimetric analysis (TGA) of the tissue-mimicking phantom.

The muscle is the largest organ in the body, representing 40% of body weight, and is composed of approximately 76% water [58]. This water is useful for metabolic processes, metabolism, mechanical function, and other functions.

Fig. 3.4 shows the response of thermal conductivity and specific heat as a function of temperature variations. An approximate reduction of 57.4% is seen in the thermal conductivity of the phantom material, which drops from 0.588 to 0.25075 W/(m·K). The TGA analysis

indicates that the phantom contains a high water fraction (76.76 wt%), which strongly influences its thermophysical properties. In addition, the specific heat decreases with increasing temperature. An analogous pattern was observed in which the specific heat capacity reduces from 3.1935 MJ/(m³·K) to 1.664 MJ/(m³·K) with the rise in temperature from 25°C to 45°C. Water possesses a high specific heat capacity of 4.2 J/g K, allowing it to absorb thermal energy through hydrogen-bond interactions [59, 60].




Figure 3.4: Thermal properties of the phantom evaluated at different temperatures. (a) Thermal conductivity (b) Specific heat capacity.

Fig. 3.4 shows the response of thermal conductivity and specific heat as a function of temperature variations. An approximate reduction of 57.4% is seen in the thermal conductivity of the phantom material, which drops from 0.588 to 0.25075 W/(m·K). The TGA analysis indicates that the phantom contains a high water fraction (76.76 wt%), which strongly influences its thermophysical properties. In addition, the specific heat decreases with increasing temperature. An analogous pattern was observed in which the specific heat capacity reduces from 3.1935 MJ/(m³·K) to 1.664 MJ/(m³·K) with the rise in temperature from 25°C to 45°C. Water possesses a high specific heat capacity of 4.2 J/g K, allowing it to absorb thermal energy through hydrogen-bond interactions [59, 60].

3.2.3 Optical Response

Optical characterization of the tissue-mimicking phantom across the visible and near-infrared (VIS–NIR) spectral range was carried out using the experimental setup illustrated in Figure 3.5.

The setup was based on a laser-driven plasma white light source (Energetiq EQ-99-FC, Wilmington, MA, USA), which provides a broadband and spectrally stable emission spanning from 190 nm to 2500 nm, ensuring sufficient optical power and spectral continuity over the wavelength range of interest. The experimental setup was based on a laser-driven plasma white-light source (Energetiq EQ-99-FC, Wilmington, MA, USA), providing broadband and spectrally stable emission from 190 to 2500 nm. The reflected optical signal was analyzed using an optical spectrum analyzer (OSA) (Yokogawa AQ6373, Tokyo, Japan). Both the light source and the OSA were coupled to the measurement system through a bifurcated fiber-optic probe (Ocean Optics, QR200-7-UV-VIS). The probe was positioned at an incidence angle of 45° relative to the phantom surface (see Fig. 3.5). This angular configuration was selected to minimize specular reflections and prevent direct back-reflection of the incident beam into the collection fiber [61].

In addition, spectral measurements were conducted using a linear wavelength scale with a nominal resolution of 5 nm, covering the spectral range from 350 nm to 1200 nm with a wavelength increment of 1 nm. Before phantom measurements, the system was calibrated using a certified diffuse reflectance standard (USRS-99-010, Labsphere, Palo Alto, CA, USA). Following calibration, diffuse reflectance spectra of the tissue-mimicking phantom samples were acquired under identical experimental conditions.




Figure 3.5: Experimental setup used for VIS–NIR diffuse reflectance measurements of the tissue-mimicking phantom.

Figure 3.6 shows the VIS–NIR spectral response of the phantom. The mean normalized reflectance spectrum (Fig. 3.6a) exhibits a decreasing trend from the near-ultraviolet toward the visible region, followed by a gradual decay toward longer wavelengths. A local reflectance maximum is observed around 590–610 nm. The corresponding apparent absorbance spectra derived from reflectance (Fig. 3.6b) show a progressive increase toward the NIR region, consistent with absorption mechanisms associated with hydrated polymer-based media. The shaded regions indicate the standard deviation across measurement positions.




Figure 3.6: VIS–NIR optical characterization of the tissue-mimicking phantom. (a) Mean normalized reflectance spectrum with standard deviation. (b) Apparent absorbance spectrum derived from reflectance. (c) Wavelength-dependent coefficient of variation of absorbance.

Spatial uniformity was quantified by computing the wavelength-dependent coefficient of

variation (CV) of the apparent absorbance, shown in Fig. 3.6c. Over most of the spectral range, the CV remains below 15%, while higher values appear near the upper wavelength limit, which is attributed to reduced signal-to-noise ratio and increased absorption in water-rich media.

To provide a band-wise quantitative comparison, the apparent absorbance was integrated over selected VIS–NIR spectral intervals. The resulting values are summarized in Table 3.2. The visible range (400–700 nm) presents the highest integrated absorbance, followed by the near-infrared bands. Similar integrated values are obtained for the 900–1000 nm and 1000–1100 nm intervals, indicating a smooth absorption baseline across the near-infrared region.

Table 3.2: Integrated apparent absorbance of the tissue-mimicking phantom over selected VIS–NIR spectral bands. Values are reported as mean \pm standard deviation across measurement positions.

Spectral band	Wavelength range (nm)	Integrated absorbance (a.u.)	CV (%)
VIS	400–700	368.56 \pm 39.10	10.61
NIR-I	700–900	265.91 \pm 30.38	11.42
Water band	900–1000	140.52 \pm 17.71	12.60
NIR-II	1000–1100	142.06 \pm 19.78	13.93

Principal component analysis (PCA) was applied to the apparent absorbance spectra to assess spectral reproducibility across measurement positions. Figure 3.7 shows the PCA score plot, where the first principal component accounts for most of the total variance, indicating that inter-point variability is dominated by global intensity-related contributions. A single measurement position is separated from the main cluster, while the remaining points form a compact distribution.




Figure 3.7: Principal component analysis of the VIS–NIR apparent absorbance spectra measured at multiple positions on the phantom surface.

3.3 Experimental Setup and Reproducibility

Using the experimental configuration described in the previous section, the thermal response of the tissue-mimicking phantom was characterized under controlled microwave excitation conditions.

At each acquisition time instant, multiple consecutive thermal frames were recorded to evaluate temperature dispersion and temporal stability. For the low-power condition, acquisitions were performed up to 390 min, while for the high-power condition the window was limited to 120 min.

The acquisition windows were selected considering both the thermal response of the phantom and the objective of evaluating its stability under prolonged microwave exposure. The 120 min duration used for the 20 dBm condition was chosen to reproduce temperature ranges typically associated with hyperthermia experiments, where temperatures between 40–45 °C are reached within relatively short exposure periods. In contrast, the longer acquisition window of 390 min for the 10 dBm condition was used to evaluate the long-term thermal behavior of the phantom

under low-power excitation. This extended observation period allowed assessing the gradual temperature stabilization and identifying potential changes associated with moisture loss in the material during prolonged heating.

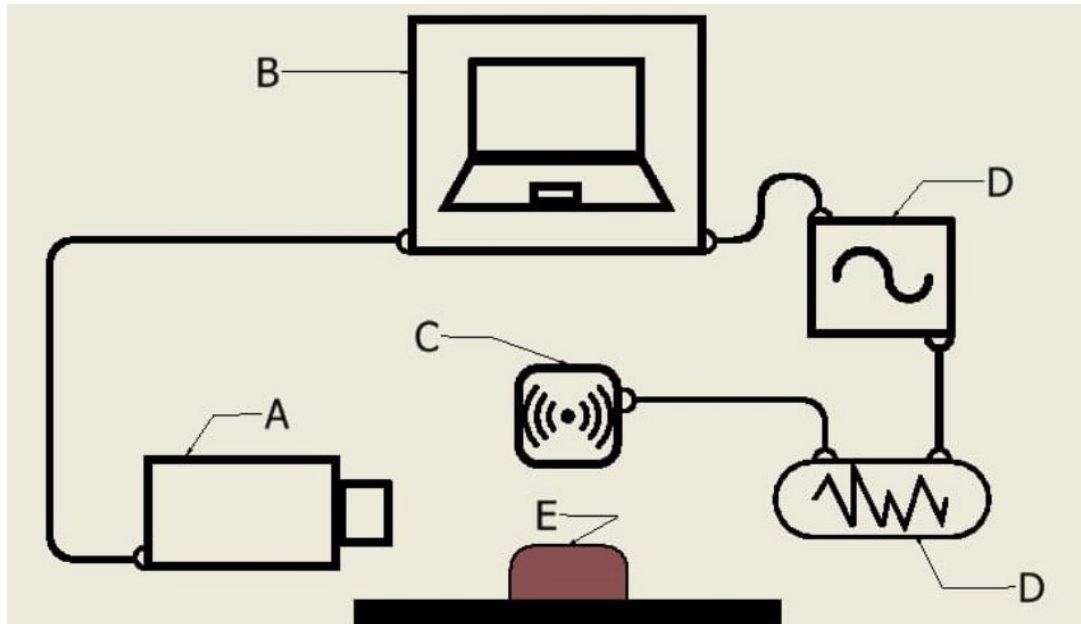


Figure 3.8: Experimental setup for microwave heating and thermal acquisition.

3.4 Thermal Characterization under Microwave Excitation

The thermal response of the tissue-mimicking phantom was evaluated under controlled microwave excitation using a software-defined radio (SDR) platform (USRP) operating at a carrier frequency of 2.45 GHz. A continuous-wave complex sinusoidal signal with a baseband frequency of 100 kHz and a sampling rate of 1 MS/s was generated and transmitted through a directive antenna oriented toward the phantom surface. The RF signal was amplified using an external power amplifier before radiation.

Thermal emission from the phantom surface was recorded using an Optris IR camera positioned at a fixed distance of 20 cm. A constant region of interest was defined and maintained throughout all measurements to ensure spatial consistency. Ambient temperature remained within the range of 21.3–22.4 °C during all experiments. Two excitation power levels were analyzed: 10 dBm and 20 dBm. For each condition, IR images were acquired at predefined time intervals in order to characterize the temporal evolution of the surface temperature.

Representative infrared images acquired at different exposure times are shown in Fig. 3.9. These images illustrate the spatial distribution of surface temperature at selected instants during the heating process under microwave excitation.

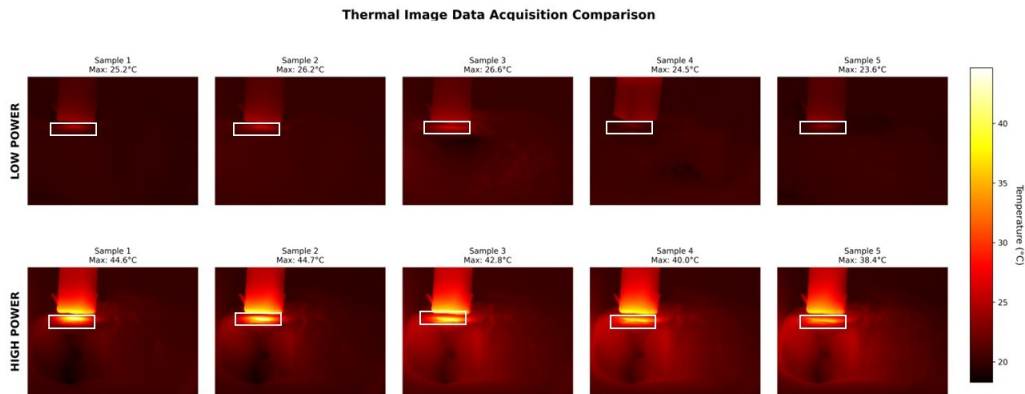


Figure 3.9: Representative infrared images of the tissue-mimicking phantom acquired at different exposure times under microwave excitation.

The temporal evolution of the surface temperature under 20 dBm excitation is presented in Fig. 3.10. The temperature increases rapidly during the initial phase of excitation, reaching approximately 45 °C at 15 min, followed by a gradual stabilization within the range of 38–45 °C over a total exposure time of 120 min.




Figure 3.10: Temporal evolution of phantom surface temperature under 20 dBm microwave excitation.

For the 10 dBm condition, the surface temperature exhibits a slower increase, rising from 23.8 °C to a maximum value of 26.6 °C at 240 min, followed by a slight decrease to 24.5 °C at 390 min, as shown in Fig. 3.11. This behavior indicates a stable thermal response under low-power excitation over extended acquisition times.




Figure 3.11: Temporal evolution of phantom surface temperature under 10 dBm microwave excitation.

Figure 3.8 corresponds to the USRP-based RF excitation platform used for system characterization and thermal response modeling. All power levels reported in this section (10–20 dBm) refer exclusively to this configuration.

3.5 THz Electromagnetic Characterization

Terahertz time-domain spectroscopy (THz-TDS) measurements were performed to provide a complementary electromagnetic characterization of the skin-mimicking phantom samples and to evaluate their temporal stability. The measurements were conducted in reflection mode using a benchtop THz spectrometer, following a consistent acquisition protocol based on a metallic mirror reference and sequential measurements of the phantom samples under unchanged experimental conditions [62, 63]. A schematic of the reflection configuration is shown in Fig. 3.12.




Figure 3.12: Schematic of the SHR reflection configuration, showing the THz beam path from the Tx antenna to the sample and back to the Rx antenna.

Two phantom slabs with identical lateral dimensions and thicknesses of 0.5 cm and 0.3 cm were analyzed. For each sample, repeated measurements were acquired at regular time intervals over several hours in order to assess time-dependent changes in the THz response. All acquisition parameters were kept constant to ensure that the observed variations could be attributed to the sample rather than to instrumental drift. Figures 3.13 and 3.14 summarize the time- and frequency-domain responses obtained for the 0.5 cm and 0.3 cm thick phantoms, respectively. In the time domain (subfigures (a)), a clear contrast is observed between the reference signal and the phantom response. The mirror reference exhibits a dominant, high-amplitude reflected pulse, whereas the phantom signals present lower amplitudes and smoother waveforms, consistent with increased absorption and scattering in the water-rich material.

In addition to amplitude attenuation, the phantom responses show a progressive temporal shift of the reflected pulse as the measurements advance. This effect is clearly visible in the zoomed views of the time-domain traces and indicates a cumulative change in the effective optical path length during the acquisition period. For the 0.5 cm thick phantom (Figure 3.13(a)), the reflected pulse remains within the fixed analysis window throughout the measurement sequence, allowing a consistent temporal analysis. In contrast, for the 0.3 cm thick phantom (Figure 3.14(a)), the pulse displacement occurs more rapidly, and after a certain time the main reflection moves outside the analysis window, limiting the range over which quantitative comparisons remain valid.

The corresponding frequency-domain spectra (subfigures (b)) show that the phantom response is mainly concentrated below approximately 2 THz and exhibits significantly lower

magnitude than the mirror reference. Although the overall spectral shape remains similar across measurements, a gradual reduction in spectral amplitude is observed over time, particularly at lower frequencies, in agreement with the time-domain observations.

On the other hand, amplitude-based metrics computed within a fixed temporal window (subfigures (c)) further confirm this behavior. For the 0.5 cm phantom, both peak-to-peak and RMS metrics evolve smoothly over time, indicating systematic changes rather than random fluctuations. For the 0.3 cm phantom, a more pronounced decrease is observed within the initial measurement interval, consistent with the faster temporal drift seen in the time-domain signals.

Finally, the temporal shift of the reflected pulse was quantified as an equivalent optical path-length change (subfigures (d)). The thicker phantom exhibits a moderate, approximately linear drift over time, whereas the thinner phantom shows a larger drift rate under identical conditions. This thickness-dependent behavior suggests that the phantom–measurement system undergoes a measurable temporal evolution in the THz regime, occurring on a shorter timescale for thinner samples. This THz characterization confirms that the phantom exhibits time-dependent electromagnetic behavior that should be considered when it is used as a reference material in thermal and microwave experiments [63]. In the context of this thesis, the THz analysis serves to support the experimental characterization of the phantom and to delimit its stability over time, without constituting the primary focus of the study.

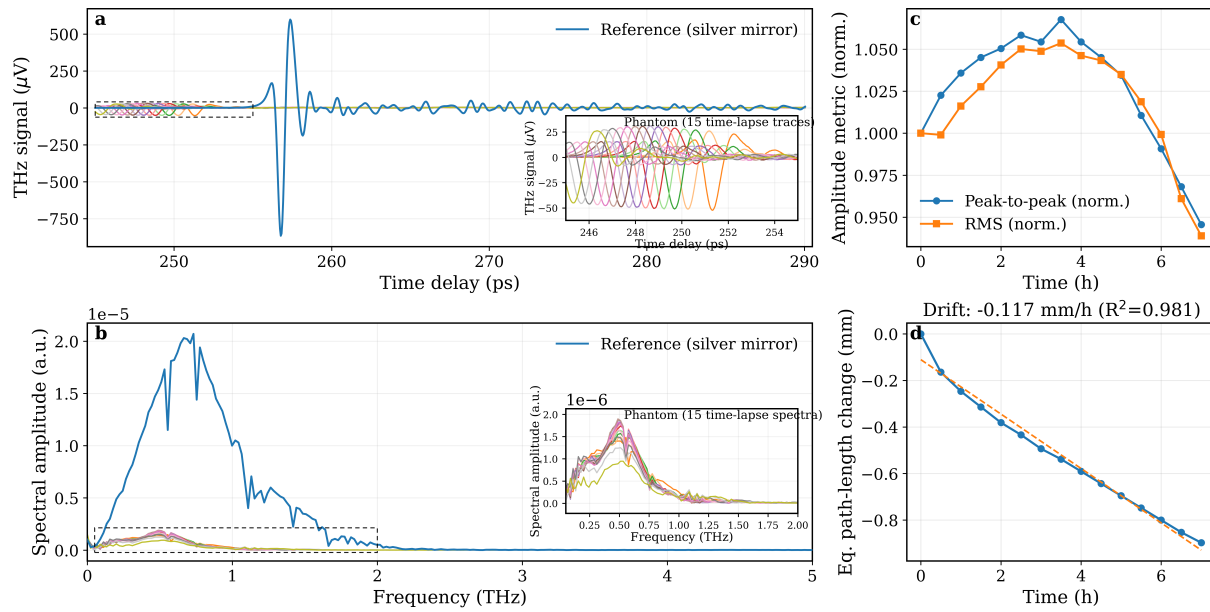


Figure 3.13: THz-TDS of the phantom (0.5 cm thickness). (a) Time-domain traces: reference (mirror) and 15 sequential acquisitions of the phantom, with a zoomed view of the window of interest. (b) Corresponding frequency-domain spectra, with a magnified view in 0.05–2 THz. (c) Normalized amplitude metrics (peak-to-peak and RMS) computed within the same time window. (d) Drift expressed as an equivalent optical path-length change estimated from the pulse shift and linear regression.

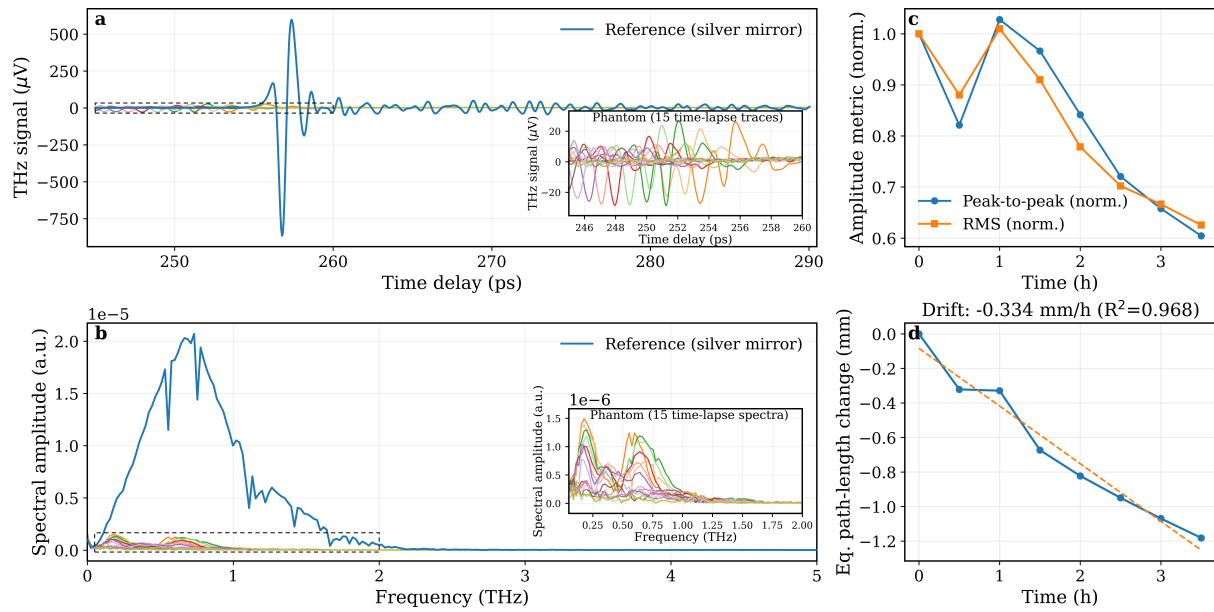


Figure 3.14: THz-TDS of the phantom (0.3 cm thickness). (a)–(b) Same as Fig. 3.13 for 15 sequential acquisitions. (c)–(d) Amplitude metrics and drift computed over the initial subset of measurements for which the pulse remains within the fixed temporal window used in the analysis.

Chapter 4

Construction and Labeling of the Database

4.1 Radiometric Validation

The calibration strategy was structured in two experimental stages, each addressing a specific requirement for the construction of reliable thermal image pairs for super-resolution analysis. The first stage focused on evaluating the inter-sensor consistency between thermal cameras of different resolution levels, with the objective of quantifying measurement dispersion and relative deviations under controlled conditions. The second stage extended this analysis by incorporating the effect of acquisition distance on radiometric accuracy, in order to identify the optimal working range that minimizes measurement variability and improves temperature reliability across sensors. The experimental setup, illustrated in Fig. 4.1, comprised an Ametek Jofra CTC-650A dry-block temperature calibrator used as the thermal reference source, the thermal cameras under evaluation, and a personal computer acting as host and data acquisition unit for both sensors.




Figure 4.1: Experimental setup used for the calibration procedure: (A) Ametek Jofra CTC–650A dry–block temperature calibrator, (B) thermal cameras under evaluation, and (C) host computer for data acquisition and control.

Initial Inter–Sensor Consistency Assessment

The first calibration experiment was designed to evaluate the consistency between a low–resolution thermal sensor (Thermal Camera HAT) and a medium–resolution sensor (Optris PI 450i). Both cameras were employed to acquire radiometric temperature measurements from a controlled thermal reference. At this stage, only temperature values were analyzed, since the objective was to assess the dispersion and relative deviation between sensors with different performance levels.

Table 4.1: Inter-sensor measurements obtained with the Optris and the Thermal Camera HAT (Waveshare) for each reference temperature. Each value corresponds to the ROI-averaged temperature extracted from five independent measurement sweeps.

T_{ref}	Camera	D1	D2	D3	D4	D5
20	Optris	20.01	20.22	20.19	20.50	20.80
	Waveshare	20.20	20.60	20.22	20.31	30.65
25	Optris	24.82	25.01	25.12	20.34	24.98
	Waveshare	25.80	25.68	25.55	25.75	25.50
30	Optris	30.20	30.12	30.15	30.22	30.28
	Waveshare	29.80	29.96	30.08	30.22	30.11
35	Optris	35.26	35.18	35.08	35.17	34.98
	Waveshare	34.95	34.87	35.52	35.12	35.28
40	Optris	40.26	40.18	40.33	40.22	40.28
	Waveshare	40.18	40.01	40.08	40.15	40.20

It is important to note that two isolated outliers are observed in Table 4.1, namely the 30.65 °C measurement obtained with the Waveshare sensor at $T_{\text{ref}} = 20$ °C and the 20.34 °C value recorded by the Optris sensor at $T_{\text{ref}} = 25$ °C. These values deviate from the surrounding measurements within the same sweep and were identified as isolated events rather than systematic sensor bias. Given that each reference temperature was measured through five independent sweeps, the remaining repetitions exhibit consistent behavior around the nominal reference value. Since the primary operating range of the proposed hyperthermia system is between 30–40 °C, these low-temperature outliers do not influence the main performance analysis. Nevertheless, they are explicitly reported to ensure transparency in the inter-sensor consistency evaluation.

For each reference temperature, five independent measurement sweeps were performed in randomized order in order to minimize biases associated with thermal transients. During each sweep, the mean temperature was extracted from a fixed region of interest defined on the calibration target. The raw values are summarized in Table 4.1. From these data, descriptive statistics were computed to evaluate both intra-sensor repeatability and inter-sensor dispersion between the Thermal Camera HAT and the Optris camera.

As shown in Fig. 4.2, both cameras exhibit an approximately linear relationship with the reference temperature in the range of 30–40 °C, with deviations below 0.3 °C for most of the evaluated points.

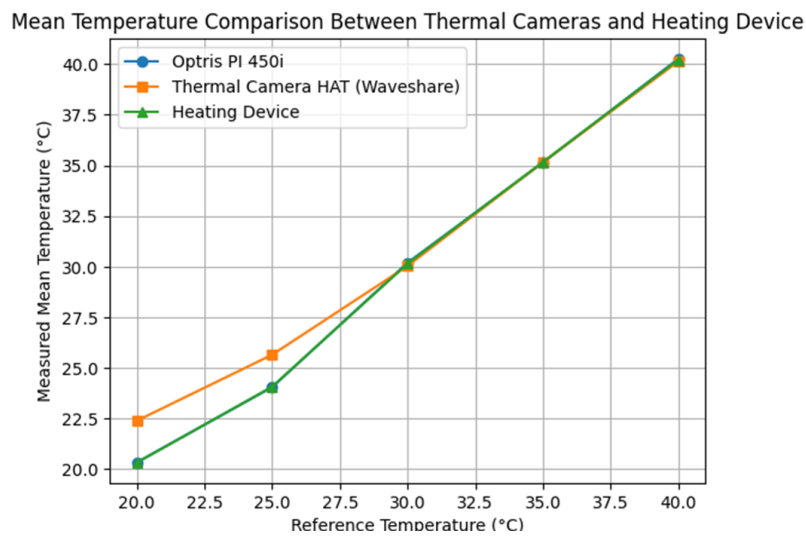


Figure 4.2: Mean temperature measured by the Optris camera and the Thermal Camera HAT as a function of the reference temperature.

Figure 4.3 reveals a marked increase in the dispersion of the Waveshare camera at 20 °C, whereas for temperatures above 30 °C both sensors present narrow and highly reproducible distributions.

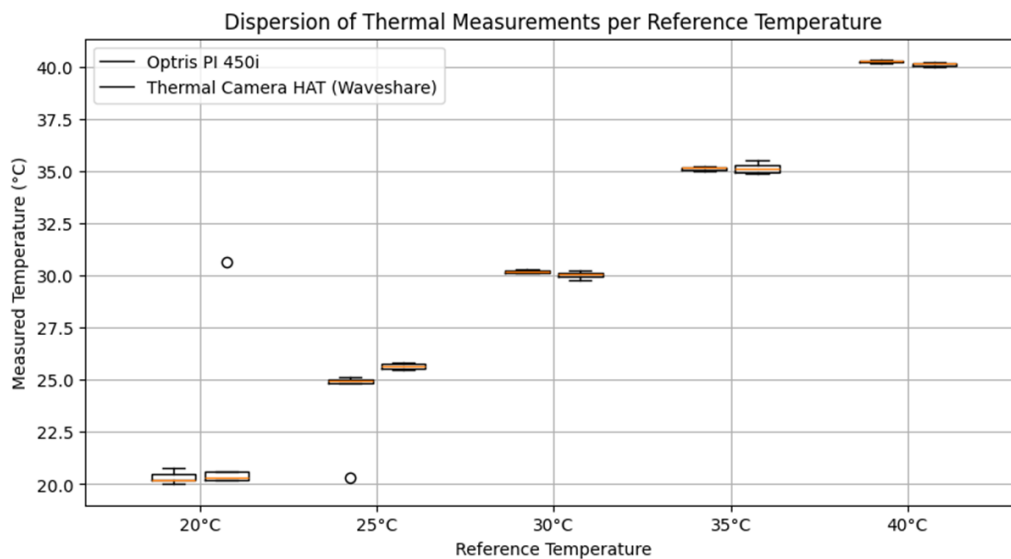


Figure 4.3: Boxplots of the thermal measurements obtained with the Optris camera and the Thermal Camera HAT for each reference temperature.

Table 4.2: Descriptive statistics of the inter-sensor measurements for the Optris and the Thermal Camera HAT (Waveshare).

T_{ref} (°C)	Camera	Mean (°C)	σ (°C)	CV (%)
20	Optris	20.344	0.309	1.52
	Waveshare	22.396	4.617	20.61
25	Optris	24.054	2.079	8.64
	Waveshare	25.656	0.128	0.50
30	Optris	30.194	0.062	0.21
	Waveshare	30.034	0.160	0.53
35	Optris	35.134	0.107	0.30
	Waveshare	35.148	0.261	0.74
40	Optris	40.254	0.057	0.14
	Waveshare	40.124	0.078	0.19

The results of the inter-sensor evaluation are summarized in Table 4.2, where the mean, standard deviation, and coefficient of variation (CV) are reported for each sensor and reference temperature. For temperatures of 30, 35, and 40 °C, both cameras exhibit CV values below 1%, with minimum values of 0.14% for the Optris at 40 °C and 0.19% for the Waveshare, indicating high stability and repeatability within the thermal range of interest for hyperthermia applications. In contrast, at 20 °C the Waveshare camera shows a pronounced increase in dispersion, reaching a CV of 20.6%, whereas the Optris remains close to 1.5%. This behavior at low temperatures is mainly attributed to instabilities in the heating system used as thermal reference, which exhibited difficulties in maintaining the setpoint in the lower regime, generating thermal transients that more severely affect the lower-quality radiometric sensor. Given that within the main operational range of the study (30–40 °C) both cameras exhibit low dispersion and consistent behavior, no additional corrective calibration between sensors was deemed necessary, and the radiometric values were directly used for the construction of the thermal super-resolution image pairs.

Distance–dependent calibration and precision analysis

Based on the results obtained in the initial inter–sensor consistency assessment, a second calibration stage was designed to analyze the effect of the camera–to–target distance on the accuracy of thermal measurements. This experiment was conducted using the same experimental setup described in Fig. 4.1, ensuring consistent conditions in terms of the thermal reference source, camera positioning, and data acquisition procedure.

This stage was carried out exclusively with the Optris and FLIR cameras, since the previous experiment showed that the Thermal Camera HAT exhibits higher dispersion and reduced radiometric stability, particularly outside its nominal operating range. The device technical documentation and the outcomes of the inter–sensor analysis enabled a more detailed characterization of its image quality limitations, which motivated its exclusion from this phase. Consequently, the objective of this experiment was to identify the distance range in which the medium– and high–resolution sensors present lower measurement variability and more consistent behavior, considering that both devices allow manual adjustment of their acquisition parameters.

Thermal images were acquired at controlled reference temperatures while varying the working distance between 5 cm, 10 cm, and 15 cm. For each configuration, multiple independent measurements were performed following the same procedure for temperature extraction from a fixed region of interest described in the previous section. Unlike the Thermal Camera HAT, the Optris and FLIR cameras provide acquisition interfaces that allow the definition of regions of interest directly on the thermal images, which facilitated a geometrically consistent selection of the analysis area and reduced the errors associated with the manual delineation applied in the previous experiment. The values obtained for each measurement round, reference temperature, and acquisition distance are summarized in Table 4.3.

Table 4.3: Thermal measurements obtained at different acquisition distances using the Optris and FLIR cameras. Each value corresponds to the average temperature extracted from the ROI for each measurement round.

Round	T_{ref} (°C)	Optris (°C)			FLIR (°C)		
		5 cm	10 cm	15 cm	5 cm	10 cm	15 cm
1	35	33.66	34.14	34.88	37.77	35.70	35.40
	40	39.57	39.72	40.23	42.60	41.50	39.10
	45	44.15	44.40	44.80	48.80	43.80	46.20
	50	49.51	48.80	49.68	53.80	51.90	49.80
2	35	34.58	33.63	35.90	37.15	35.23	36.08
	40	40.58	38.91	40.80	42.17	41.05	39.75
	45	42.85	43.73	43.99	49.55	44.68	45.63
	50	50.91	49.99	50.96	52.85	50.93	50.50
3	35	34.61	33.31	35.58	38.88	36.72	36.35
	40	40.62	40.72	41.37	41.74	40.36	39.70
	45	43.02	43.33	43.65	47.46	44.58	47.45
	50	50.17	49.71	48.84	55.34	52.92	50.97
4	35	34.55	34.82	35.47	38.60	35.07	34.59
	40	38.72	40.45	39.26	43.66	42.33	38.17
	45	43.59	44.90	45.26	48.16	42.50	47.24
	50	50.84	47.57	48.99	55.06	50.92	48.86
5	35	32.99	35.12	34.01	37.19	35.15	36.21
	40	39.11	38.63	39.80	43.58	42.29	39.59
	45	45.45	45.15	43.94	47.93	44.47	46.79
	50	48.14	49.50	50.23	54.72	53.24	50.84

For each distance and temperature condition, five independent measurement scans were recorded following the same randomized protocol used in the previous experiment. The extracted temperature values were subsequently analyzed in terms of dispersion and deviation with respect to the imposed setpoint. To illustrate the type of radiometric information acquired, Figs. 4.4 and 4.5 present representative examples of the thermal mosaics obtained with the Optris and FLIR cameras, respectively, for different temperature levels and acquisition distances. These mosaics provide a visual representation of the spatial variations in the thermal distribution and highlight the effect of the working distance on the definition of the region of interest and on the stability of the radiometric response.




Figure 4.4: Thermal image mosaic acquired with the Optris camera at different reference temperatures and working distances.




Figure 4.5: Thermal image mosaic acquired with the FLIR camera at different reference temperatures and working distances.

The quantitative results of the dispersion for each combination of camera, acquisition distance, and reference temperature are summarised in Table 4.4. In general, measurements

acquired at a distance of 5 cm exhibit the highest standard deviation values, with coefficients of variation ranging between 1.9% and 2.4% for both cameras, indicating increased radiometric instability under this configuration. This behaviour is attributed to geometric and optical effects associated with the extreme proximity of the sensor to the thermal target, as well as to the increased sensitivity to small variations in focus, alignment, and local thermal gradients. In contrast, acquisitions performed at 10 cm and 15 cm show a consistent reduction in dispersion, with coefficients of variation typically below 2%, which reflects improved repeatability and measurement stability within this operational range.

Table 4.4: Descriptive statistics of the thermal measurements obtained with the Optris and FLIR cameras at different acquisition distances.

Setpoint (°C)	Camera	Distance	Mean (°C)	σ (°C)	CV (%)
35	Optris	5 cm	34.078	0.727	2.13
35	Optris	10 cm	34.204	0.767	2.24
35	Optris	15 cm	35.168	0.745	2.12
35	FLIR	5 cm	37.918	0.796	2.10
35	FLIR	10 cm	35.574	0.686	1.93
35	FLIR	15 cm	35.726	0.732	2.05
40	Optris	5 cm	39.720	0.858	2.16
40	Optris	10 cm	39.686	0.918	2.31
40	Optris	15 cm	40.292	0.826	2.05
40	FLIR	5 cm	42.750	0.851	1.99
40	FLIR	10 cm	41.506	0.839	2.02
40	FLIR	15 cm	39.262	0.663	1.69
45	Optris	5 cm	43.812	1.049	2.39
45	Optris	10 cm	44.302	0.768	1.73
45	Optris	15 cm	44.328	0.674	1.52
45	FLIR	5 cm	48.380	0.813	1.68
45	FLIR	10 cm	44.006	0.910	2.07
45	FLIR	15 cm	46.662	0.750	1.61
50	Optris	5 cm	49.914	1.143	2.29
50	Optris	10 cm	49.114	0.969	1.97
50	Optris	15 cm	49.740	0.881	1.77
50	FLIR	5 cm	54.354	1.021	1.88
50	FLIR	10 cm	51.982	1.084	2.09
50	FLIR	15 cm	50.194	0.873	1.74

Additionally, the FLIR camera exhibits larger systematic deviations at 5 cm, particularly at higher setpoints, whereas the Optris camera maintains a more homogeneous variability across all

evaluated distances. Nevertheless, within the 10–15 cm range both cameras present comparable and sufficiently low dispersion levels for the hyperthermia applications considered in this work.

Although the reduction in the coefficient of variation indicates improved repeatability at 10–15 cm, the selection of the operational distance was not based solely on dispersion metrics. For each camera, the absolute bias and the RMSE relative to the primary thermal reference T_{ref} were also evaluated in order to jointly consider precision and accuracy. The 10 cm configuration provides a simultaneous reduction in variability and absolute deviation for both sensors when compared to the 5 cm condition. Therefore, the working distance was selected by minimizing the combined radiometric error with respect to the reference standard rather than relying exclusively on the coefficient of variation.

These results support the conclusion that no additional intersensor corrective calibration procedures are required. Instead, it is sufficient to explicitly acknowledge the magnitude of the error associated with the geometric capture configuration and to restrict thermal data acquisition to distances equal to or greater than 10 cm. This condition was adopted as an experimental criterion in all subsequent stages, as it ensures an objective reduction in dispersion and improved consistency of the radiometric data employed in the thermal super-resolution processing pipeline.

It is important to clarify that, in this stage, neither the Optris nor the FLIR camera was used as a reference standard for the other. The dry-block calibrator remained the primary thermal reference throughout the entire experimental procedure. Both cameras were evaluated independently but under identical controlled conditions in order to analyze the influence of acquisition distance on measurement stability.

The objective of this experiment was not to establish an intersensor hierarchy, but to determine the distance range in which each device exhibits reduced dispersion and acceptable deviation relative to T_{ref} . Therefore, the comparison presented is parallel rather than hierarchical, and all performance interpretations are implicitly grounded on the imposed reference temperature.

4.1.1 Inter-Sensor Comparative Analysis

Based on the calibration experiments described in the previous sections, an additional comparative analysis was performed to explicitly evaluate the agreement between thermal sensors operating at different spatial and radiometric resolutions. This analysis aims to consolidate the statistical information obtained from repeated measurements and to assess whether an explicit inter-sensor calibration procedure is required prior to database construction.

Figure 4.6 presents the cross-comparison between the Optris camera and the Thermal Camera HAT for the reference temperatures considered in the initial calibration stage. Each data point corresponds to the mean temperature obtained from five independent measurement sweeps at a given setpoint, while the horizontal and vertical error bars represent the standard deviation of the measurements acquired by each sensor, respectively.




Figure 4.6: Cross-comparison between Optris and Thermal Camera HAT measurements. Markers indicate mean values across five repeated measurements per setpoint, while error bars represent the corresponding standard deviation for each sensor. The solid line corresponds to the linear regression fit, and the dashed line indicates the ideal identity relation.

The dispersion represented by the error bars provides a measure of the repeatability of each sensor under controlled thermal conditions. Most data points are located close to the identity line, indicating a linear correspondence between the temperatures measured by both devices over the evaluated range.

A linear regression analysis was performed using the mean temperature values obtained at each reference setpoint. The resulting coefficient of determination was $R^2 = 0.978$, showing that the Thermal Camera HAT measurements exhibit a strong linear dependence on the Optris reference values. It is important to emphasize that the regression analysis was performed using the averaged temperature values at each reference setpoint (five data points), ensuring methodological consistency between the graphical representation and the statistical model.

The inter-sensor deviation was further quantified using the root mean square error (RMSE),

which yielded a value of 1.02 °C. For consistency with the regression analysis, the RMSE was computed using the mean temperature obtained at each reference setpoint rather than the 25 individual measurements. Recalculation using the reported mean values yields an RMSE of approximately 1.17 °C, with minor differences attributable to rounding effects. If all individual measurements—including the two isolated outliers identified previously—are considered, the RMSE increases to approximately 2.27 °C. However, such computation would overweight repeated measurements and artificially amplify the influence of transient deviations. Therefore, the RMSE derived from averaged setpoint values is retained as the most statistically coherent indicator of inter-sensor agreement. This metric represents the typical magnitude of the difference between both sensors across all measurement points.

The systematic offset between sensors was evaluated through the bias, defined as the mean difference between the Thermal Camera HAT and Optris measurements,

$$\text{Bias} = \frac{1}{N} \sum_{i=1}^N (T_{\text{HAT},i} - T_{\text{Optris},i}) .$$

The computed bias was 0.29 °C, indicating a small and approximately constant offset, with no evidence of temperature–dependent deviation within the analyzed range.

Figure 4.7 presents the cross–comparison between the Optris and FLIR cameras for measurements acquired at a working distance of 10 cm. For each reference temperature, the plotted points correspond to the mean value obtained from five independent measurement sweeps, while the horizontal and vertical error bars represent the corresponding standard deviations of the Optris and FLIR measurements, respectively. The use of averaged values ensures a comparable representation between sensors by reducing the influence of random fluctuations and short–term thermal transients, while the dispersion information is explicitly preserved through the error bars.




Figure 4.7: Cross-comparison between Optris and FLIR temperature measurements acquired at a working distance of 10 cm. Each point represents the mean temperature obtained from five independent measurement sweeps for a given setpoint. Horizontal and vertical error bars indicate the corresponding standard deviations of the Optris and FLIR measurements, respectively. The solid line represents the linear regression between both sensors.

The proximity of the data points to the identity line indicates a consistent linear relationship between both cameras over the evaluated temperature range. The coefficient of determination ($R^2 = 0.923$) confirms that most of the variability in the FLIR measurements can be explained by the corresponding Optris values. The root mean square error ($RMSE = 1.65\text{ °C}$) quantifies the overall deviation between sensors, while the bias (1.44 °C), computed as the mean difference between FLIR and Optris measurements, indicates a systematic overestimation of temperature by the FLIR camera under the analyzed conditions. Recalculation of the RMSE using the mean temperature values reported for each reference setpoint at 10 cm yields a value of approximately 1.84 °C . The slight discrepancy with respect to the initially reported value is explained by the use of non-rounded internal data during the original regression analysis. As in the previous comparison, the error metric was computed using averaged values to preserve coherence between

regression, bias, and dispersion indicators.

In order to provide an integrated assessment of the statistical behavior of the evaluated thermal cameras, a normalized radar chart was constructed using metrics derived from the calibration experiments. It is emphasized that this analysis does not aim to perform an inter-sensor calibration or to apply corrective transformations to the measurements. Instead, the objective is to compare the statistical performance of each device under controlled and repeatable experimental conditions, considering accuracy, dispersion, and consistency with respect to the imposed thermal reference.

The radar chart, shown in Fig. 4.8, summarizes five complementary indicators: the coefficient of determination (R^2), the root mean square error (RMSE), the bias, the coefficient of variation (CV), and the standard deviation.

Each metric was normalized using a global min–max scaling across all evaluated cameras. For a given metric, the minimum and maximum values observed among the three sensors were used to rescale the data to the interval [0,1]. This normalization was performed per metric (not per camera) in order to preserve the comparative interpretation of the chart. For those indicators where lower values correspond to better performance (RMSE, absolute bias, CV, and standard deviation), the normalized values were inverted so that higher values in the radar representation consistently indicate improved statistical behavior.

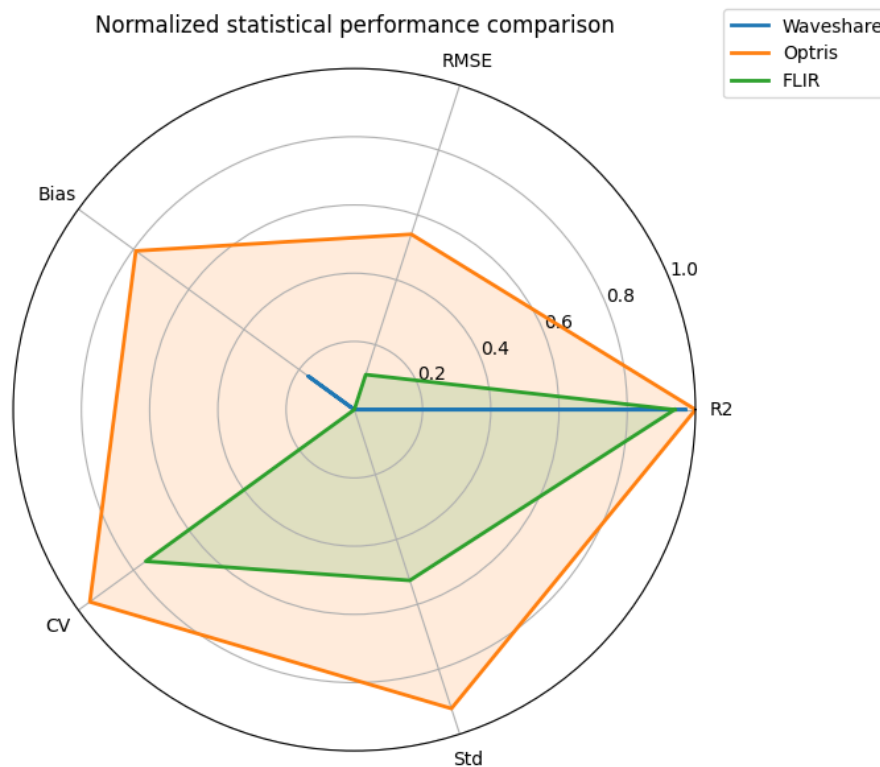


Figure 4.8: Normalized radar chart comparing the statistical performance of the Waveshare, Optris, and FLIR thermal cameras based on R^2 , RMSE, bias, coefficient of variation, and standard deviation.

From the normalized representation, the FLIR camera exhibits a balanced performance across all evaluated metrics, characterized by a high coefficient of determination, reduced root mean square error, and a bias close to zero. This combination indicates a strong linear agreement with the reference temperature, limited systematic deviation, and controlled variability. The Optris camera also shows a high R^2 , reflecting a consistent linear response; however, its larger bias and dispersion metrics suggest the presence of systematic deviations and increased variability under the evaluated conditions. In contrast, the Thermal Camera HAT (Waveshare) presents lower dispersion indicators, such as reduced standard deviation and coefficient of variation, but a diminished coefficient of determination and higher global error, which indicate a reduced ability to track the reference temperature accurately across the evaluated range.

The radar chart confirms that each camera exhibits a distinct trade-off between correlation, accuracy, and variability. These results support the methodological decision adopted throughout this work: rather than applying inter-sensor calibration corrections, the statistical characteri-

zation is used to delimit the operational ranges and experimental conditions under which each device provides sufficiently consistent measurements. Within the temperature interval and acquisition distances selected for subsequent experiments, the observed levels of dispersion and bias are considered acceptable, and the thermal data are therefore employed directly for database construction and super-resolution analysis.

4.2 Controlled Acquisition Protocol

The thermal database was constructed following a controlled and reproducible acquisition procedure that ensured uniform experimental conditions across all materials and configurations. The setup was designed to capture the surface thermal response of biological and inert samples subjected to radiofrequency (RF) excitation within the 2.45 GHz. The experimental system consisted of two subsystems: an RF heating unit and a multi-camera thermal imaging unit, both operating under constant laboratory conditions to minimize variability. Ambient temperature and relative humidity were maintained at 16 °C and 78%, respectively, throughout all sessions.

The RF subsystem included a signal generator (Rohde & Schwarz SMCV100B), a 6 W linear power amplifier (Mini-Circuits ZVE-6W-83+), a three-port circulator (Pasternack PE8432), and a 50 Ω matched load (Mini-Circuits KARN-50-18+). The applicator antenna was a modified folded dipole specifically designed for superficial hyperthermia, operating at a resonance frequency of 2.5 GHz. The antenna was positioned parallel to the sample surface, maintaining a constant spacing of 5 mm to optimize electromagnetic coupling. Under these conditions, RF power was varied between 0 and 5.5 dBm, with an extension up to 6.5 dBm for low-permittivity materials such as foam and dry wood.

Figure 4.9 shows the schematic representation of the complete system, which includes the RF generation chain and the thermal acquisition stage. The diagram illustrates the integration of the signal generator (A), power amplifier (B), circulator (C), matched load (D), and the directional antenna (E) used to irradiate the samples. The infrared cameras (F) were placed at different observation angles and distances to obtain multi-perspective radiometric data under identical heating conditions [13].




Figure 4.9: Schematic representation of the RF heating and thermal imaging system. The configuration includes: (A) RF signal generator, (B) microwave amplifier, (C) circulator, (D) matched load, (E) directional antenna, and (F) infrared cameras positioned at multiple viewing angles. Adapted from [13].

It is important to note that the configuration shown in Fig. 4.9 differs from the USRP-based setup described in Section 3.4. The present arrangement corresponds to the RF heating system employed for database acquisition. Due to differences in the signal source and power amplifier, the available excitation range was limited to 0–5.5 dBm at the generator output. Therefore, the power values reported in this chapter are not directly comparable to those presented in Section 3.4.

Thermal data were acquired using the three radiometric infrared cameras described in Section 2.3.3. Each device was calibrated before acquisition to ensure radiometric accuracy and operated in absolute temperature mode. The emissivity was fixed at 0.98 for all cameras, a standard value for biological tissues, to maintain uniform calibration. To prevent cross-heating effects between acquisitions, a 15–20 minute stabilization period was introduced after each test to allow the sample to reach thermal equilibrium.

It is important to clarify that the three infrared cameras (HR, MR, and LR) did not operate simultaneously. The acquisition was performed in a fully sequential manner, following a fixed order: first the high-resolution (HR) camera, then the medium-resolution (MR) camera, and finally the low-resolution (LR) camera. To minimize temporal inconsistencies associated

with continuous RF excitation, images were captured only after the sample reached a thermal steady-state condition. The stabilization criterion was defined as a significant reduction in the heating rate, such that incremental temperature changes occurred slowly over time. Under these conditions, the temporal offset between consecutive acquisitions was limited to the time required to trigger and store each frame, resulting in temperature differences typically below 0.5 °C between cameras. These small deviations are attributable to sequential capture and inherent sensor dispersion rather than to dynamic thermal evolution. For each geometric configuration (angle and distance), the three cameras were always acquired consecutively before modifying any experimental parameter, ensuring consistency across modalities.

Samples were positioned on a flat, plastic-coated surface to ensure stable contact and homogeneous field coupling. Organic specimens (pork, beef, and chicken tissues) were used fresh, stored at -5 °C , and allowed to reach ambient temperature before exposure to RF radiation. Inert materials (foam, wood, and phantoms) were preconditioned to control surface humidity; porous samples were soaked in water for 12 hours prior to measurement to achieve consistent moisture content. Each sample was cut into $5\text{ cm} \times 5\text{ cm} \times 6\text{ cm}$ blocks to maintain comparable volumes across materials.

The acquisition geometry and camera placement are depicted in Figure 4.10, which presents the experimental configuration adopted during the imaging process. The antenna (C) was placed directly above the sample (B), maintaining a fixed 5 mm separation. The infrared cameras (A1, A2) were mounted on tripods (E1, E2) and positioned at predefined angles of 0° , 30° , and 60° , and distances of 20 cm or 30 cm from the sample. A black curtain (D) was used to minimize reflections and isolate the measurement field from external infrared radiation [13].

Each acquisition sequence followed a structured protocol to ensure repeatability, as described below:

1. Placement and alignment of the sample on the measurement platform to ensure stable contact and uniform coupling with the RF field.
2. Adjustment of the RF generator power according to the experimental plan and material properties.




Figure 4.10: Spatial configuration used during data acquisition. Cameras (A1, A2) correspond to the HR and MR thermal systems mounted on tripods (E1, E2). The antenna (C) was placed above the sample (B) at a constant 5 mm distance, with a black curtain (D) to reduce reflections. Adapted from [13].

3. Configuration of the camera position and focus based on the selected distance and observation angle.
4. Image acquisition once the surface temperature reached a steady state. Thermal frames were recorded every five minutes under continuous RF exposure to evaluate temperature stabilization and repeatability.

The acquisition matrix systematically combined all parameter configurations: three viewing angles (0° , 30° , 60°), two distances (20 cm, 30 cm), and multiple power levels (0–5.5 dBm, or up to 6.5 dBm for porous materials). This design ensured a diverse and representative dataset covering multiple spatial resolutions and heating dynamics. All acquisitions were labeled according to their material, camera resolution, power level, distance, and viewing angle.

In addition to thermal acquisition, RGB images were recorded for each sample to provide complementary visual information and support multimodal analysis. RGB data were acquired using the rear camera of a commercial smartphone (Motorola G84), which was employed as a low-cost optical imaging device. The camera was operated in standard photo mode, and one RGB image was captured for each thermal acquisition condition.

The RGB camera was positioned manually at the same nominal distances and observation angles used for the thermal cameras, ensuring geometric consistency between modalities. The acquisition was performed under identical illumination conditions for all samples, and no automatic scene-dependent adjustments were manually modified during capture. This approach ensured that the RGB images corresponded spatially and contextually to the thermal data without introducing additional variability.

The RGB images were used as auxiliary information to support multimodal alignment, visual interpretation of material features, and subsequent processing stages involving thermal–optical data fusion. For each sample, a single RGB image was associated with the corresponding thermal measurements, maintaining a one-to-one correspondence within the dataset.

To ensure database consistency, acquisitions were performed exclusively under controlled laboratory conditions, with an ambient temperature of 16 °C and a relative humidity of 78%. Only images obtained within these environmental parameters were included in the dataset. Blurred images were discarded, as the camera focus was adjusted manually, and occasionally, human error resulted in out-of-focus frames. In these cases, additional captures were made. However, frames with reduced sharpness were excluded to preserve spatial integrity without affecting the associated temperature measurements. Furthermore, at the end of certain experimental sessions, some biological samples exhibited a significant loss of moisture accompanied by a decrease in surface temperature due to dehydration during prolonged heating. Images acquired under these conditions were excluded, as the physical integrity of the sample was no longer preserved. Acquisitions made outside of controlled laboratory conditions or without a stable steady-state temperature were not included in the final database. These criteria ensured the radiometric reliability and structural consistency of all preserved samples.

4.2.1 Ethical Considerations

This study did not involve human participants or live animal experimentation. The biological tissue used for experimental validation consisted of animal meat purchased from authorized commercial establishments where it is routinely processed and distributed for human consumption. No animals were sacrificed specifically for research purposes. All experimental procedures

were conducted under controlled laboratory conditions and in accordance with institutional safety and biosafety regulations. No personal data were collected during this research.

4.3 Structure and Validation of the Database

The thermal image database was designed to provide a structured, reproducible, and extensible framework for analyzing temperature responses of different materials under controlled RF exposure. Each dataset element contains quantitative temperature values recorded on a per-pixel basis, along with metadata describing the acquisition conditions such as resolution, distance, power, and viewing angle. The structure was developed to facilitate its use in computer vision applications, super-resolution modeling, and future multimodal registration studies.

Database Composition and Labeling

The dataset is composed of both biological and inert samples, each labeled according to a unique alphanumeric identifier that ensures traceability throughout acquisition and processing. Table 4.5 lists all materials included in the database and their associated labels. Organic tissues were selected to represent biological variability in thermal and dielectric properties, while inert materials served as stable references to assess the repeatability of RF-induced heating.

Table 4.5: Materials included in the database and their respective labels. Adapted from [13].

Material Type	Label
Phantom	PH
Pork	PK
Beef Rib	BR
Foam	FM
Wood	WD
Chicken	CH
Pork Skin	PS
Chicken Breast	CB
Beef	BF

Each file is named following a standardized convention that encodes the experimental conditions: `Material_Resolution_Tstart_Tend_Distance_Angle.csv`. For example,

the file `Pork_MR_19_55_30_0.csv` corresponds to a pork sample recorded with the medium-resolution (MR) camera, where the surface temperature increased from 19°C to 55°C, acquired at a distance of 30 cm and an angle of 0°. This nomenclature ensures full traceability and allows automated data parsing during processing and model training.

The entire dataset is organized into three resolution-based directories: `/HR/`, `/MR/`, and `/LR/`. Each directory includes two subfolders: `/Original/` (containing the unprocessed thermal images) and `/Augmented/` (containing synthetic variations generated through the augmentation pipeline). This hierarchical organization simplifies selective loading by resolution or material type and supports systematic benchmarking of computer vision algorithms.

Figure 4.11 shows the general organization of the dataset and its hierarchical directory tree.

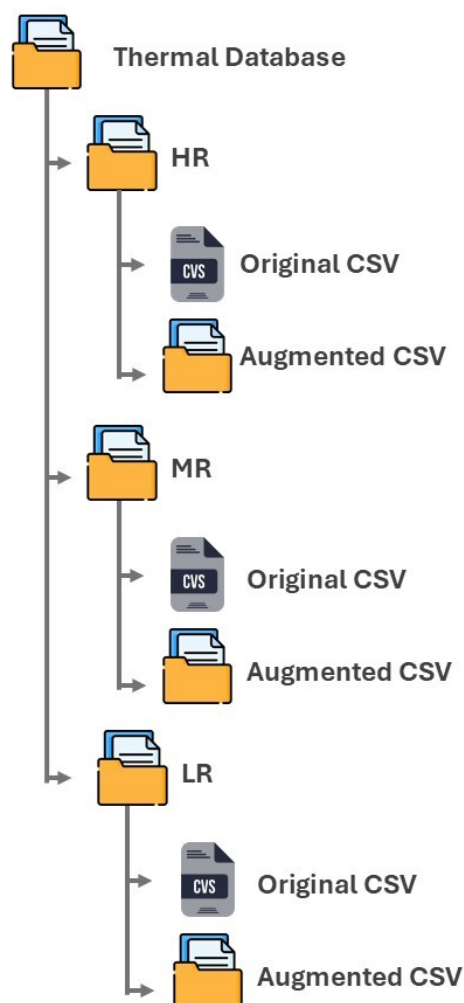


Figure 4.11: Hierarchical structure of the thermal image database organized by spatial resolution (HR, MR, LR). Each directory contains original and augmented subfolders. Adapted from [13].

In addition to the thermal data, the dataset includes a separate directory dedicated to RGB images, labeled /RGB/. This folder contains a total of 1200 RGB images, with one image captured per sample under the same experimental conditions as the corresponding thermal acquisition. Each RGB file follows the same naming convention used for the thermal data, ensuring direct correspondence between modalities. This one-to-one naming strategy enables straightforward pairing of RGB and thermal samples during multimodal processing, registration, and data fusion tasks.

Data Augmentation Pipeline

To increase the robustness and generalization capacity of future machine learning models, a structured data augmentation strategy was implemented for all three resolution levels. The pipeline introduced controlled variability in geometric and photometric domains, maintaining the physical coherence of the temperature distribution.

For the high-resolution (HR) images, only minimal transformations were applied to preserve fidelity, such as random rotations within $\pm 5^\circ$ (probability 0.15) and horizontal flips (probability 0.5). For medium-resolution (MR) data, moderate transformations were introduced, including rotations within $\pm 10^\circ$ ($p=0.25$), horizontal flips ($p=0.5$), light translations and scaling ($p=0.2$), and Gaussian noise ($p=0.1$). For low-resolution (LR) data, more aggressive transformations simulated the variability typical of low-cost thermal sensors: rotations within $\pm 15^\circ$ ($p=0.4$), salt-and-pepper noise ($p=0.2$), Gaussian thermal noise ($p=0.25$), and intensity jitter ($p=0.2$). This design ensured the dataset remained physically meaningful while representing a broader range of acquisition conditions.

Figures 4.12–4.14 illustrate representative examples of the augmented data for each resolution level.

After augmentation, the database expanded from 3,600 to 4,200 images, distributed evenly across the three resolution levels. Each level contains approximately 1,200 original thermal frames and 200 augmented images. Figure 4.15 summarizes the number of samples per material type, ensuring balanced representation among biological and inert categories.

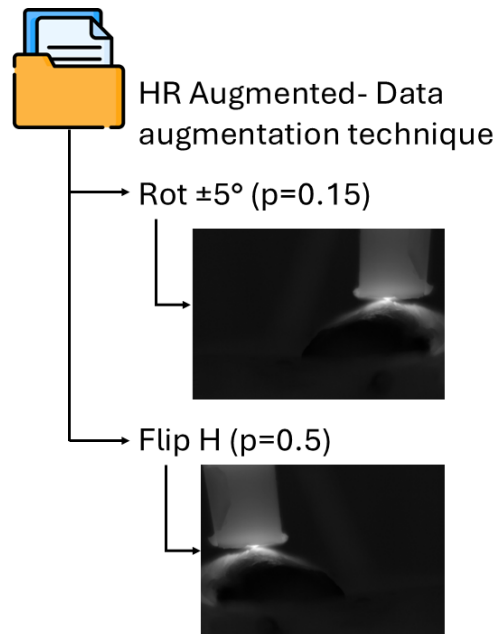


Figure 4.12: High-resolution (HR) thermal images after minimal augmentations preserving structural integrity. Adapted from [13].

Validation and Quality Assurance

The validation of the database involved both experimental and structural assessments. Experimentally, each acquisition session was performed three under identical conditions to verify temperature reproducibility and evaluate sensor stability. Experimental repetitions under identical acquisition conditions demonstrated high reproducibility of the heating profiles across both biological and inert samples. No significant deviations were observed in the recorded temperature-time curves between repetitions, indicating consistency in the acquisition protocol. Structurally, all files were reviewed to confirm correct labeling and metadata encoding, ensuring the absence of missing or corrupted entries.

The thermal response variability among materials was quantified by measuring the temperature difference between initial and final steady-state conditions under a fixed input power of 5.5 dBm. Table 4.6 summarizes the average temperature rise (ΔT) for each material type.

These results confirmed the reproducibility of the heating response, demonstrating clear differentiation between organic and inert materials. The reproducible temperature rise profiles validate the thermal stability of the acquisition protocol and the robustness of the resulting dataset. Figure 4.16 provides a representative example of images obtained for a pork skin

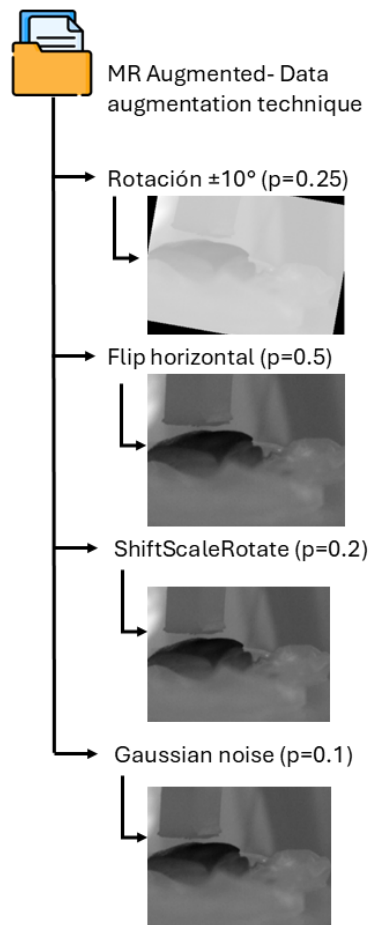


Figure 4.13: Medium-resolution (MR) thermal images with moderate augmentations, including rotation and noise injection. Adapted from [13].

Table 4.6: Initial and final surface temperatures for each material at 5.5 dBm RF power. Ambient temperature ranged between 16–19°C. Adapted from [13].

Material	Initial (°C)	Final (°C)	ΔT (°C)
Phantom	17.3	28.2	10.9
Pork	18.1	50.5	32.4
Beef Rib	16.9	52.0	35.1
Foam	17.2	31.0	13.8
Wood	16.8	30.4	13.6
Chicken	17.5	46.8	29.3
Pork Skin	18.0	49.1	31.1
Chicken Breast	17.4	47.6	30.2
Beef	17.8	51.3	33.5

sample under standard conditions (30 cm, 30° angle, 45°C surface temperature).

For each thermal acquisition, a corresponding RGB image of the sample was also recorded




Figure 4.14: Low-resolution (LR) thermal images with aggressive augmentations simulating noisy or degraded conditions. Adapted from [13].

using a visible-spectrum camera under the same geometric conditions (distance and viewing angle). These RGB images provide complementary structural and textural information of the sample surface and are used exclusively for multimodal alignment, registration, and data fusion

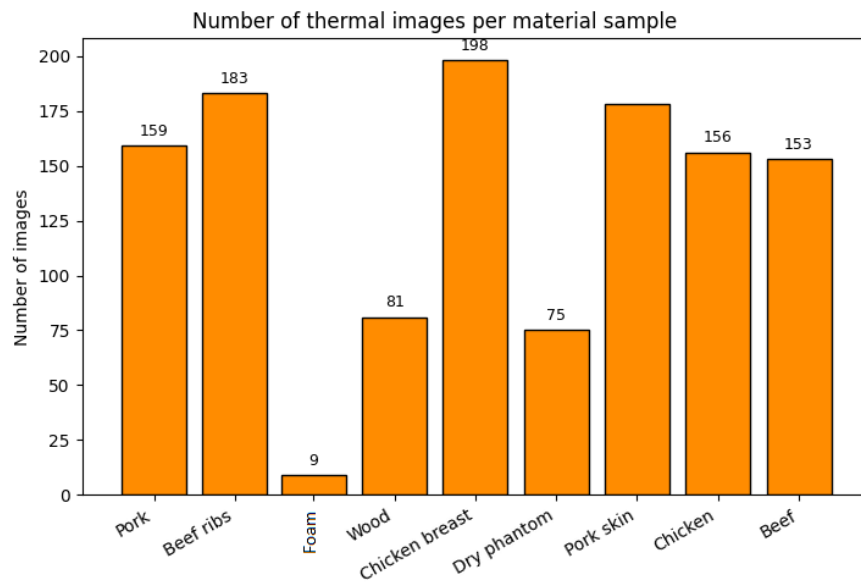


Figure 4.15: Distribution of thermal images by material type across all resolution levels. Adapted from [13].

purposes. The RGB images are not employed for temperature estimation but serve as a reference modality to support spatial correspondence with the thermal data.

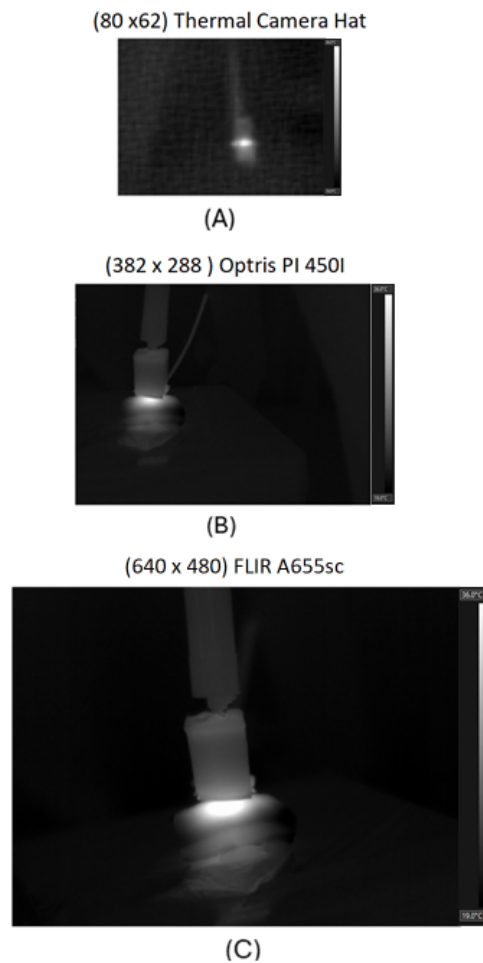


Figure 4.16: Example of thermal images recorded with the three infrared cameras under controlled conditions. (A) Thermal Camera HAT (80×60), (B) Optris PI 450i (382×288), and (C) FLIR A655sc (640×480). The sample corresponds to pork skin heated up to 45°C at a 30° viewing angle. Adapted from [13].

Chapter 5

Thermal Image Processing

5.1 Noise Reduction Techniques

Noise reduction was addressed through a combination of acquisition constraints and digital post-processing, with the objective of minimizing radiometric variability while preserving the true thermal distribution of the samples.

First, geometric restrictions were imposed during image capture by defining a minimum acquisition distance of 10 cm between the cameras and the sample. This configuration showed lower radiometric dispersion compared to shorter distances. In addition, all images were recorded only after the system reached thermal steady state. The stability criterion was defined using the temporal temperature variation measured on the region of interest, requiring the condition $|dT/dt| < 0.02 \text{ }^\circ\text{C}/\text{min}$ during a continuous observation interval of at least 3 minutes. This condition ensured that the heating dynamics had significantly slowed down and that the temperature field remained approximately constant during the sequential acquisition of the images, thereby reducing temporal variability associated with local thermal gradients.

Since the cameras provide manual focusing, fine focus adjustment was performed prior to each acquisition session in order to maximize the sharpness of the region of interest. This procedure reduced spatial blur, which introduces high-frequency components that manifest as structural noise in the thermal maps.

The thermal images were then converted to grayscale to ensure that only radiometric in-

formation was processed, removing color components introduced by the visualization palettes of each manufacturer. This normalization step enabled a consistent treatment of data acquired with cameras operating at different resolutions and internal configurations.

Finally, a two-dimensional Gaussian filter was applied to each image to attenuate high-frequency spatial noise while preserving the macroscopic thermal distribution. The filter was implemented using a kernel size of 5×5 pixels and a standard deviation of $\sigma = 0.8$. These parameters were selected empirically after evaluating different configurations in order to reduce sensor noise without smoothing relevant thermal gradients within the region of interest. Larger kernel sizes or higher variance values were avoided because they introduce excessive spatial smoothing, which can distort temperature boundaries and reduce the spatial fidelity of the thermal field.

5.2 Multimodal Registration and Alignment

The alignment between thermal and visible images constitutes a multimodal registration problem in which direct intensity correspondence cannot be assumed, since each sensor relies on different physical measurement principles. In this context, the implemented registration strategy was designed to integrate global geometric transformations with local deformation models, together with similarity metrics that are robust to inter-modal discrepancies. Figure 5.1 presents a representative example of the registration process, showing the original RGB image, the independently acquired thermal image, and the result obtained after applying the proposed transformation model.




Figure 5.1: Example of the multimodal registration process. From left to right: RGB image, unaligned thermal image, and the result after applying the geometric registration model.

As a first step, a differentiable affine transformation was applied to correct global misalignments associated with translations, rotations, anisotropic scaling, and shear effects. This model is grounded on the concept of spatial transformers introduced by Jaderberg et al. [64], which enables the estimation of geometric parameters within a fully differentiable optimization framework. This stage compensates for errors derived from variations in relative camera placement, changes in viewing angle, and manual focus adjustments.

Subsequently, local misalignments were modeled using Thin Plate Splines, a non-rigid registration technique that represents the deformation field as a smooth function interpolated from a discrete set of control points. This formulation introduces a regularization term based on the minimization of bending energy, which ensures spatial continuity and prevents abrupt distortions. Its application allowed the correction of residual errors associated with perspective differences, parallax effects, and optical distortions between the thermal and visible imaging systems.

Since thermal and visible intensities are not directly comparable, the alignment process was driven by similarity metrics that are independent of the radiometric domain. In particular, mutual information was employed as the objective function following the classical formulation of Viola and Wells [65], due to its ability to quantify statistical dependence between heterogeneous distributions. In addition, a gradient correlation term was incorporated to reinforce structural correspondence between edges and contours present in both modalities.

To prevent non-physical deformations and ensure geometric coherence, a regularization term was included to penalize abrupt variations in the displacement field, promoting smooth and locally continuous transformations. This type of constraint is consistent with modern deformable registration frameworks reported in recent literature, such as the VoxelMorph approach [66].

Finally, the applied transformation was defined as the composition of an affine model and a non-rigid deformation field, enabling the simultaneous correction of global positioning errors and local discrepancies between the thermal and visible modalities.

The quality of the alignment was assessed through visual inspection of structural correspondences between both modalities, particularly along object contours, edges, and reference elements present in the scene. A fully quantitative metric was not adopted as a primary valida-

tion criterion because thermal and visible images represent different physical measurements and exhibit differences in spatial resolution and radiometric distribution. Under these conditions, direct pixel-wise comparison metrics may not accurately reflect geometric alignment quality. Therefore, the evaluation focused on verifying the spatial coherence between the registered images and the preservation of relevant structural features required for the subsequent multimodal fusion stage.

5.2.1 Multimodal Fusion of Thermal and RGB Images

Once the geometric registration between thermal and visible images was completed, the multimodal fusion stage was performed. The objective of this step is to integrate the radiometric information provided by the infrared sensor with the high-resolution structural content of the RGB image. This integration enables the generation of hybrid representations in which the spatial distribution of temperature is preserved while improving the geometric definition of the scene.

Figure 5.2 illustrates the general processing flow, which starts from the initial thermal and RGB images, continues with the multimodal registration stage, and concludes with the application of the selected fusion strategy.

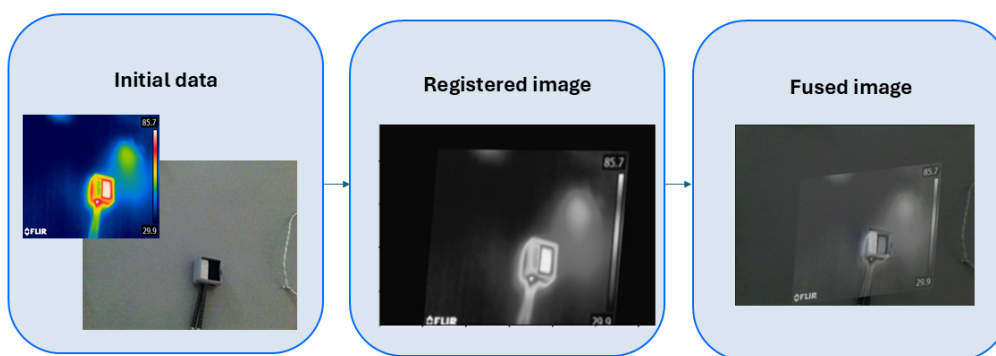


Figure 5.2: Multimodal fusion workflow. The process starts from the initial thermal and RGB images, applies geometric registration, and subsequently performs the selected fusion strategy.

To assess the effect of the integration strategy, three fusion methods were implemented: wavelet-based decomposition, multiscale pyramid fusion, and direct blending. Figure 5.3

presents a visual comparison of the results obtained with each approach.

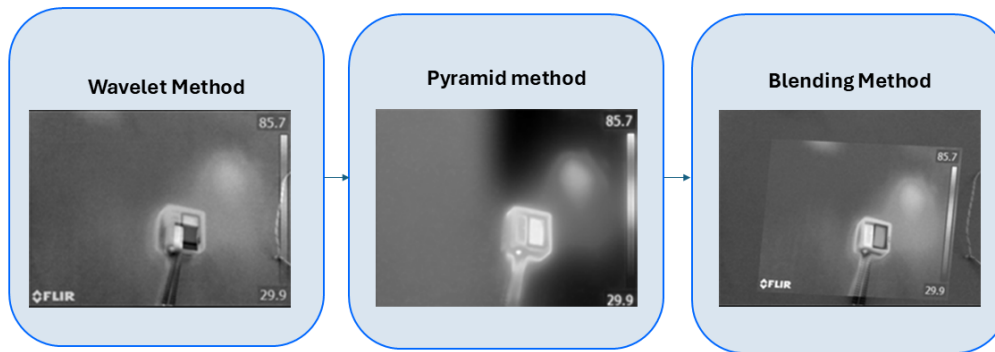


Figure 5.3: Comparison of the implemented multimodal fusion methods. The image shows the results obtained using wavelet fusion, multiscale pyramid fusion, and blending, all applied to the same registered image pair.

The results reveal distinct behaviors among the three evaluated fusion methods. The wavelet-based approach exhibited limitations in consistently integrating the geometric information from the RGB image with the thermal data, particularly along the applicator contours and in regions with abrupt intensity transitions. This behavior resulted in local misalignments and poorly defined edges, which restrict its ability to coherently represent both modalities within a single image.

In the case of multiscale pyramid fusion, a better hierarchical integration of information was observed; however, the reconstruction process introduced a smoothing effect that led to partial loss of detail in specific regions of interest. In particular, areas near the applicator present a reduction in structural sharpness, which is associated with the progressive filtering applied at each pyramid level.

By contrast, the direct blending method achieved the best visual correspondence between the thermal and RGB images. This strategy preserved both the scene geometry and the localization of high-temperature regions, maintaining spatial coherence and edge definition. Consequently, blending was selected as the fusion strategy employed in the subsequent processing stages, as it provides a consistent representation of both modalities without introducing noticeable distortions.

Chapter 6

Integration of the Thermal Acquisition System with Closed–Loop Control

6.1 System Integration and Experimental Setup

The thermal acquisition platform was integrated into a closed–loop control system in which radiofrequency (RF) excitation, temperature sensing, and control algorithms operate as a unified framework. Figure 6.1 presents a schematic overview of the complete system, which consists of a USRP - 2922 software–defined radio used as the RF signal source, a linear RF power amplifier, a 2.45 GHz applicator antenna, a host computer for real–time processing and control, and a Thermal Camera HAT employed as the primary feedback sensor for temperature regulation.




Figure 6.1: Schematic representation of the integrated closed-loop thermal system. The architecture includes: (A) USRP - 2922 SDR acting as RF generator, (B) RF linear power amplifier, (C) host computer executing the control algorithms, (D) 2.45 GHz applicator antenna, (E) thermal camera for monitoring, and (F) biological or phantom sample.

The RF excitation is generated digitally by the NI USRP - 2922, which enables flexible control of frequency, amplitude, and modulation parameters. The generated signal is amplified by the RF power amplifier and delivered to the applicator antenna, positioned at a fixed distance of approximately 5 mm from the sample surface in order to ensure stable electromagnetic coupling. The host computer coordinates the transmission parameters and executes the closed-loop control algorithms in real time.

The physical implementation of the experimental platform is shown in Figure 6.2. A custom mechanical mounting structure was developed to guarantee repeatable positioning of all elements, including the antenna, the Thermal Camera HAT, and the auxiliary thermal cameras. The FLIR and Optris cameras were incorporated as reference sensors to validate the radiometric measurements obtained with the Thermal Camera HAT and to assess consistency across devices.

The Thermal Camera HAT was selected as the primary feedback sensor in the closed-loop control architecture because the objective of the system is to evaluate thermal regulation using low-cost sensing devices whose measurements may present larger dispersion compared with

laboratory-grade instruments. In this framework, the HAT camera provides the temperature signal used by the control algorithm, while the FLIR and Optris cameras are maintained as reference sensors to verify the consistency of the thermal measurements obtained during the experiments. This configuration allows the control system to operate using the same sensing conditions expected in low-cost thermal acquisition scenarios, while still enabling independent radiometric verification of the results through higher-performance instruments.




Figure 6.2: Photograph of the experimental mounting system. The structure holds the RF applicator antenna (c), the Thermal Camera HAT used for closed-loop control (b), and the FLIR and Optris cameras (a) employed as reference instruments for validation.

All sensors were aligned to observe the same region of interest on the sample surface, allowing direct comparison of temperature measurements across modalities. This configuration enabled the evaluation of the closed-loop control strategies under identical thermal conditions while ensuring traceability and validation of the acquired data.

6.2 Closed-Loop Thermal Control Evaluation

6.2.1 Control strategies and experimental methodology

Based on the experimental setup illustrated in Fig. 6.1, an initial open-loop characterization of the thermal plant was performed. The system was driven at its maximum available RF excitation in order to observe its intrinsic thermal dynamics without feedback action, and temperature data were recorded until steady-state conditions were reached. The resulting open-loop thermal response is presented in Fig. 6.3. In order to mitigate high-frequency disturbances caused by electromagnetic interference and sensor noise, the acquired signal was processed using a low-pass filtering stage, preserving the dominant dynamics of the process while suppressing spurious fluctuations. The filtered response was then used to obtain an approximate first-order model of the plant through the system identification tools provided by MATLAB, following the methodology described in the MathWorks documentation [67] and the classical treatment of PID tuning for thermal processes [68].




Figure 6.3: Open-loop thermal response of the system obtained under maximum RF excitation, after low-pass filtering.

The tuning procedure was carried out by specifying the first-order plant dynamics and

enforcing stability margins consistent with thermal systems. The PID Tuner provided the optimal gains that guarantee a compromise between rise time and robustness, resulting in $K_p = 0.7882$, $K_i = 7.826 \times 10^{-4}$ and $K_d = 11.31$. These parameters were fixed and applied identically to the three control strategies in order to ensure a fair comparison. The conventional PID controller follows the standard error–based formulation that combines proportional, integral and derivative actions to regulate the RF gain applied to the transmitter [69]. In this architecture, the control signal is computed directly from the instantaneous temperature error and is limited only by the physical saturation bounds of the actuator.

The small value obtained for the integral gain ($K_i = 7.826 \times 10^{-4}$) is consistent with the thermal dynamics of the experimental setup and with the physical constraints of the acquisition system. During the experiments, the thermal camera was positioned close to the sample surface in order to maintain sufficient visibility of the region of interest when using the low–cost sensor. In this configuration, the proximity between the camera and the RF applicator occasionally introduced perturbations associated with electromagnetic activity during the heating stage. Additionally, the thermal process itself evolves slowly, with characteristic response times on the order of 8 to 12 minutes before reaching the steady region. Under these conditions, a reduced integral action limits the accumulation of the control signal and contributes to maintaining stable closed–loop behavior during the slow temperature evolution observed in the system.

However, the RF heating process exhibits a strongly nonlinear behavior during the initial transient due to thermal inertia and electromagnetic coupling effects, which prevents the plant from presenting a quasi–linear input–output relationship in this region. For this reason, a hybrid PID strategy was implemented, in which the control architecture incorporates a saturation stage that applies a constant actuation signal whenever the temperature error exceeds a predefined threshold [70]. In this work, the saturation region was defined as a band of 8 °C below the setpoint. The value of 8 °C was determined experimentally during the implementation of the control strategies. Several switching thresholds were tested between the constant RF actuation stage and the PID regulation stage. When the PID controller was activated closer to the setpoint, the system frequently produced positive overshoot. A threshold of 8 °C reduced this behavior and provided the most stable response among the tested values. While the process remains outside this

band, the RF gain is kept constant to accelerate the heating phase and avoid integrator windup. Once the measured temperature enters this region, the controller automatically switches to the conventional PID law, allowing finer regulation around the operating point.

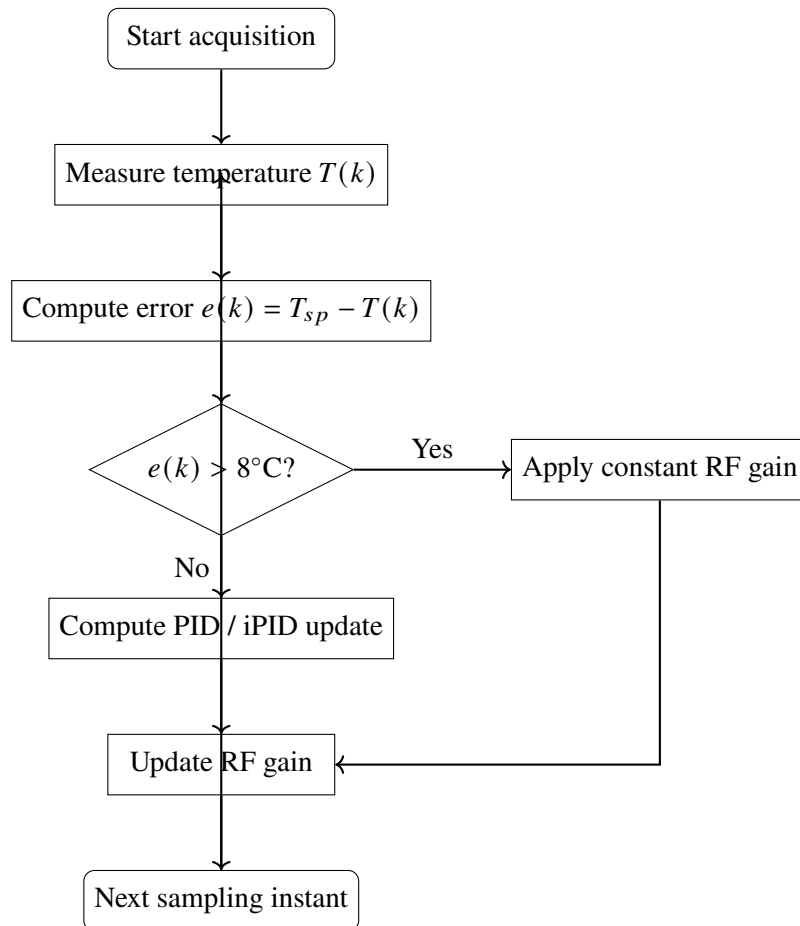


Figure 6.4: Flowchart of the hybrid PID and hybrid incremental PID control architectures. When the temperature error exceeds the regulation band, a constant RF actuation is applied. Inside the 8 °C region below the setpoint, the controller switches to PID or iPID regulation.

The control logic illustrated in Fig. 6.4 summarizes the switching mechanism implemented in both hybrid strategies. In addition, a hybrid incremental PID (iPID) controller was implemented using the same architectural logic, but updating the control action in an incremental form rather than computing its absolute value [71]. This approach improves numerical stability and robustness to measurement noise by basing the actuation on variations of the error signal. As in the hybrid PID case, the iPID controller operates under constant actuation outside the regulation band and transitions to incremental closed-loop control only when the temperature approaches

the reference, thus combining fast transient response with smooth steady-state regulation. All three controllers were implemented using the same gain values obtained from the open-loop identification stage, ensuring that any performance differences are exclusively due to the control structure and not to parameter tuning. The entire methodology was applied identically to phathom and pork meat samples in order to evaluate the robustness of the strategies with respect to different thermal properties of the material

6.2.2 Results and comparative analysis




Figure 6.5: Closed-loop thermal regulation scheme used in all experiments, including RF excitation, thermal sensing, control algorithm and actuation stages.

The closed-loop thermal regulation scheme adopted in this work is illustrated in Fig. 6.5. For each control strategy — conventional PID, hybrid PID and hybrid incremental PID — two independent experimental sweeps were performed. The first set of tests employed a tissue-mimicking phantom sample with a temperature setpoint of 35 °C, while the second set used ex-vivo pork meat acquired from a local food supplier, with a setpoint of 40 °C. This material

was obtained exclusively for experimental purposes and did not involve animal experimentation, ensuring compliance with ethical considerations. All experiments were conducted under identical environmental conditions, with an ambient temperature of 21 °C and a relative humidity of 50%. Using the same protocol and controller gains in all cases, the resulting thermal responses are presented and compared in the following subsections.

The feedback variable used by the control system corresponds to the maximum temperature detected within the region of interest (ROI) of the thermal image. At each sampling instant, the thermal frame is analyzed and the highest pixel value inside the ROI is selected as the control variable. This criterion was adopted because, in RF hyperthermia applications, the primary safety constraint is the prevention of local overheating. By regulating the maximum temperature rather than the spatial average, the controller directly constrains the most critical thermal point on the sample surface and reduces the risk of localized hot spots.

Under the hybrid iPID strategy, the maximum temperature at steady state was regulated within 35.1–35.2 °C for the phantom (setpoint 35 °C, relative error < 0.3%) and within a comparable margin around 40 °C for the pork sample (relative error < 0.2%). Because the control variable corresponds to the maximum pixel temperature, these values define the upper thermal bound within the ROI. Consequently, an explicit verification of spatial thermal uniformity is not required for safety purposes. Although the RF antenna produces a naturally non-uniform heating pattern, no point inside the ROI can exceed the regulated maximum temperature. In this sense, the control strategy inherently limits the worst-case thermal condition, ensuring that the imposed temperature bound is respected across the entire observed region.




Figure 6.6: Closed-loop thermal response of the conventional PID controller using the phantom sample. From top to bottom: maximum temperature evolution, temperature error, and RF gain applied to the actuator.

The closed-loop response obtained with the conventional PID controller using the phantom sample is shown in Fig. 6.6. The temperature increases gradually but in an irregular manner, with noticeable oscillations that persist throughout the experiment. Although the system exhibits a

clear heating trend, it is unable to reach the desired reference, remaining more than 10 °C below the setpoint even after 1800 s of operation, which reveals a clear limitation of the conventional PID strategy for this thermal process. The temperature error, depicted in the second subplot, confirms this limitation. The error decreases only marginally, from approximately 12 °C to 10 °C, and never enters the regulation band. This behavior indicates that the controller is unable to provide sufficient corrective action during the transient phase, which is a direct consequence of the strong thermal inertia and the nonlinear electromagnetic heating dynamics of the process.

The third subplot reveals the root cause of this poor performance. The control gain exhibits a highly oscillatory pattern with amplitudes between 0 and 4 dB throughout the entire experiment. This chattering behavior is produced by the derivative action amplifying measurement noise and by the integrator reacting to a slowly varying error signal, which results in an ineffective modulation of the RF power. As a consequence, the controller operates in a quasi-limit cycle without driving the system towards the desired operating point.

These observations indicate that the conventional PID controller is not adequate for this application, as it is unable to cope with the strong nonlinear behavior present during the initial heating stage and does not deliver enough actuation to effectively drive the process toward the reference. As a result, the system remains trapped in a low-temperature regime. This drawback motivates the adoption of hybrid control strategies, which are introduced to overcome the initial nonlinear region and are analyzed in the following section.

In view of the limitations observed with the conventional PID controller, the following results present the closed-loop response achieved with the hybrid PID strategy, evaluate the effect of the constant actuation stage on the thermal behavior of the system, and compare it with the previous system.

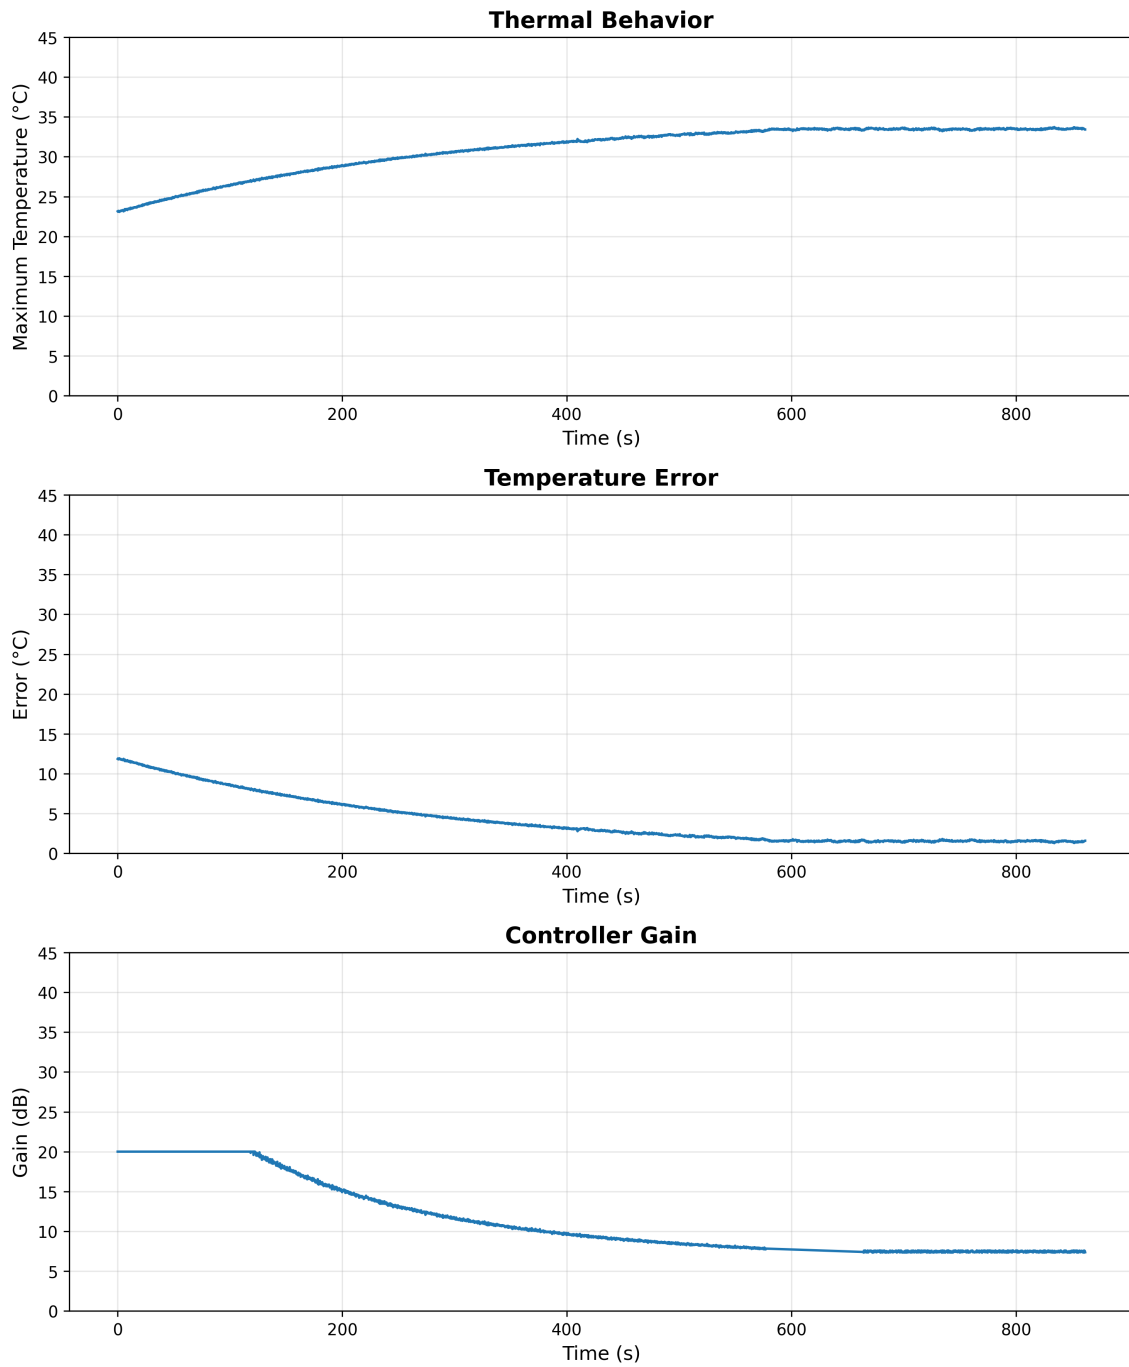


Figure 6.7: Closed-loop response using the hybrid PID controller with the phantom sample.

From top to bottom: thermal response, temperature error and controller gain.

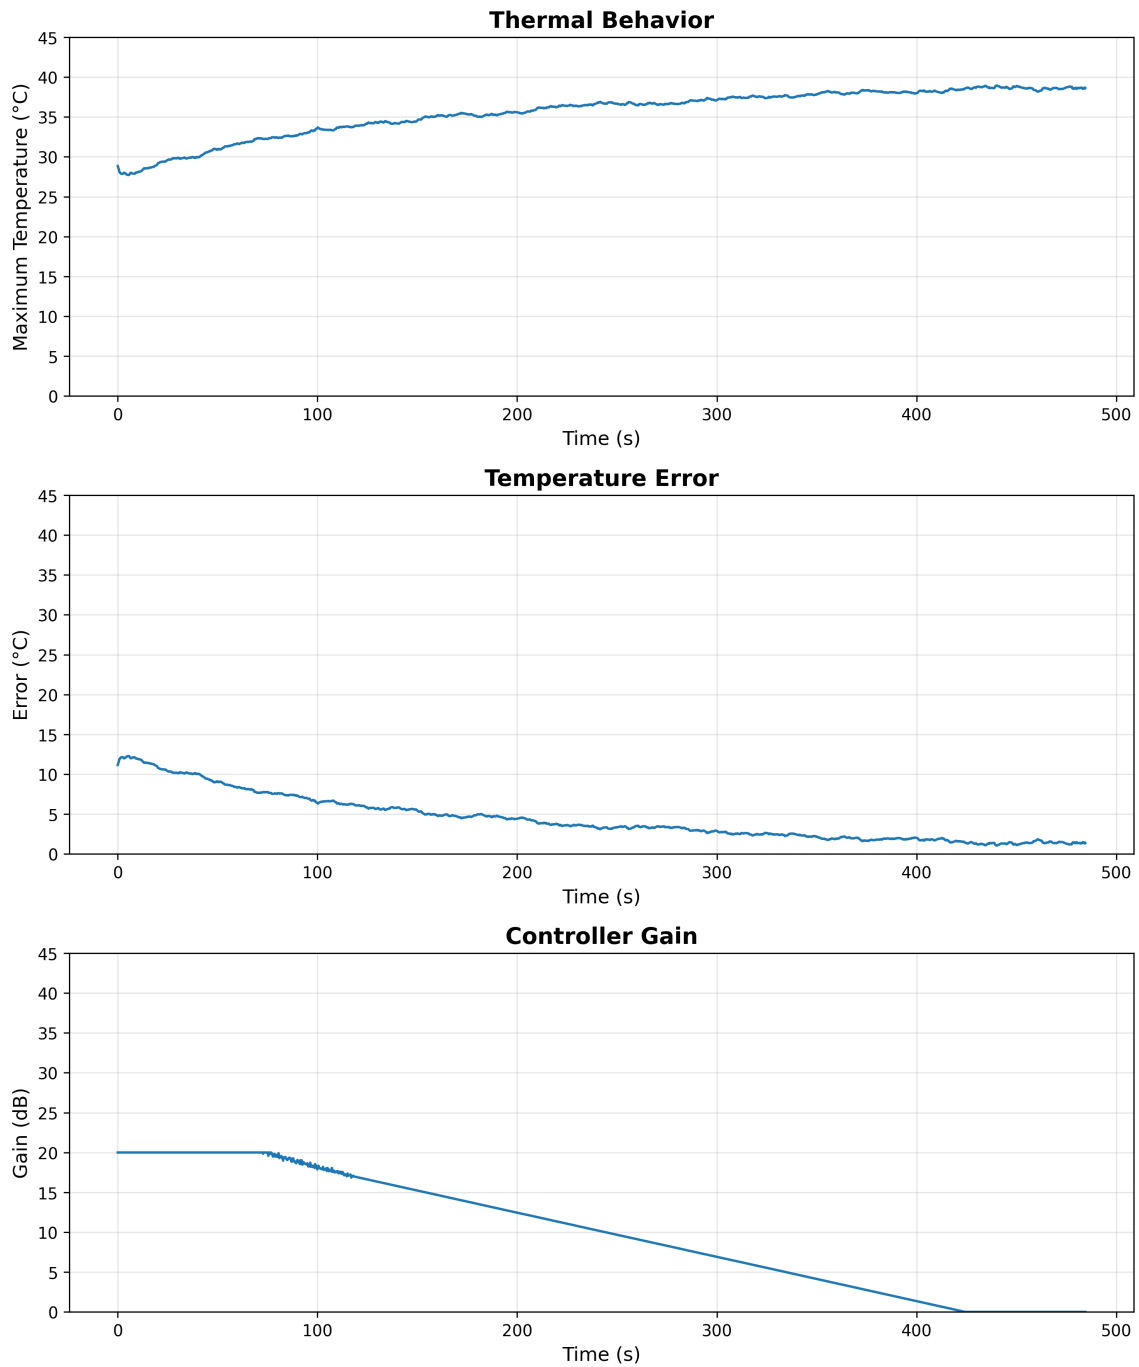


Figure 6.8: Closed-loop response using the hybrid PID controller with the pork meat sample. From top to bottom: thermal response, temperature error and controller gain.

The closed-loop responses obtained with the hybrid PID controller are shown in Fig. 6.7 for the phantom sample and in Fig. 6.8 for the pork sample. In both cases, the inclusion of the constant actuation stage allowed the system to overcome the initial nonlinear region that prevented the conventional PID controller from initiating the heating process.

In the phantom experiment, the controller stabilized the temperature around 34 °C, which took approximately 800 s to reach the steady-state region. In contrast, when the same controller gains were applied to the pork sample, the closed-loop response stabilized near 38 °C with a reduced stabilization time of approximately 500 s. This difference is mainly attributed to the moisture content of the material. Although the phantom was prepared under conditions intended to reproduce thermal behavior in a controlled manner, its longer preparation time caused a partial reduction in surface moisture, which in turn reduced its effective microwave absorption and resulted in a slower thermal response.

It is important to note that the open-loop identification stage was carried out exclusively with the phantom sample, and that the same controller parameters were subsequently applied to the pork experiments without the need for readjustment. Despite this, the hybrid PID strategy showed consistent first-order behavior in both materials, confirming the robustness of the proposed approach with respect to variations in thermal properties.

Although the hybrid PID controller provided a clear improvement in the transient behavior with respect to the conventional PID, a steady-state offset of approximately 2% was consistently observed in both experimental conditions. This deviation is evident from the difference between the stabilized temperature and the imposed setpoint. In addition, small oscillations in the control gain were detected during the regulation phase, suggesting a sensitivity of the absolute control law to the slow thermal dynamics of the process and to measurement noise. These effects limit the achievable accuracy of the hybrid PID and justify the introduction of an incremental hybrid PID formulation, which is analyzed in the following section.

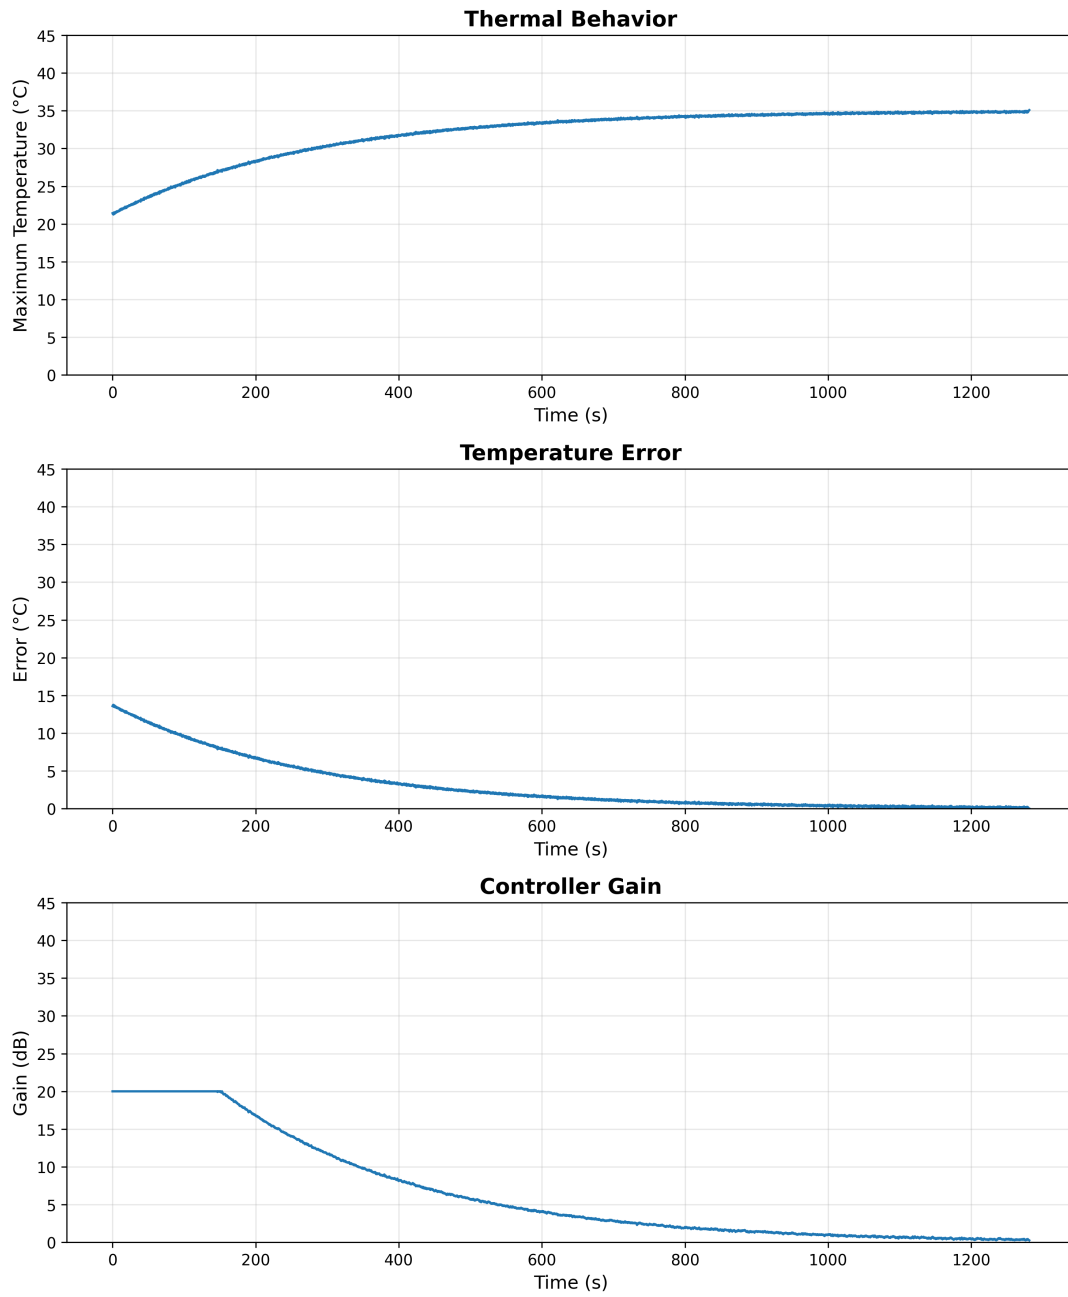


Figure 6.9: Closed-loop thermal response using the hybrid incremental PID controller with the phantom sample. From top to bottom: temperature evolution, temperature error and controller gain.

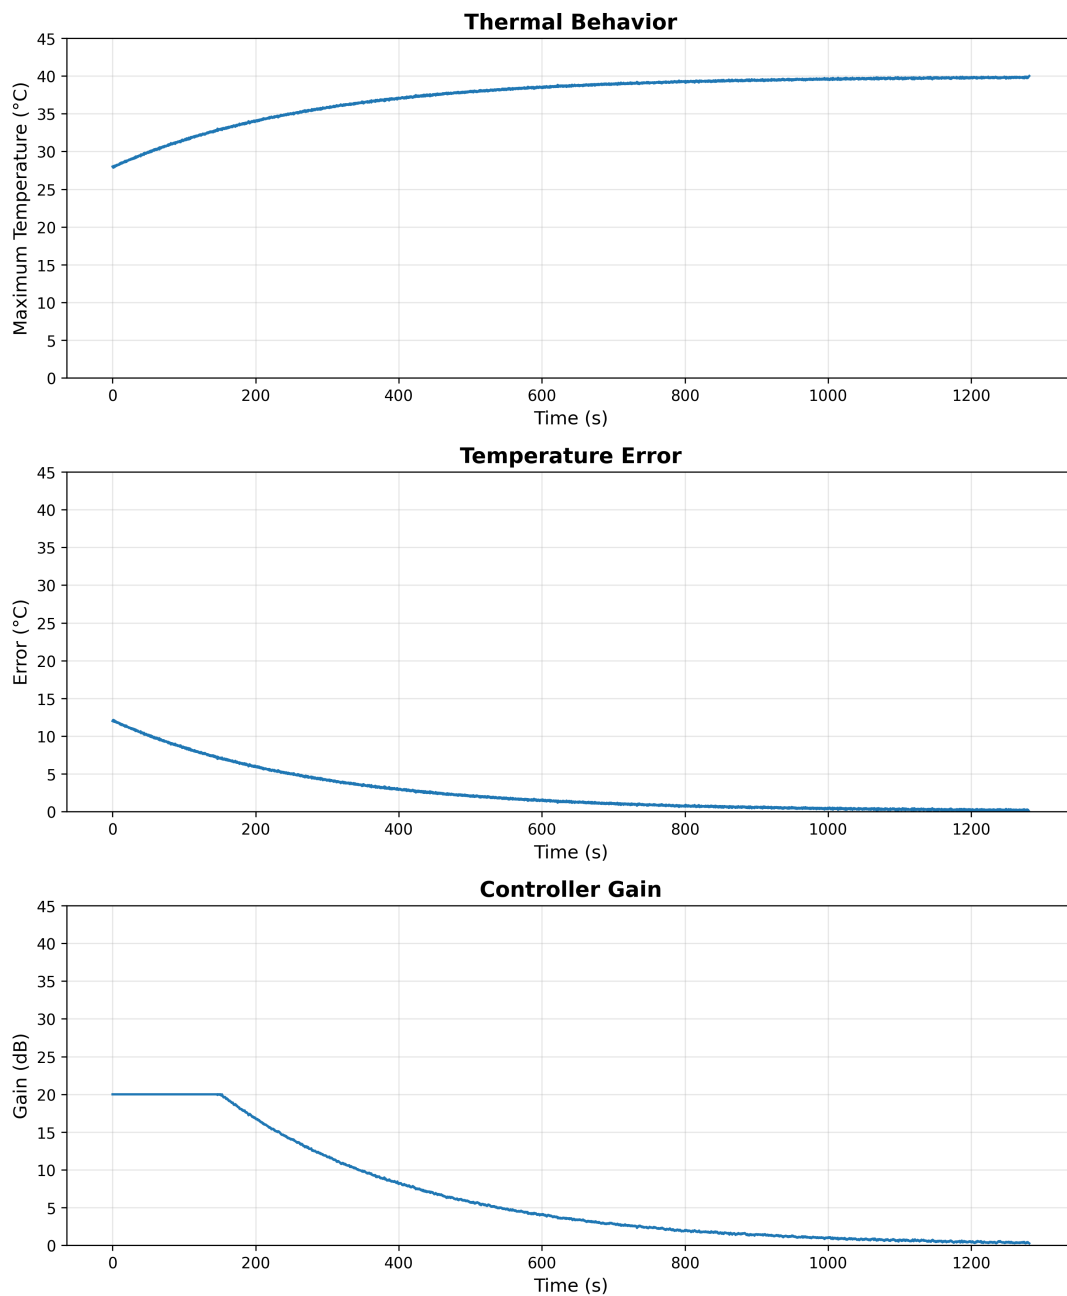


Figure 6.10: Closed-loop thermal response using the hybrid incremental PID controller with the pork meat sample. From top to bottom: temperature evolution, temperature error and controller gain.

The closed-loop response obtained with the hybrid incremental PID controller using the phantom sample is presented in Fig. 6.9. The temperature converges to the reference value of 35 °C after approximately 1300 s, exhibiting small oscillations around the setpoint, typically within the range 35.1–35.2 °C. The corresponding temperature error remains below 0.3% during

steady–state operation, indicating a high regulation accuracy. The controller gain shows smooth variations with no abrupt changes, reflecting the improved numerical stability provided by the incremental formulation.

The closed–loop response obtained with the A similar behavior is observed in the experiment conducted with pork meat, as shown in Fig. 6.10. The temperature reaches the reference value of 40 °C in approximately 1300 s and remains tightly bounded around the setpoint with negligible steady–state error. As in the phantom case, the error stays below 0.2%, while the control gain evolves smoothly without oscillatory patterns, confirming the robustness of the iPID hybrid strategy against variations in the thermal properties of the material.

The joint analysis of the three control schemes allows a clear performance hierarchy to be established. The conventional PID controller is unable to adequately excite the plant during the initial phase, due to the strong nonlinearity of the RF heating process and the high thermal inertia of the system, which prevents the setpoint from being reached even after long operating periods.

The hybrid PID controller in its absolute formulation overcomes this initial limitation by incorporating a constant actuation stage; however, it exhibits a steady–state error on the order of 2% and increased variability in the control signal, revealing a persistent sensitivity to slow system dynamics and measurement noise.

In contrast, the hybrid iPID controller effectively combines the initial excitation action with high–precision incremental regulation, achieving a steady–state error below 0.3% in both experimental scenarios. Although this approach requires a longer settling time, approximately 1300 s, the level of stability attained, together with the reduction of oscillations and the smoothness of the control signal, positions it as the most suitable strategy for RF thermal heating applications where accuracy and process repeatability are critical.

Chapter 7

Machine Learning Models for Thermal Super-Resolution

Three deep learning architectures were implemented to address the problem of thermal image super-resolution. These models follow different architectural principles, including convolutional-based deterministic processing, generative modeling, and probabilistic learning strategies. In all cases, the architectures were structured to incorporate thermal information at multiple spatial scales—low resolution (LR), medium resolution (MR), and high resolution (HR)—allowing a unified multi-scale learning framework throughout the super-resolution process.

All models were trained and evaluated using a common dataset derived from the acquisition protocol described in previous chapters. The dataset was randomly partitioned into training (70%), validation (15%), and testing (15%) subsets at the frame level. The split was performed after multimodal registration and radiometric calibration, ensuring that each sample corresponded to an aligned LR–MR–HR triplet. The assignment of samples to subsets was completed prior to model training.

During acquisition, the global geometry of the scene remained fixed under controlled laboratory conditions. The random frame-based partition therefore preserves geometric consistency across subsets while introducing radiometric variability through independent temperature sweeps. Although spatial configuration is maintained, each frame corresponds to a distinct thermal state, resulting in different pixel-wise temperature distributions. Under this protocol,

no frame was shared between training, validation, and testing subsets.

All deep learning models presented in this chapter were trained on a workstation equipped with an NVIDIA RTX 6000 Ada Generation GPU, 128 GB of RAM, and an Intel Xeon W3-2425 processor running at 3.00 GHz. The system architecture was 64-bit. These hardware specifications define the computational conditions under which the models were implemented and trained, ensuring reproducibility of the experimental setup.

7.1 Implemented Architectures

7.1.1 Multi-Resolution Convolutional Neural Network

The multi-resolution convolutional neural network (MR-CNN) was implemented as a fully convolutional architecture designed to process thermal images acquired at different spatial resolutions within a multimodal acquisition framework. The network receives low-resolution (LR) and medium-resolution (MR) thermal images as inputs and produces a high-resolution (HR) thermal image as output.

Prior to inference, all thermal inputs undergo radiometric calibration. Additionally, a multimodal registration stage is performed using the corresponding RGB images as spatial references. In this stage, LR, MR, and HR thermal images are geometrically aligned to a common coordinate system by estimating the transformation parameters from the RGB–thermal image pairs. This procedure ensures spatial consistency across resolutions and modalities before network processing.

The processing pipeline follows a deterministic feedforward structure. First, LR and MR thermal images are independently normalized and resized to a common reference grid. Feature extraction is performed through separate convolutional branches, each composed of stacked convolutional layers with fixed kernel sizes. The extracted feature maps are concatenated along the channel dimension and propagated through a fusion block that integrates multi-resolution thermal information.

Subsequently, the fused representation is passed to an upsampling stage, where spatial resolution is progressively increased using interpolation-based resizing followed by convolutional

refinement. The final convolutional layer maps the reconstructed feature representation to a single-channel HR thermal image, preserving the original temperature scale.

The complete computational flow of the MR-CNN architecture is summarized in Fig. 7.1.




Figure 7.1: Flow diagram of the multi-resolution convolutional neural network. LR and MR thermal images are processed through independent convolutional branches, fused at the feature level, and reconstructed into an HR thermal output.

The network is trained using paired thermal samples in different resolutions. The loss function is computed between the reconstructed HR image and the corresponding reference HR measurement. In this architecture, no adversarial training or probabilistic modelling is incorporated.

7.1.2 Multimodal Generative Adversarial Network

The second architecture is based on a multimodal generative adversarial network (GAN) for thermal image super-resolution. The model operates using thermal data acquired at three spatial resolutions (LR, MR, and HR), together with an RGB image used as complementary spatial information.

The generator receives the thermal inputs at multiple resolutions through independent encoding branches. Each branch extracts resolution-specific thermal features using convolutional layers. In parallel, the RGB image is processed by a dedicated encoder to extract spatial and structural information. The encoded thermal features from the LR, MR, and HR branches are then combined with the RGB-derived features at a fusion stage.

The fused representation is propagated through a decoding stage composed of successive upsampling and convolutional layers, producing a reconstructed HR thermal image. This reconstruction process integrates thermal information from multiple resolutions and spatial cues from the RGB modality.

The discriminator network receives either real HR thermal images or reconstructed outputs from the generator. It is implemented as a patch-based classifier that evaluates local regions of the input image to distinguish between real and generated thermal patterns. The adversarial interaction between the generator and discriminator constrains the reconstructed thermal images to follow the statistical distribution of the reference HR data.

A schematic flow diagram of the multimodal GAN architecture, including the generator and discriminator pathways, is shown in Fig. 7.2.




Figure 7.2: Flow diagram of the multimodal generative adversarial network.

7.1.3 Progressive Variational Autoencoder

This architecture corresponds to a progressive variational autoencoder (VAE) formulated to model thermal images at multiple spatial resolutions. The encoder follows a hierarchical structure in which low-resolution (LR), medium-resolution (MR), and high-resolution (HR) thermal images are processed through successive encoding stages. Each stage reduces the spatial dimensionality while increasing the abstraction level of the extracted thermal features.

Prior to encoding, all thermal images undergo radiometric calibration and multimodal geometric registration using the corresponding RGB images as spatial references. The registration stage estimates the geometric transformations from RGB–thermal image pairs, ensuring spatial alignment across LR, MR, and HR thermal representations before they are introduced into the VAE.

The encoder outputs the parameters of a probabilistic latent distribution, defined by a mean vector and a variance vector. Latent samples are obtained using the reparameterization strategy, enabling gradient-based optimization during training. The latent representation encodes statistical information shared across the different spatial resolutions.

The decoder mirrors the encoder structure and performs a staged reconstruction process.

Starting from the latent representation, thermal images are reconstructed sequentially at LR, MR, and HR resolutions. Each decoding stage refines the spatial content of the previous level, enforcing consistency across resolutions.

The training objective combines reconstruction losses computed independently at each resolution level with a Kullback–Leibler divergence term that regularizes the latent distribution. This formulation constrains the latent space to follow a predefined prior distribution and ensures coherence between the reconstructed thermal representations.

Figure 7.3 presents the flow diagram of the progressive variational autoencoder, illustrating the hierarchical encoding of LR, MR, and HR thermal images, the probabilistic latent space, and the staged decoding process.

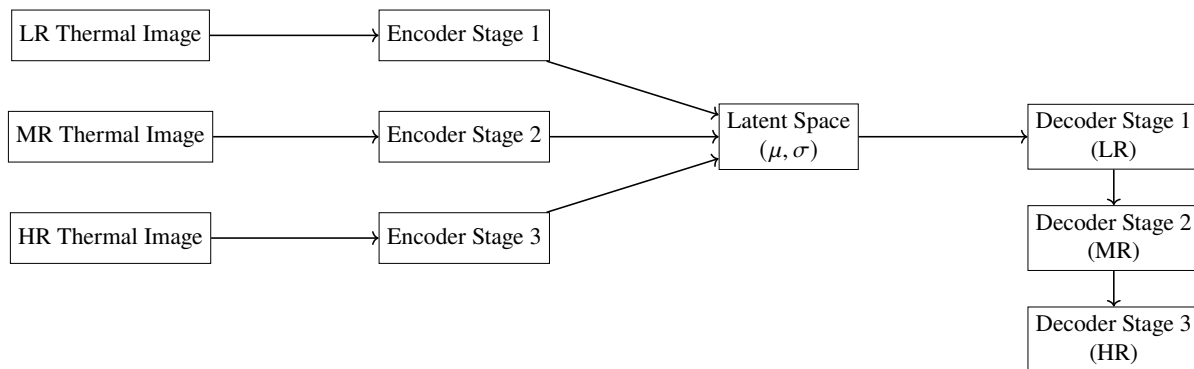


Figure 7.3: Flow diagram of the progressive variational autoencoder for multi-resolution thermal image modeling.

7.2 Training, Validation, and Testing

All three architectures were trained under identical experimental conditions in order to ensure comparability of results. The same dataset partition (70% training, 15% validation, 15% testing), preprocessing pipeline, radiometric calibration procedure, and registration strategy were applied across models. Differences in performance therefore arise exclusively from architectural and optimization characteristics rather than from variations in data handling or training protocol.

For supervised optimization, the high-resolution (HR) thermal images were treated as deterministic ground truth after radiometric calibration and inter-sensor validation. The instrumental uncertainty quantified in Chapter 4 was not explicitly incorporated into the loss functions as

a probabilistic noise model or weighting factor. However, the inter-sensor RMSE values previously reported (on the order of 1–2 °C) define an intrinsic measurement uncertainty of the acquisition system. This uncertainty establishes a practical lower bound for the achievable temperature-domain reconstruction error (Temp-L1), meaning that part of the residual error observed during evaluation may originate from measurement variability rather than exclusively from model limitations.

The dataset partition was performed using a random image-level split (70% training, 15% validation, 15% testing). All acquisitions were conducted under controlled laboratory conditions over the same physical samples, implying that the geometric structure of the scene remains largely consistent across subsets. Variability within the dataset is introduced through controlled modifications of three experimental factors: sample temperature, camera-to-object distance, and viewing angle.

Under this protocol, geometrically similar images may be present across training, validation, and testing subsets. However, because the objective of this study is not cross-object generalization but resolution enhancement under controlled acquisition variations, the random distribution ensures that all experimental conditions are represented in each subset. Therefore, model evaluation reflects reconstruction performance across the defined thermal and geometric configurations rather than memorization of isolated scenarios.

7.2.1 Multi-resolution convolutional neural network training

The training of the proposed MR-CNN was carried out using paired low-resolution (LR), medium-resolution (MR), and high-resolution (HR) thermal images acquired under controlled laboratory conditions.

The network was trained for 50 epochs following a supervised learning strategy. A composite loss function was employed, combining pixel-wise reconstruction loss, perceptual constraints, and a temperature-domain error term to ensure radiometric consistency. Throughout training, several indicators were monitored to assess convergence and learning stability, including total loss, peak signal-to-noise ratio (PSNR), structural similarity (SSIM), and the absolute temperature error expressed as an L1 norm (Temp-L1).

In addition to quantitative metrics, qualitative evaluation was performed by visually inspecting representative thermal samples at different stages of training. Figure 7.4 illustrates examples of LR inputs, the corresponding super-resolved (SR) outputs generated by the network, and the HR reference images. These samples allow direct observation of how spatial structures and thermal gradients evolve through the super-resolution process. While the LR inputs exhibit limited spatial detail and strong quantization effects, the SR outputs progressively recover spatial continuity and thermal patterns that are more consistent with the HR targets, without introducing visually implausible artifacts.




Figure 7.4: Representative qualitative samples showing low-resolution input (LR), super-resolved output (SR), and high-resolution reference (HR) thermal images obtained during training.

The evolution of the training metrics is presented in Fig. 7.5. Over the 50 training epochs, the total loss decreases from 4.15 to 2.58, indicating a consistent reduction of the optimization objective throughout the learning process. This trend reflects the progressive minimization of both reconstruction and temperature-domain error components as the network parameters are updated. Simultaneously, the peak signal-to-noise ratio (PSNR) increases from 17.43 dB to 20.45 dB, evidencing an incremental improvement in spatial reconstruction accuracy. While PSNR values in thermal imagery are typically lower than those reported for visible-spectrum data, their monotonic increase confirms the convergence of the network toward a more accurate super-resolved representation.

The temperature-domain error follows a comparable trend throughout the training process. The Temp-L1 metric decreases from 1.38 °C at the initial epoch to 0.78 °C at the final epoch, remaining below 1 °C during the later stages of training. This behavior indicates that

the optimization process progressively reduces absolute temperature deviations between the reconstructed and reference images, without introducing discontinuities in the error evolution.

No abrupt changes, oscillatory patterns, or metric degradation are observed during training, suggesting numerically stable optimization under the selected training configuration. Validation and testing samples were obtained from independent acquisitions within the same experimental setup, allowing the monitored metric trends to reflect the network response to unseen thermal data rather than direct reuse of training samples.

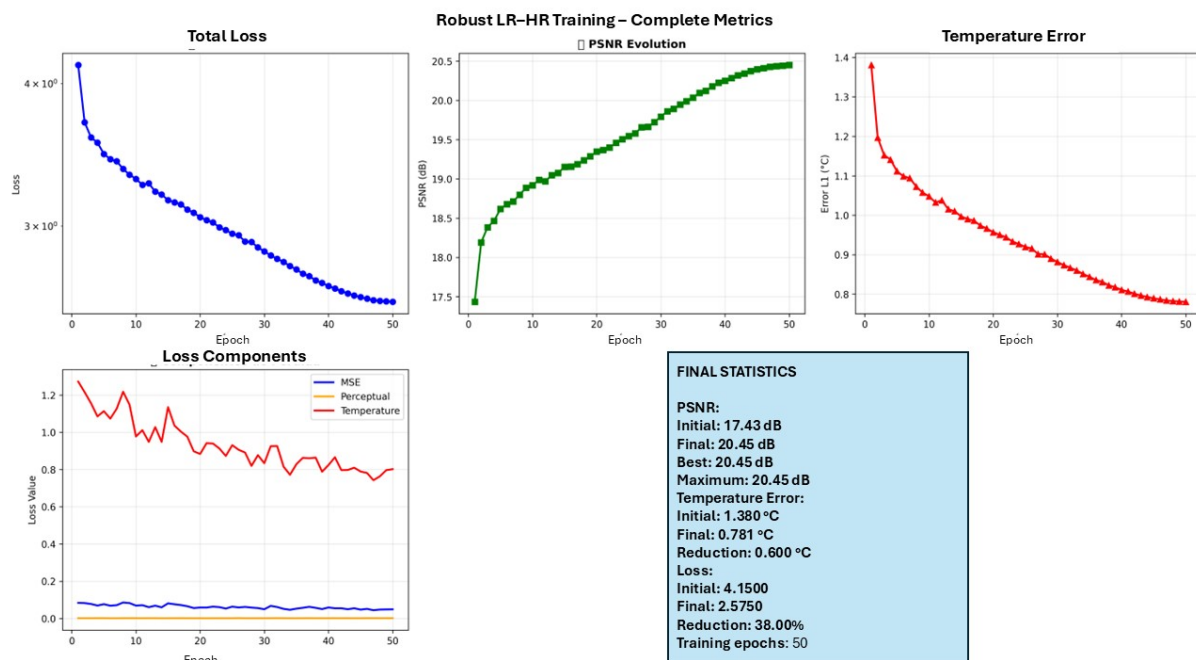


Figure 7.5: Evolution of training metrics for the MR-CNN, including total loss, PSNR, and temperature error (Temp-L1) over 50 epochs.

7.2.2 Multimodal generative adversarial network training

The training of the proposed multimodal generative adversarial network (GAN) was performed using paired thermal images at low resolution (LR), medium resolution (MR), and high resolution (HR), together with the corresponding RGB images acquired under identical experimental conditions. Prior to training, all modalities were spatially aligned to guarantee pixel-wise correspondence.

The training strategy follows an adversarial learning scheme in which a generator and a discriminator are optimized simultaneously. The generator reconstructs high-resolution thermal

images from multimodal inputs, while the discriminator distinguishes between real HR thermal images and generated outputs. This interaction enforces both temperature-domain consistency and realistic spatial structure in the reconstructed images.

The network was trained for a maximum of 15 epochs. The generator loss combines a pixel-wise L1 reconstruction loss with an adversarial loss term. The L1 component constrains absolute temperature deviations with respect to the HR reference images, while the adversarial term promotes realistic local thermal patterns. The discriminator is trained using a patch-based classification loss applied to real and generated HR samples.

During training, several metrics were monitored to evaluate convergence and stability. These metrics include the total generator loss, the individual loss components (L1 and adversarial), the discriminator loss, the peak signal-to-noise ratio (PSNR), and the absolute temperature error expressed as an L1 norm (Temp-L1). The evolution of these metrics over the training epochs is shown in Fig. 7.6.

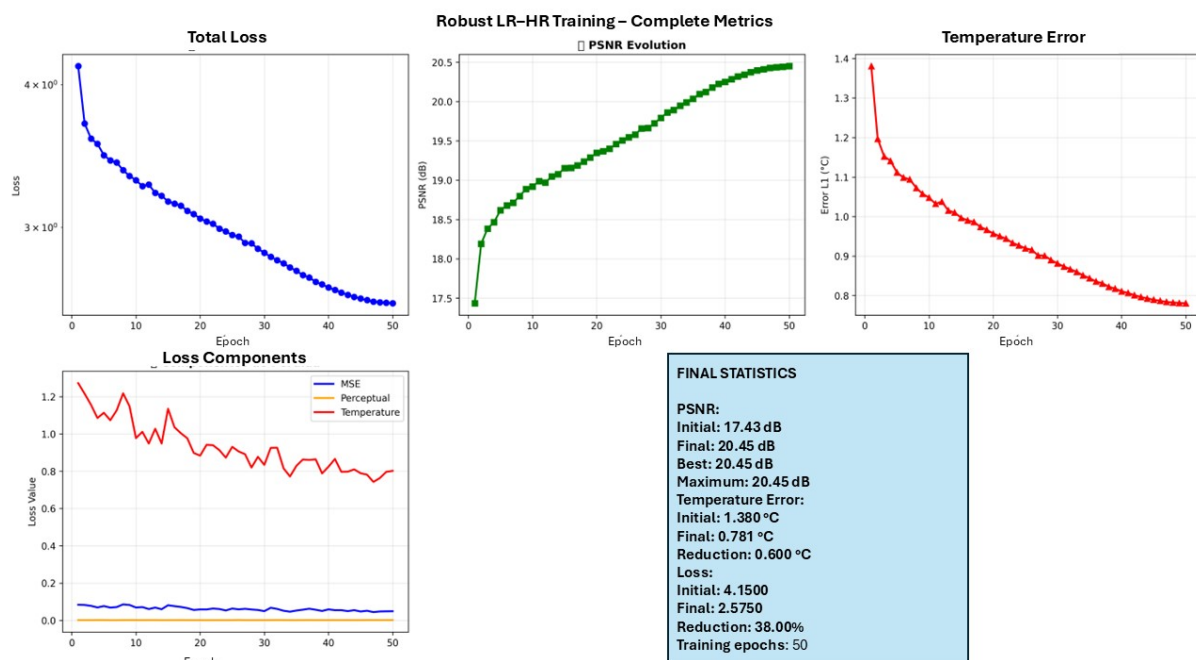


Figure 7.6: Evolution of training metrics for the multimodal GAN, including total loss, loss components, PSNR, and temperature error (Temp-L1) over the training epochs.

In the initial training epochs, the reduction of the reconstruction loss is accompanied by an increase in PSNR and a decrease in temperature-domain error. This behavior indicates that the generator learns consistent spatial and radiometric features from the multimodal inputs.

However, after approximately the 13th epoch, the PSNR evolution becomes non-monotonic and small oscillations appear in the adversarial loss terms. This behavior indicates the onset of adversarial instability, where further optimization leads to abrupt parameter updates in the generator.

To complement the numerical analysis, qualitative inspection of the reconstructed samples was performed. Representative examples of LR inputs, super-resolved (SR) outputs, and HR reference images are shown in Fig. 7.7.




Figure 7.7: Representative qualitative samples showing low-resolution input (LR), super-resolved output (SR), and high-resolution reference (HR) thermal images generated by the multimodal GAN during training.

The qualitative samples corresponding to later epochs reveal a progressive loss of local thermal contrast and reduced spatial variability, producing visually flatter thermal reconstructions. This effect is consistent with the observed PSNR degradation and reflects the dominance of the adversarial loss, which suppresses high-frequency thermal variations.

The stopping criterion was defined by the loss of monotonic behavior in PSNR and the onset of oscillations in the adversarial loss components. Beyond this point, no further improvements in radiometric consistency were observed, while subtle degradation in spatial detail became apparent. For this reason, early stopping was adopted as a practical criterion to stabilize the learning process and avoid over-smoothing effects, ensuring preservation of the thermal information.

7.2.3 Progressive multimodal variational autoencoder training

The training procedure of the progressive multimodal variational autoencoder (VAE) was carried out using paired low-resolution (LR) and high-resolution (HR) thermal images acquired under controlled laboratory conditions, covering a temperature range from 11.8 °C to 54.2 °C. As the architectural design and multimodal data flow were described previously, this subsection focuses on the training configuration, monitored variables, and learning behavior observed during optimization.

The model was trained for a total of 8 epochs using an unsupervised learning strategy with progressive resolution enhancement. The optimization objective was defined by a variational loss function composed of a reconstruction term, a Kullback–Leibler (KL) divergence regularization term, and a temperature-domain consistency constraint. This loss formulation jointly constrains image reconstruction fidelity, latent space regularization, and preservation of radiometric consistency during training. Although the KL divergence term is part of the optimization objective, its evolution is not reported explicitly, as the analysis focuses on the observable reconstruction behavior rather than latent space statistics.

Figure 7.8 summarizes the evolution of the monitored training metrics. The PSNR increased from 18.93 dB during the initial epoch to 28.13 dB by epoch 8. The progression of this metric reflects the adjustment of the encoder–decoder parameters and the latent distribution under the variational objective, providing a quantitative reference for tracking reconstruction consistency throughout training.

The training process was intentionally limited to 8 epochs based on the observed behavior of the monitored metrics. Beyond this point, PSNR variations between consecutive epochs became marginal, indicating that further optimization produced diminishing changes in reconstruction behavior under the imposed loss constraints. This criterion was adopted to avoid unnecessary parameter updates without measurable changes in reconstruction consistency.

Qualitative inspection was performed using representative reconstruction samples obtained at epoch 5. Figure 7.9 presents LR inputs alongside their corresponding VAE-generated reconstructions. These samples enable visual examination of spatial structure recovery and continuity

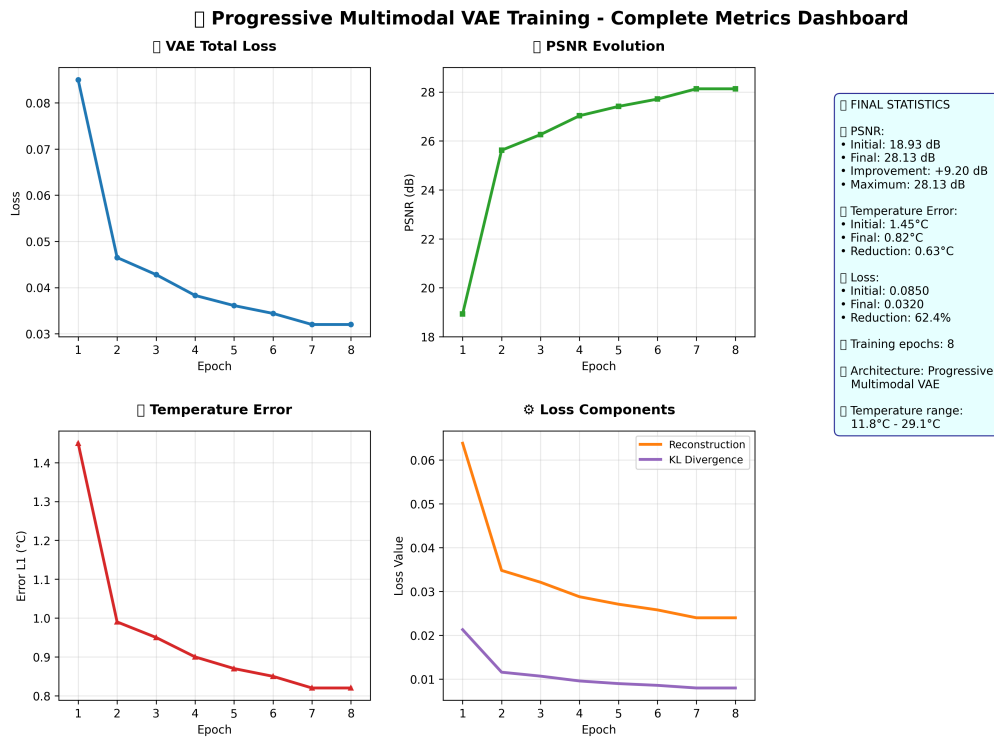


Figure 7.8: Evolution of reconstruction-related metrics during progressive VAE training, including PSNR across epochs.

of temperature gradients at an intermediate stage of the training process, complementing the numerical monitoring of reconstruction behavior.




Figure 7.9: Representative reconstruction samples obtained during progressive VAE training at epoch 5, showing LR inputs and corresponding reconstructed thermal images.

7.3 Comparative Performance Analysis

The experimental acquisition system provided radiometrically stable and repeatable thermal measurements, which ensured consistency across datasets and allowed a controlled comparison between models. However, this same stability constrained the spatial variability of the data, since the antenna and sample configuration remained fixed throughout all acquisitions. Consequently, the super-resolution problem does not evaluate the ability of the models to generalize across different spatial configurations, but rather their capacity to extract additional spatial information from a scene with limited structural variation. An additional constraint is introduced by the resolution gap between sensors, as the transition from low to medium resolution does not follow a progressive scaling factor such as $\times 2$ or $\times 4$, but instead represents an abrupt change. Under these conditions, the reconstruction task moves beyond simple interpolation, as the missing

spatial content cannot be directly inferred from neighboring pixels, imposing a common upper limit on the spatial fidelity achievable by all evaluated methods. Although temperature sweeps were introduced to increase thermal variability, the global geometry of the scene remained largely unchanged in this setup. Variations were mainly associated with viewing angle and acquisition distance. In practice, this leads to rapid convergence during training and early saturation of performance metrics. Architectures relying on pixel-wise losses or variational regularization exhibit this effect more clearly, as continued optimization beyond certain epochs does not introduce new structural information but instead reinforces repetitive spatial patterns.

To complement the qualitative analysis, a statistical evaluation of the quantitative metrics was performed over the independent testing subset. For each model, PSNR, SSIM, and Temp-L1 were computed individually for every test sample. Reported values correspond to the mean and standard deviation across the testing set, allowing assessment of inter-sample variability and reconstruction stability.

Table 7.1: Final quantitative performance over the independent testing subset (mean \pm standard deviation).

Model	PSNR (dB)	SSIM	Temp-L1 ($^{\circ}$ C)
MR-CNN	20.45 ± 0.52	0.8698 ± 0.018	0.78 ± 0.09
Multimodal GAN	21.29 ± 0.41	0.885 ± 0.022	0.92 ± 0.11
Progressive VAE	28.13 ± 0.34	0.945 ± 0.015	0.82 ± 0.08

The relatively low standard deviations observed across all architectures are consistent with the controlled acquisition protocol and limited geometric variability of the scene. Inter-sample dispersion remains below 0.6 dB in PSNR and below 0.12 $^{\circ}$ C in Temp-L1 for all models, indicating stable reconstruction behavior across the testing subset.

The temperature-domain results should be interpreted in light of the instrumental behavior characterized in Chapter 4. The inter-sensor analysis reported RMSE values between approximately 1.02 $^{\circ}$ C and 1.84 $^{\circ}$ C under controlled conditions. These values reflect the typical level of disagreement that can exist between thermal devices within the same experimental setup.

The Temp-L1 errors obtained by the super-resolution models (0.78–0.92 °C) are therefore within the same order of magnitude as the previously quantified inter-sensor deviations. This indicates that part of the residual reconstruction error may originate from measurement variability inherent to the acquisition system rather than solely from model limitations. For this reason, differences between models in the temperature domain that remain below the characterized instrumental deviation cannot be unambiguously attributed to model performance alone.

From a comparative perspective, the three evaluated strategies exhibit markedly different behaviors in terms of spatial informativeness rather than perceived visual smoothness. The convolutional baseline produces reconstructions in which the regions of interest remain spatially localized, but the recovered patterns appear diffuse and lack internal structure. While the overall thermal footprint of the antenna and the sample can be identified in the reconstructed images, localized gradients and boundaries are not sufficiently resolved to provide additional interpretable thermal information. As a result, the output remains spatially coherent but information-poor. The variational autoencoder exhibits an even stronger tendency toward spatial homogenization. The imposed latent regularization suppresses high-frequency variations, leading to reconstructions that are globally consistent in temperature but spatially flattened. Although this behavior reduces noise and improves average temperature accuracy, it effectively removes localized thermal features. Consequently, the reconstructed images convey limited actionable information, as spatial variations associated with heat concentration or distribution are largely smoothed out. In contrast, the multimodal GAN consistently recovers localized spatial structures that are meaningful in a thermal context. By leveraging RGB guidance and adversarial learning, the generator incorporates complementary spatial cues that allow the reconstruction of localized gradients and discontinuities corresponding to physically plausible thermal variations. These structures are not merely sharper edges in a visual sense, but regions where thermal contrast conveys additional information about heat distribution and spatial interaction. This distinction is particularly relevant in thermographic analysis, where informative content is associated with localized temperature variation rather than global smoothness. Although the GAN reconstructions may exhibit slightly higher local variability than those produced by the variational model, this variability reflects the recovery of spatial information rather than noise amplification. Under

the constraints imposed by the acquisition setup, the GAN therefore provides the most informative reconstructions, as it balances spatial detail with acceptable radiometric consistency. In comparison, the CNN and VAE outputs remain dominated by diffuse or averaged representations that, while stable, contribute limited additional insight for thermal interpretation. An additional constraint is imposed by the acquisition protocol. Due to hardware limitations in both the initial and revised systems, thermal images at different spatial resolutions were acquired sequentially rather than simultaneously. Although identical geometric conditions and controlled thermal excitation were maintained, temporal offsets between acquisitions introduce radiometric differences associated with sensor response dynamics and residual thermal transients. Consequently, LR, MR, and HR images correspond to closely related but not temporally synchronized thermal states. This sequential acquisition introduces a limitation in the super-resolution task, as the models are trained to learn mappings between images that are physically consistent but not identical. Under these conditions, a portion of the reconstruction error is attributable to acquisition-related effects rather than model behavior. In this setting, the multi-modal GAN shows lower sensitivity to radiometric misalignment, while the CNN and VAE are more affected by repeated spatial patterns and temporal inconsistencies. These factors define the operational bounds of super-resolution performance within the experimental framework and support the interpretation of the comparative results.

Chapter 8

Scientific Outputs and Derived Publications

During the development of this thesis, several scientific outputs directly related to the project were generated, including journal articles and conference presentations, which document the methodology, experimental results, and technical developments achieved.

Journal Articles:

- Ortiz-Santana, P. (co-author) et al., “*Capturing Material Thermal Responses under RF-Induced Hyperthermia: An Experimental Imaging Database*”, F1000Research, 2025, 14:1379.
- Reyes-Vera, E. et al., “*Multiphysical Characterization of a Tissue-Mimicking Phantom: Composition, Thermal Behavior, and Broadband Electromagnetic Properties from Visible to Terahertz and Microwave Frequencies*”, under review (Paula Ortiz-Santana, co-author).

Conference Presentations:

- “*Evaluation of Classical Image Processing Techniques for the Correction, Registering, and Fusion of RGB and Thermal Images*”, 11a Semana Internacional de Ciencia, Tecnología e Innovación, Cúcuta, Colombia, Oct. 2024. Role: Presenter.
- “*Sistema de Adquisición y Control para el Análisis Térmico Inducido por Radiación de Microondas*”, VIII Congreso Nacional de Ingeniería Física, Popayán, Colombia, Oct.

2025. Role: Presenter.

- “*Base de Imágenes Térmicas para Aplicar Técnicas de Aprendizaje Automático y Mejorar su Estudio*”, Xpoilers, Bogotá D.C., Colombia, May. 2025. Role: Presenter.
- “*User Interface for Displaying and Controlling Temperature Based on Thermal Cameras and RGB Cameras for Hyperthermia Therapies*”, 11a Semana Internacional de Ciencia, Tecnología e Innovación, Cúcuta, Colombia, Oct. 2024. Role: Presenter.

Chapter 9

Conclusions and Future Work

9.1 Conclusions

This thesis focuses on the experimental evaluation of an integrated acquisition, processing, and control system aimed at improving the spatial resolution and temperature consistency of thermal images obtained from low- and medium-cost infrared sensors. While the proposed hypothesis anticipated that the combination of controlled multimodal acquisition, noise reduction techniques, and machine learning models trained with proprietary data would lead to substantial improvements, the experimental results indicate that these improvements were partial and strongly dependent on both hardware constraints and processing stages.

Quantitatively, the proposed framework achieved a reduction in temperature error (RMSE) from 1.45 °C to 0.82 °C under controlled acquisition conditions, indicating a measurable improvement expressed directly in degrees Celsius rather than solely perceptual enhancement.

The general objective of developing an integrated thermal image acquisition and processing system was achieved. The proposed system successfully combined thermal and RGB sensors, a controlled acquisition protocol, signal conditioning, and post-processing algorithms, enabling a structured evaluation of spatial and thermal performance. In particular, the acquisition architecture, database construction, and closed-loop control implementation fulfilled their intended objectives and provided a solid experimental platform for analysis and comparison.

One of the main strengths of this work lies in the acquisition and control architecture, where

a host PC acted as the central unit for synchronization, thermal monitoring, control execution, and data storage. This architecture enabled stable closed-loop temperature regulation and reliable data collection under controlled conditions. The control experiments demonstrated that conventional PID and hybrid iPID strategies are capable of regulating temperature in RF-induced hyperthermia scenarios within a defined operating range. The inclusion of the hybrid iPID approach was motivated by the observed behavior of the conventional PID when applied to a system that is only partially linear and exhibits material-dependent dynamics. Despite the presence of nonlinear behavior in the RF actuator, mainly associated with power–temperature coupling, the iPID-based strategies showed improved stability compared to conventional tuning approaches. In this sense, the control stage and the overall system integration can be considered one of the most consistent outcomes of the thesis.

In contrast, the results obtained from the machine learning-based image enhancement stage did not fully meet the initial expectations. Although certain architectures produced measurable improvements in spatial metrics and visual structure, the enhanced thermal images did not consistently preserve temperature information with the desired accuracy. The improvements observed were therefore incremental rather than definitive, indicating that the learning-based models were limited by both the quality of the input data and the complexity of the multimodal registration process. Despite these limitations, the machine learning stage was retained as part of the system in order to systematically evaluate its actual contribution under realistic experimental constraints rather than idealized conditions.

It is important to emphasize that the implemented architectures were computationally robust and demanded substantial training resources. Each deep learning model required approximately four days of training on average, with longer convergence times observed in probabilistic models such as the Variational Autoencoder (VAE). In some cases, training extended beyond this average duration depending on hyperparameter configuration and convergence stability. This computational cost represents a practical limitation of the proposed approach in terms of scalability and iterative experimentation.

A critical factor influencing these results is the non-simultaneous nature of the thermal and RGB acquisition. The lack of strict temporal synchronization increased the difficulty of the reg-

istration stage, particularly under dynamic thermal conditions. For this reason, the registration process required additional constraints and manual validation steps to avoid spatial misalignment that could propagate errors into the learning-based enhancement stage, and introduced uncertainty that could not be fully compensated by post-processing. Additionally, the large resolution gap between the high-megapixel RGB camera and the low-resolution thermal sensor significantly increased computational complexity and made feature alignment more sensitive to small spatial errors. Rather than simplifying the learning task, this imbalance imposed strong constraints on both the registration algorithms and the subsequent super-resolution models.

This resolution imbalance also increased memory consumption and overall computational load during training, reinforcing the dependence of the framework on high-performance GPU hardware.

The inter-sensor comparison confirmed that low-cost thermal cameras exhibit higher noise, bias, and variability compared to reference systems such as FLIR and Optris cameras. Although calibration and controlled acquisition allowed these effects to be characterized, they could not be entirely eliminated. As a result, the machine learning models implicitly learned sensor-specific distortions along with structural features, which partially explains the limited generalization and the incomplete preservation of thermal accuracy after enhancement.

The constructed database remains a central contribution of this work. Its controlled nature, inclusion of multiple materials and resolutions, and direct linkage to the acquisition hardware enabled a fair comparison of machine learning architectures. Importantly, this database made it possible to evaluate not only visual improvement but also temperature-related performance, revealing limitations that would not be evident from perceptual metrics alone.

9.2 Recommendations and Future Research Lines

- Expand the thermal image database by incorporating a wider variety of materials, temperature ranges, and acquisition conditions, in order to improve the generalization capability of machine learning models.
- Design machine learning architectures specifically tailored for thermal super-resolution

that explicitly incorporate temperature-domain constraints, prioritizing thermal information preservation over purely visual enhancement.

- Evaluate simultaneous acquisition strategies between thermal and RGB sensors to reduce the complexity of the multimodal registration stage and minimize spatial alignment errors.
- Analyze the impact of reducing the resolution gap between sensors by incorporating intermediate thermal cameras or alternative acquisition configurations, with the aim of simplifying model training and validation.
- Explore nonlinear, adaptive, or model-based control strategies that explicitly account for the nonlinear dynamics of the RF actuator and material-dependent thermal behavior.
- Migrate the acquisition, processing, and control system to embedded or edge-computing platforms to enable real-time operation without reliance on a host PC.

Bibliography

- [1] Q. Liu, M. Li, W. Wang, S. Jin, H. Piao, Y. Jiang, N. Li, and H. Yao, “Infrared thermography in clinical practice: a literature review,” *European Journal of Medical Research* 2025 30:1, vol. 30, pp. 33–, 1 2025. [Online]. Available: <https://link.springer.com/article/10.1186/s40001-025-02278-z>
- [2] E. F. J. Ring and K. Ammer, “The technique of infrared imaging in medicine,” in *Infrared Imaging: A Casebook in Clinical Medicine*. IOP Publishing, 2015.
- [3] S. Bagavathiappan, T. Saravanan, J. Philip, T. Jayakumar, B. Raj, R. Karunanithi, T. M. Panicker, M. Korath, and K. Jagadeesan, “Infrared thermal imaging for detection of peripheral vascular disorders,” *Journal of Medical Physics*, vol. 34, pp. 43–47, 1 2009. [Online]. Available: https://journals.lww.com/jomp/fulltext/2009/34010/infrared_thermal_imaging_for_detection_of.7.aspx
- [4] B. B. Lahiri, S. Bagavathiappan, T. Jayakumar, and J. Philip, “Medical applications of infrared thermography: A review,” *Infrared Physics Technology*, vol. 55, pp. 221–235, 7 2012. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1350449512000308>
- [5] R. E. Rivadeneira, A. D. Sappa, and B. X. Vintimilla, *Thermal Image Super-resolution: A Novel Architecture and Dataset*, 2020. [Online]. Available: <https://orcid.org/0000-0003-2468-0031>
- [6] Z. Wang, J. Chen, and S. C. H. Hoi, “Deep learning for image super-resolution: A survey,” *arXiv preprint arXiv:1902.06068*, 2019. [Online]. Available: <http://arxiv.org/abs/1902.06068>

- [7] R. E. Rivadeneira, A. D. Sappa, B. X. Vintimilla, S. Nathan, P. Kansal, A. Mehri, B. Ardakani, A. Dalal, A. Akula, D. Sharma, S. Pandey, B. Kumar, J. Yao, R. Wu, K. Feng, N. Li, Y. Zhao, H. Patel, V. Chudasama, K. Prajapati, A. Sarvaiya, K. P. Upla, K. Raja, R. Ramachandra, C. Busch, F. Almasri, T. Vandamme, O. Debeir, N. B. Gutierrez, Q. H. Nguyen, and W. J. Beksi, “Thermal image super-resolution challenge – pbvs 2021,” <https://pbvs-workshop.github.io/challenge.html>, 2021, pBVS 2021 Workshop Challenge.
- [8] C. Arnold, P. Jouvét, and L. Seoud, “Swinfusr: An image fusion-inspired model for rgb-guided thermal image super-resolution,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. IEEE Computer Society, 4 2024, pp. 3027–3036. [Online]. Available: <https://arxiv.org/pdf/2404.14533>
- [9] Y. Jiang, Y. Liu, W. Zhan, W. Flores-Fuentes, J. C. Rodríguez-Quinonez, J. E. Miranda-Vega, Y. Jiang, Y. Liu, W. Zhan, and D. Zhu, “Improved thermal infrared image super-resolution reconstruction method based on multimodal sensor fusion,” *Entropy*, vol. 25, no. 6, p. 914, 2023.
- [10] X. Chen, L. Chen, L. Chen, P. Chen, G. Sheng, X. Yu, and Y. Zou, “Modeling thermal infrared image degradation and real-world super-resolution under background thermal noise and streak interference,” *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 34, pp. 6194–6206, 7 2024.
- [11] V. M. Oliveira and A. H. J. Moreira, “Edge ai system using a thermal camera for industrial anomaly detection,” in *Artificial Intelligence and IoT for Sustainable Development*. Springer, 2022, pp. 172–187.
- [12] A. Hakim and R. N. Awale, “Thermal imaging - an emerging modality for breast cancer detection: A comprehensive review,” *Journal of Medical Systems*, vol. 44, p. 136, 8 2020.
- [13] P. Ortiz-Santana, J. Herrera-Ramirez, E. Reyes-Vera, C. Furnieles-Chipagra, J. Botero-Valencia, and J. Araque, “Capturing material thermal responses under rf-induced hyperthermia: An experimental imaging database,” *F1000Research*, vol. 14, p. 1379, 12 2025. [Online]. Available: <https://f1000research.com/articles/14-1379>

- [14] M. W. Dewhirst, Z. Vujaskovic, E. Jones, and D. Thrall, “Re-setting the biologic rationale for thermal therapy,” *International Journal of Hyperthermia*, vol. 21, pp. 779–790, 12 2005.
- [15] M. H. Falk and R. D. Issels, “Hyperthermia in oncology,” *International Journal of Hyperthermia*, vol. 17, pp. 1–18, 1 2001.
- [16] M. Urano and E. Double, *Interstitial Hyperthermia: Physics, Biology and Clinical Aspects*. CRC Press, 1 2023.
- [17] K. Stowe, “Blackbody radiation,” *An Introduction to Thermodynamics and Statistical Mechanics*, pp. 438–456, 6 2007. [Online]. Available: <https://www.cambridge.org/core/books/an-introduction-to-thermodynamics-and-statistical-mechanics/blackbody-radiation/6CDBDF783FB84361A2F51062C21FF9E2>
- [18] S. M. Stewart and R. B. Johnson, *Blackbody Radiation*. CRC Press, 9 2016.
- [19] R. Usamentiaga, P. Venegas, J. Guerediaga, L. Vega, J. Molleda, and F. G. Bulnes, “Infrared thermography for temperature measurement and non-destructive testing,” *Sensors (Basel, Switzerland)*, vol. 14, p. 12305, 7 2014. [Online]. Available: <https://pmc.ncbi.nlm.nih.gov/articles/PMC4168422/>
- [20] G. Kramm and N. M. Olders, “Planck’s blackbody radiation law: Presentation in different domains and determination of the related dimensional constants,” *Journal of the Calcutta Mathematical Society*, 1 2009. [Online]. Available: <https://arxiv.org/pdf/0901.1863>
- [21] J. M. Marr and F. P. Wilkin, “A better presentation of planck’s radiation law,” *American Journal of Physics*, vol. 80, pp. 399–405, 3 2012. [Online]. Available: <http://arxiv.org/abs/1109.3822><http://dx.doi.org/10.1119/1.3696974>
- [22] G. Kramm, F. Herbert, and J. W. Goethe, “Heuristic derivation of blackbody radiation laws using principles of dimensional analysis,” 1 2008. [Online]. Available: <https://arxiv.org/pdf/0801.2197>

- [23] E. Daneshvar, G. Finlayson, M. H. Brill, and R. Deeb, “Introducing a temperature adjustment to make wien’s law a more accurate approximation of planckian blackbody radiation in the visible range,” *Optics Express*, vol. 33, p. 11956, 3 2025.
- [24] L. Reggiani and E. Alfinito, “Stefan-boltzmann law revisited,” 10 2021. [Online]. Available: <https://arxiv.org/pdf/2112.12090>
- [25] E. Barreira, R. M. Almeida, and M. L. Simões, “Emissivity of building materials for infrared measurements,” *Sensors 2021, Vol. 21, Page 1961*, vol. 21, p. 1961, 3 2021. [Online]. Available: <https://www.mdpi.com/1424-8220/21/6/1961/html><https://www.mdpi.com/1424-8220/21/6/1961>
- [26] N. OGASAWARA, J. SAITO, and H. YAMADA, “Infrared emissivity measurement based on polarized reflection characteristics of non-transmissive insulators,” *Mechanical Engineering Journal*, vol. 11, pp. 24–00 115, 2024.
- [27] M. Planck, *The Theory of Heat Radiation*, 1914.
- [28] M. F. Modest, *Radiative Heat Transfer*. Academic Press, 2013.
- [29] FLIR Systems, *Infrared Thermography Handbook*, 2011.
- [30] S. P. Garaba, T. Acuña-Ruz, and C. B. Mattar, “Hyperspectral longwave infrared reflectance spectra of naturally dried algae, anthropogenic plastics, sands and shells,” *Earth System Science Data*, vol. 12, pp. 2665–2678, 11 2020.
- [31] F. Desta, M. Buxton, and J. Jansen, “Fusion of mid-wave infrared and long-wave infrared reflectance spectra for quantitative analysis of minerals,” *Sensors (Basel, Switzerland)*, vol. 20, p. 1472, 3 2020. [Online]. Available: <https://pmc.ncbi.nlm.nih.gov/articles/PMC7085633/>
- [32] B. B. Turgut, G. G. Artan, and A. Bek, “A comparison of mwir and lwir imaging systems with regard to range performance,” p. 38, 4 2018. [Online]. Available: <https://open.metu.edu.tr/handle/11511/40859>

- [33] P. Burgholzer, “Thermodynamic limits of spatial resolution in active thermography,” *International Journal of Thermophysics*, vol. 36, pp. 2328–2341, 4 2015. [Online]. Available: <http://arxiv.org/abs/1502.01044><http://dx.doi.org/10.1007/s10765-015-1890-7>
- [34] J. G. Sun, “Analysis of pulsed thermography methods for detect depth prediction,” *Journal of Heat Transfer*, vol. 128, pp. 329–338, 4 2006.
- [35] C. Maierhofer, R. Arndt, M. Röllig, C. Rieck, A. Walther, H. Scheel, and B. Hillemeier, “Application of impulse-thermography for non-destructive assessment of concrete structures,” *Cement and Concrete Composites*, vol. 28, pp. 393–401, 4 2006.
- [36] L. Vergani, C. Colombo, and F. Libonati, “A review of thermographic techniques for damage investigation in composites,” *Frattura ed Integrità Strutturale*, vol. 8, pp. 1–12, 2014.
- [37] E. L. Dereniak and G. D. Boreman, “Infrared detectors and systems,” p. 561, 1996. [Online]. Available: <https://www.wiley.com/en-us/Infrared+Detectors+and+Systems-p-9780471122098>
- [38] A. Rogalski, “Infrared detectors: status and trends,” *Progress in Quantum Electronics*, vol. 27, pp. 59–210, 1 2003. [Online]. Available: <https://www.sciencedirect.com/science/article/abs/pii/S0079672702000241>
- [39] M. Vollmer and K. P. Möllmann, “Infrared thermal imaging: Fundamentals, research and applications,” *Infrared thermal imaging: Fundamentals, research and applications*, pp. 1–769, 10 2017. [Online]. Available: <https://onlinelibrary.wiley.com/doi/book/10.1002/9783527693306>
- [40] R. G. Driggers, M. H. Friedman, and J. W. Devitt, “Introduction to infrared and electro-optical systems,” 2022.
- [41] W. Minkina and S. Dudzik, “Infrared thermography: Errors and uncertainties,” *Infrared Thermography: Errors and Uncertainties*, pp. 1–192, 10 2009. [Online]. Available: <https://onlinelibrary.wiley.com/doi/book/10.1002/9780470682234>

- [42] J. Zapata-Londoño, J. Botero-Valencia, A. Torres, E. Reyes-Vera, and R. Hernández-García, “Automated 3d phenotyping of maize plants: Stereo matching guided by deep learning,” *Agriculture*, vol. 15, no. 24, 2025. [Online]. Available: <https://www.mdpi.com/2077-0472/15/24/2573>
- [43] J. Zapata-Londoño, J. Botero-Valencia, V. García-Pineda, E. Reyes-Vera, and R. Hernández-García, “A comprehensive review of optical and ai-based approaches for plant growth assessment,” *Agronomy*, vol. 15, no. 8, 2025. [Online]. Available: <https://www.mdpi.com/2073-4395/15/8/1781>
- [44] J. Soto-Perdomo, E. Reyes-Vera, J. Montoya-Cardona, and P. Torres, “Experimental dataset of tunable mode converter based on long-period fiber gratings written in few-mode fiber: Impacts of thermal, wavelength, and polarization variations,” *Data*, vol. 9, no. 1, 2024. [Online]. Available: <https://www.mdpi.com/2306-5729/9/1/10>
- [45] Y. Yu, B. G. Lee, M. Pike, Q. Zhang, and W. Y. Chung, “Deep learning-based rgb–thermal image denoising: review and applications,” *Multimedia Tools and Applications*, vol. 83, no. 4, pp. 11 613–11 641, 2024.
- [46] Z. Behrouzkia, Z. Joveini, B. Keshavarzi, N. Eyvazzadeh, and R. Z. Aghdam, “Hyperthermia: How can it be used?” *Oman Medical Journal*, vol. 31, pp. 89–97, 3 2016.
- [47] N. Datta, S. G. Ordóñez, U. Gaipl, M. Paulides, H. Crezee, J. Gellermann, D. Marder, E. Puric, and S. Bodis, “Local hyperthermia combined with radiotherapy and/or chemotherapy: Recent advances and promises for the future,” *Cancer Treatment Reviews*, vol. 41, pp. 742–753, 11 2015.
- [48] C. Furnieles, C. Zapata-Hernández, J. Botero-Valencia, E. Reyes-Vera, and J. Araque, “Material characterization for phantoms in microwave sar evaluation,” in *2024 IEEE 1st Latin American Conference on Antennas and Propagation (LACAP)*, 2024, pp. 1–2.
- [49] H. Sahinbas, M. Rosch, and M. Demiray, “Temperature measurements in a capacitive system of deep loco-regional hyperthermia,” *Electromagnetic Biology and Medicine*, vol. 36, pp. 248–258, 7 2017.

- [50] J. H. Kim, E. W. Hahn, and S. A. Ahmed, "Combination hyperthermia and radiation therapy for malignant melanoma," *Cancer*, vol. 50, pp. 478–482, 8 1982.
- [51] J. Overgaard and M. Overgaard, "Hyperthermia as an adjuvant to radiotherapy in the treatment of malignant melanoma," *International Journal of Hyperthermia*, vol. 3, pp. 483–501, 1987. [Online]. Available: <https://www.tandfonline.com/doi/abs/10.3109/02656738709140422>
- [52] J. Overgaard, S. M. Bentzen, D. G. Gonzalez, M. C. Hulshof, G. Arcangeli, O. Dahl, and O. Mella, "Randomised trial of hyperthermia as adjuvant to radiotherapy for recurrent or metastatic malignant melanoma," *The Lancet*, vol. 345, pp. 540–543, 3 1995. [Online]. Available: [https://www.thelancet.com/action/showFullText?pii=S0140673695904638https://www.thelancet.com/action/showAbstract?pii=S0140673695904638https://www.thelancet.com/journals/lancet/article/PIIS0140-6736\(95\)90463-8/abstract](https://www.thelancet.com/action/showFullText?pii=S0140673695904638https://www.thelancet.com/action/showAbstract?pii=S0140673695904638https://www.thelancet.com/journals/lancet/article/PIIS0140-6736(95)90463-8/abstract)
- [53] P. Wust, B. Hildebrandt, G. Sreenivasa, B. Rau, J. Gellermann, H. Riess, R. Felix, and P. Schlag, "Hyperthermia in combined treatment of cancer," *Lancet Oncology*, vol. 3, pp. 487–497, 8 2002. [Online]. Available: [https://www.thelancet.com/action/showFullText?pii=S1470204502008185https://www.thelancet.com/action/showAbstract?pii=S1470204502008185https://www.thelancet.com/journals/lanonc/article/PIIS1470-2045\(02\)00818-5/abstract](https://www.thelancet.com/action/showFullText?pii=S1470204502008185https://www.thelancet.com/action/showAbstract?pii=S1470204502008185https://www.thelancet.com/journals/lanonc/article/PIIS1470-2045(02)00818-5/abstract)
- [54] A. A. Skaptsov, S. O. Ustalkov, A. H. M. Mohammed, O. A. Savenko, A. S. Novikova, E. A. Kozlova, and V. I. Kochubey, "Fabrication and characterization of biological tissue phantoms with embedded nanoparticles," *Journal of Physics: Conference Series*, vol. 917, p. 042003, 11 2017. [Online]. Available: <https://iopscience.iop.org/article/10.1088/1742-6596/917/4/042003>
- [55] A. Dabbagh, B. J. J. Abdullah, N. H. A. Kasim, and C. Ramasindarum, "Reusable heat-sensitive phantom for precise estimation of thermal profile in hyperthermia application," *International Journal of Hyperthermia*, vol. 30, pp. 66–74, 2 2014. [Online]. Available: <http://www.tandfonline.com/doi/full/10.3109/02656736.2013.854930>

- [56] J. Song, X. Li, S. Zhang, Y. Li, L. Chen, Z. Wang, G. Cheng, Q. Yu, and Y. Han, “Ginkgo biloba leaf extract: A natural and eco-friendly stabilizer for enhancing the thermal and oxidative stability of polyethylene,” *Journal of Polymer Research*, vol. 31, 11 2024.
- [57] T. Kurose, T. Takahashi, and K. Koyama, “A new process to make a porous ptfе structure from aqueous ptfе dispersion with the help of hydrogel,” Tech. Rep., 2004.
- [58] I. Lorenzo, M. Serra-Prat, and J. C. Yébenes, “The role of water homeostasis in muscle function and frailty: A review,” *Nutrients*, vol. 11, p. 1857, 8 2019.
- [59] K. Mostakim, M. R. Akbar, M. A. Islam, and M. K. Islam, “Integrated photovoltaic-thermal system utilizing front surface water cooling technique: An experimental performance response,” *Heliyon*, vol. 10, 2 2024.
- [60] D. Lilley, P. Yu, J. Ma, A. Jain, and R. Prasher, “Thermal fluids with high specific heat capacity through reversible diels-alder reactions,” *iScience*, vol. 25, 1 2022.
- [61] J. Botero-Valencia, E. Reyes-Vera, E. Ospina-Rojas, and F. Prieto-Ortiz, “A portable tool for spectral analysis of plant leaves that incorporates a multichannel detector to enable faster data capture,” *Instruments*, vol. 8, no. 1, 2024. [Online]. Available: <https://www.mdpi.com/2410-390X/8/1/24>
- [62] E. Reyes Vera, J. Montoya-Cardona, V. García Pineda, A. Valencia-Arias, N. Gomez-Cardona, F. Lopez-Giraldo, and E. Gonzalez-Valencia, “Advances in THz spectroscopy: A systematic review of material characterization applications,” *ASEAN Journal of Science and Engineering*, vol. 6, no. 2, pp. 173–208, 2026.
- [63] S. Saxena, C. Bench, D. Garg, P. Boardman, M. Mrnka, H. Penketh, N. Stone, and E. Hendry, “Limitations of effective medium models for tissue phantoms in the THz frequency range,” *Scientific Reports*, vol. 14, p. 22968, 2024.
- [64] M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu, “Spatial transformer networks,” *Advances in Neural Information Processing Systems*, vol. 2015-January, pp. 2017–2025, 6 2015. [Online]. Available: <https://arxiv.org/pdf/1506.02025>

- [65] P. Viola and W. M. Wells, "Alignment by maximization of mutual information," *International Journal of Computer Vision*, vol. 24, pp. 137–154, 1997. [Online]. Available: <https://link.springer.com/article/10.1023/A:1007958904918>
- [66] G. Balakrishnan, A. Zhao, M. R. Sabuncu, J. Guttag, and A. V. Dalca, "Voxelmorph: A learning framework for deformable medical image registration," *IEEE Transactions on Medical Imaging*, vol. 38, pp. 1788–1800, 9 2019. [Online]. Available: <http://arxiv.org/abs/1809.05231><http://dx.doi.org/10.1109/TMI.2019.2897538>
- [67] MathWorks, "System identification for pid control - matlab & simulink," 2023, available at MathWorks Documentation.
- [68] K. J. Astrom and T. Hagglund, *PID Controllers: Theory, Design, and Tuning*, 2nd ed. ISA, 1995.
- [69] K. J. Åström and T. Hägglund, *PID Controllers: Theory, Design, and Tuning*, 2nd ed. Research Triangle Park, NC: ISA, 1995.
- [70] K. J. Hunt and D. Sbarbaro, "Neural networks for nonlinear internal model control," *IEE Proceedings D - Control Theory and Applications*, vol. 138, no. 5, pp. 431–438, 1991.
- [71] K. H. Ang, G. Chong, and Y. Li, "Pid control system analysis, design, and technology," *IEEE Transactions on Control Systems Technology*, vol. 13, no. 4, pp. 559–576, 2005.