

Article

3D Thermal Imaging System with Decoupled Acquisition for Industrial and Cultural Heritage Applications

Ivo Campione ^{1,*}, Francesca Lucchi ¹, Nicola Santopuoli ² and Leonardo Seccia ³

¹ Department of Industrial Engineering, University of Bologna, Via Fontanelle 40, 47121 Forlì (FC), Italy; f.lucchi@unibo.it

² Department of History, Representation and Restoration of Architecture (Italian acronym DSDRA), Sapienza University of Rome, Piazza Borghese 9, 00186 Rome, Italy; nicola.santopuoli@uniroma1.it

³ Department of Mathematics and CAILab (Archeo-Engineering Laboratory), University of Bologna, Via Fontanelle 40, 47121 Forlì (FC), Italy; leonardo.seccia@unibo.it

* Correspondence: ivo.campione@unibo.it

Received: 18 December 2019; Accepted: 22 January 2020; Published: 23 January 2020

Abstract: Three-dimensional thermography is a recent technique—with various fields of application—that consists of combining thermography with 3D spatial data in order to obtain 3D thermograms, high information objects that allow one to overcome some limitations of 2D thermograms, to enhance the thermal monitoring and the detection of abnormalities. In this paper we present an integration methodology that can be applied to merge data acquired from a generic thermal camera and a generic laser scanner, and has the peculiarity of keeping the two devices completely decoupled and independent, so that thermal and geometrical data can be acquired at different times and no rigid link is needed between the two devices. In this way, the stand-alone capability of each device is not affected, and the data fusion is applied only when necessary. In the second part, the real effectiveness of our approach is tested on a 3D-printed object properly designed. Furthermore, one example of an application of our methodology in the cultural heritage field is presented, with an eye to preservation and restoration: the integration is applied to a marble statue called Madonna with the Child, a fine work of the Florentine sculptor Agostino di Duccio (1418–1481). The results suggest that the method can be successfully applicable to a large set of scenarios. However, additional tests are needed to improve the robustness.

Keywords: 3D thermography; thermal imaging; laser scanning; integration methodology; extrinsic calibration; decoupled acquisition; heritage conservation

1. Introduction

All bodies at temperature above absolute zero (0 K) emit electromagnetic radiation. If the temperature of a body has the same magnitude as the ambient temperature, then the emission is mainly relegated to the infrared (IR) range of the spectrum and can be sensed and displayed by a thermal camera as a false-color image, called a thermogram. In addition, by knowing a series of parameters such as the inspected object emissivity and the apparent reflected temperature (i.e., the ambient temperature), an approximated temperature map of the object can be computed in output. Infrared thermography is a non-invasive (non-contact and non-destructive) imaging method, which makes it a widely applicable technique. For example, it has a vast range of applications in research and industry, building and infrastructure, electrical installation inspection, microsystems engineering, but also in biology, medicine, life sciences and cultural heritage [1,2].

Like thermal imaging, 3D reconstruction techniques are nowadays widespread in many different fields and are commonly used to acquire the object geometry and to provide easy 3D

documentation. They are a powerful tool to improve the identification, monitoring, conservation and restoration of objects and structures [3].

The utility of the integration is largely due to the fact that the ability to combine both temperature and geometric data together can lead to several advantages: it enhances and speeds up the interpretation of the results; it offers the possibility to select the region of interest by taking into account the geometry; it allows the easy segmentation of the 2D data from the background. One of the most important advantages, however, is that it allows one to overcome a significant limitation of 2D thermograms, namely the systematic error in the measured temperature due to the dependence of the emissivity on the viewing angle [4–6].

Several works in the literature have shown the strong potential of 3D thermal mapping (commonly known in the literature as 3D thermography). The thermal data are obtained with a thermal camera, whereas the way in which 3D data are acquired varies: a concise overview can be found in the work of G. Chernov et al. [4]. For example, in the medical field some applications of 3D thermography are the work of X. Ju et al. [7] and the one of K. Skala et al. [8]. In [7], the process of 3D capture relied upon stereo photogrammetry, whereas in [8] the system consisted of a high-resolution offline 3D laser scanner and a real-time low-resolution 3D scanner, both paired together with a thermal imaging camera, for human body 3D thermal model comparison and analysis. In [9], S. Vidas and P. Moghadam presented “HeatWave”, a handheld, low-cost 3D thermography system, which allows non-experts to generate detailed 3D surface temperature models for energy auditing. The core technology of this device is obtained by combining a thermal camera and an RGB-D sensor (depth sensing device coupled with an RGB camera). In several other latest applications, such as [10–12], the spatial data were obtained by using a depth camera such as the Microsoft Kinect, which has become one of the top choices for 3D thermography, because of its large versatility and the capability to be exploited to perform real-time integrations. In [13], the integration was carried out on the data acquired by two smartphones arranged in a stereo configuration and a thermal camera. In [14], a fully automatic system that generated 3D thermal models of indoor environments was presented; it consisted of a mobile platform equipped with a 3D laser scanner, an RGB camera and a thermal camera.

In the cultural heritage field, spatial and multispectral data have usually been fused together for documentation reasons, historical studies, restoration plans and visualization purposes; several examples can be found in [15–19].

One advantage of 3D thermal models is that, for each 3D point, one can compute the so-called viewing angle (i.e., the angle between the surface normal vector in that point and the vector joining the point and the optical center). This information can be used to correct the error in the temperature caused by the dependence of the emissivity on the viewing angle. Indeed, for a given material, the emissivity is usually not constant, but depends on several factors, such as the surface condition, the wavelength, the temperature, the presence of concavities and the viewing angle (a viewing angle-dependent emissivity is often called “directional emissivity”). A detailed explanation of how these factors affect the emissivity can be found in [1] (pp. 35–45). Whereas the role of many of these factors can be in general considered negligible, the dependence on the viewing angle is normally relevant, and can bias the results, as outlined, for example, in [20] and [5]. Therefore, by knowing both the directional emissivity and viewing angle, it is possible to correct the temperature accordingly. Examples can be found in [4,6] and moreover in [21], where the internal reflections due to concave surfaces of a complex test setup were also taken into account.

It is worth noticing that in these publications the different sensors were rigidly linked together (mainly for calibration purposes and real-time data integration), and the trend was to strengthen this physical union (until obtaining, in the final form, a unique device such as “HeatWave” [9]). However, in some cases, it can be more convenient to keep the two devices decoupled and independent. This is especially true in outdoor surveys, where there is often the need to perform the thermographic analysis at a specific day-time or night-time and/or weather conditions, which requires high versatility (e.g., for the assessment of the damages and energy efficiency of the building envelope [18]). Laser scanners, on the other hand, can be bulky and heavy; their handling and the regulation

of their position and orientation (usually are mounted on a tripod) may be time consuming and requires caution. Fixing a thermal camera to this type of scanner would make them even more difficult to regulate, and the easiness of handling of the thermal camera would be compromised. Furthermore, the two devices may have very different optics, which make the optimal distance of acquisition distinct. Conversely, with a decoupled acquisition, the integration can be applied in a flexible way, namely only when it is useful, based on previously recorded data, and does not affect the stand-alone capability of each device. In the literature, works based on this approach are uncommon. One exception is the work of A. G. Krefer et al. [22], which consisted of a method for generating 3D thermal models with decoupled acquisition, which relies on structure from motion and particle swarm optimization. Our paper focuses as well on a decoupled type of integration, but it differs from [22] in several aspects, such as the calibration method, the data fusion technique and the management of the superimposition of multiple thermograms.

This paper is organized as follows: the system architecture is outlined in Section 2, Section 3 and Section 4 explain the geometrical calibration and data fusion procedure adopted, respectively, and Section 5 presents the results. In the first experimental case of Section 5, the effectiveness of our approach is tested on a 3D-printed object properly designed. The second example of application belongs to the cultural heritage field. In the last few years, we have carried out a many-sided research project aimed at preserving and restoring the ancient sanctuary of Santa Maria delle Grazie, built toward the middle of the 15th century, in the place of Fornò near the city of Forlì (Italy). In particular, we have applied thermal imaging and laser scanning both to the building at large and to the ornamental elements. One example, presented in this paper, is the application to a marble statue called Madonna with the Child, an admirable work of the Florentine sculptor Agostino di Duccio (1418–1481), made up of four superimposed blocks. Up to the year 2000, this sculpture was in a niche on the entrance arch to the prothyrum of the sanctuary. Afterwards, since it showed clear signs of deterioration, especially due to rainwater and air pollution, it was carefully restored and then moved permanently to a great hall in the Bishop's palace in Forlì, where our surveys were performed.

2. System Architecture

In this section, the system architecture is presented, with the scope to define precisely the system components and the followed workflow.

2.1. System Components

The experimental set-up consisted of a thermal camera Testo 882 and a triangulation laser scanner Konica Minolta Vivid 9i, both shown in Figure 1. The Testo 882 has an FPA detector type with 320×240 pixel (but image output up to 640×480 thanks to the super-resolution feature), a FOV of $32^\circ \times 23^\circ$ and a range of detected temperature switchable between -20°C and $+100^\circ\text{C}$ and 0°C and $+350^\circ\text{C}$ (accuracy $\pm 2\%$ of reading for both). The Konica Minolta specifications are 307,200 pixels, three interchangeable lens of focal length 25 mm, 14 mm and 8 mm, and a weight of approximately 15 kg.



Figure 1. (a) Thermal camera Testo 882. (b) Laser scanner Konica Minolta Vivid 9i.

2.2. Integration Process Workflow

The workflow includes the processes for combining spatial data, acquired by the laser scanner in the form of a point cloud (spatial coordinates and normal vectors in each point) and the two-dimensional temperature map provided by the thermal camera, available as a temperature matrix. Before proceeding to the actual integration, it was necessary to have a post-processed point cloud, i.e., the registration of several range maps (by the well-known ICP algorithm) and all the post-processing operations have to be done previously.

There are three main processes characterizing the workflow (see Figure 2): acquisition, geometrical calibration (subdivided into intrinsic and extrinsic calibration) and data fusion. The emphasis of this work was both on the extrinsic calibration and on the data fusion process, which were carried out by Matlab programming.

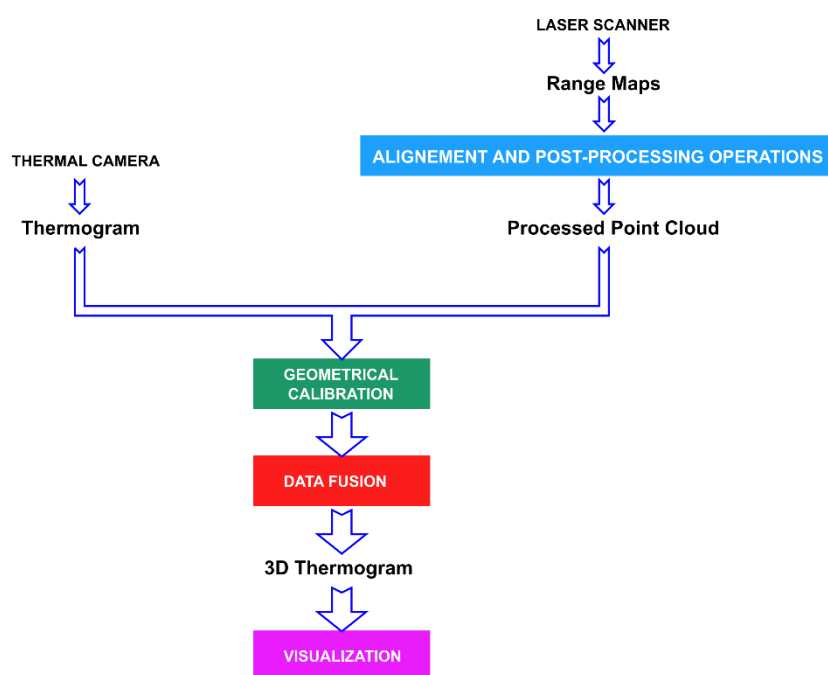


Figure 2. High-level workflow of the integration process.

3. Geometrical Calibration

When a thermal camera is involved, the term “calibration” could refer to two different types of calibration—the geometrical calibration and the radiometric calibration. The latter (which was not the object of this work) consists of a procedure that models the relationship between the digital output of the camera and the incident radiation. Hereinafter, if the term “calibration” is used, “geometrical calibration” is intended.

The geometrical calibration is a fundamental process used to compute intrinsic and extrinsic parameters of the system components.

Intrinsic parameters model the optic imaging process and are divided into two groups—four parameters defining the camera matrix, based on the pinhole camera model, and other parameters defining radial and tangential camera distortions caused by the physical properties of lenses [23].

There are six extrinsic parameters which describe the relative pose of the sensors, i.e., the transformation from world coordinates to camera coordinates. In the vast majority of methods, the intrinsic parameters are computed first, and then they are exploited to determine extrinsic ones.

For the determination of both intrinsic and extrinsic parameters, the best-known techniques involve the use of a target with a chessboard pattern, a method first presented by Z. Zang in [24]. Note that, in the case of thermal cameras, the pattern has to be detected in the IR spectrum. Several solutions have been presented for this purpose, involving, for example, a heated chessboard pattern

or marker detectable in IR; a more detailed state of art can be found in [10]. In the last few decades, however, some work has been done in order to develop the so-called automatic or targetless extrinsic calibration methods, that exploit natural scenes' features (see [11,25–27]).

Concerning the intrinsic calibration, we adopted the method developed by S. Vidas et al. in [28], which relies on a mask-based approach and does not require specialized equipment.

Regarding the extrinsic calibration, since our aim was to keep the two devices unlinked, the classical approach (consisting of acquiring several frames of the target in a fixed configuration of the two sensors) did not appear very suitable. For cases in which the identification of some homologous point is easily feasible, the extrinsic parameters are computed by means of a manual selection routine of homologous points, and then by exploiting the Matlab function “estimateWorldCameraPose”, which solves the perspective- n -point (PnP) [29]. However, this method might fail or achieve a very low accuracy in scenarios in which homologous points cannot be clearly identified. To address this limitation and to be able to consider more various scenarios, we developed an alternative method exploiting the detection of the object silhouette in the thermogram. In the literature, some silhouette-based calibration methods can be found, for instance in [30] and [31]. In [30], the pose of a known object is estimated through a hierarchical silhouette matching and unsupervised clustering, whereas in [31] the calibration of a stereo camera system was carried out by defining and minimizing a function that computes the distance between viewing cones, which are known from the silhouettes. The developed method has some similarities with the one presented in [11], in which the extrinsic parameters were obtained by minimizing an objective function that measures the alignment between edges extracted from RGB-D images (obtained by a depth camera) and the ones extracted from thermograms. Like that method, ours is suitable for thermograms in which the object contour can be easily identified, and so the object can be extracted from the background, a fairly common situation when dealing with thermograms. Differently from [11], however, the object contour detection is only needed for the thermogram, and the quantity evaluated is not the alignment between thermal and depth edges, but rather the “degree of filling” of the projected 3D points inside the thermal edge, which translates to a different definition of the objective function to minimize. Another difference is that this method takes as its input the full 3D representation of the object, and not the single view representation (often referred to as 2.5D [32]). This method can deal efficaciously with initial parameter values significantly different from the ones searched for, which makes it particularly suitable for a decoupled type of integration.

The extrinsic calibration method here developed can be divided into the following principal parts:

- 1) extraction of the object contour from the thermogram;
- 2) creation of a matrix \bar{M}_1 (with dimension equal to the resolution of the thermogram) obtained in this way: starting from the contour extracted in 1), a series of internal and external subsequent layers are created, and the same numerical value is associated to each pixel of each layer, following a specific function f_i defined later in this section;
- 3) projection of all the 3D points on the image plane with an initial set of extrinsic parameters, and the creation of a second matrix \bar{M}_2 composed of binary values, 1 if on that pixel at least one 3D point has been projected, 0 otherwise. Furthermore, in order to have a more uniform projection, the pixels sufficiently close to each pixel with values 1 are set to 1 if 0 (to avoid non uniform situations, especially for sparse point clouds);
- 4) minimization of the objective function of Equation (1), with a global minimization technique:

$$f_{obj} = f_{obj}(\bar{R}, \bar{t}) = -\text{sum}\{\bar{M}_1 \circ \bar{M}_2(\bar{R}, \bar{t})\} \quad (1)$$

where \circ is the Hadamard product, sum is the operator that sums all the elements in a matrix, \bar{R} is the rotation matrix (computed as a function of the three Euler angles) and \bar{t} is the translation vector.

The Euler angles and the components of the vector \bar{t} that make the function f_{obj} minimum are the searched values, namely the extrinsic parameters. It is to be outlined that the matrix \bar{M}_1 , once

computed, remains fixed in the optimization process described in Equation (1) and so only one contour extraction operation is required.

The point (1) can be done either in an automatic way (exploiting some well-known edge detection or background extraction algorithm) or manually. The background extraction also permits us to avoid the possible gross error of assigning to the 3D points, if the corresponding pixels are very near to the edges, the background temperature.

For the point (2), the function f_i has to be properly chosen, so to assure that the objective function minimum corresponds to the desired situation, namely that the 3D projected points completely fill up the zone inside the contour without overflowing outside the contour.

The function f_i utilized is of Gaussian type:

$$f_i(x) = c_1 e^{-c_2 x^2} \quad \text{if } x \geq 0 \quad (\text{layer } x \text{ inside the contour or the contour itself if } x = 0)$$

$$f_i(x) = c_1 e^{-c_2 x^2} - c_3 \quad \text{if } x < 0 \quad (\text{layer } x \text{ outside the contour})$$

($x \in \mathbb{Z}$, c_1, c_2, c_3 positive constants).

An example of the application of the function f_i is shown in Figure 3.

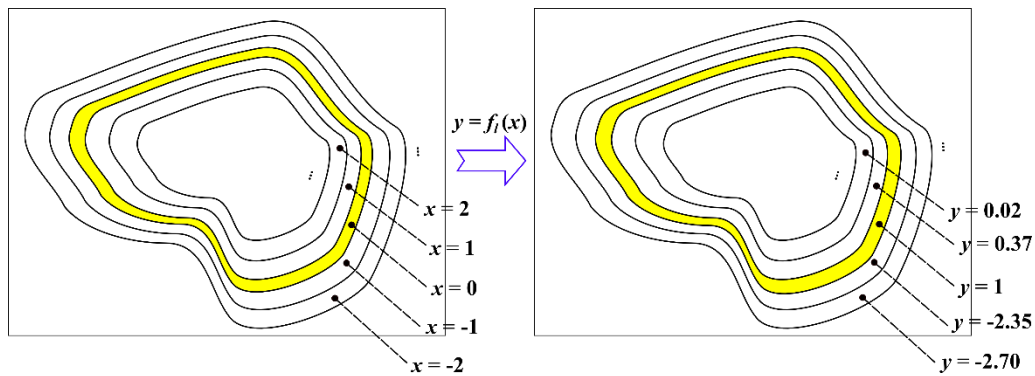


Figure 3. Example of the application of the function f_i on the layers map (with $c_1 = c_2 = 1$, $c_3 = e$). The object contour is outlined in yellow.

The c_3 constant is important, because it permits us to shift the curve of the external layers so that the function f_i has negative values there. This is necessary to avoid situations in which the minimum of the f_{obj} corresponds to the filling of zones incorporating the contour, but it is larger than the contour itself (in other words, to avoid that the projected points “overflow” the contour). As can be clearly seen in Figure 4, the function increases when approaching the object contour ($x = 0$) from both negative and positive values of x . This is another requirement, that the function f_i has to meet to grant the proper convergence of the function f_{obj} . Apart from these requirements, one has a certain freedom in the choice of the function f_i . In fact, the aforementioned Gaussian type function is the result of several experimental evaluations of possible functions. However, different choices of f_i can be further investigated in future works, in order to improve the speed and robustness of this calibration method. To minimize the objective function (1), with six independent variables (three Euler angles and three components of the translation vector \bar{t}), the Matlab optimization toolbox is exploited, in particular the function “GlobalSearch”.




Figure 4. Graph of the function f_i (considering 500 internal and external layers) with the actual constant values utilized ($c_1 = 0.01$, $c_2 = 0.001$, $c_3 = 0.011$).

4. Data Fusion

In this section, the data fusion process is examined, with particular emphasis on the ray-casting technique and the multiple thermograms handling. It is assumed that all the calibration parameters have already been computed, following the steps described in the previous sections.

4.1. Ray-Casting Technique

Since the complete point cloud of the object may be utilized, it is a common situation that only a subset of all the 3D points is seen by the thermal camera when each thermogram is acquired. The process of spotting the points seen by the thermal camera is normally fulfilled by a ray-casting technique, which allows us, among other things, to handle occlusions. When working with point clouds, one common approach is to first convert the original point cloud into a different data type, such as voxels or a mesh representation. Then, to perform the visibility check, one can start from each pixel of the thermal image and move with discrete steps on a specific line, stopping when the desired voxel (or mesh triangle) is reached ([14,33]).

The approach we propose, however, does not need this kind of type conversion and works with data in form of point cloud. Starting from all the 3D points, two types of exclusion are applied in order to obtain a subset, representing all and only the 3D points seen by the camera. In the next sections, the scheme of these two exclusions will be analyzed, but first it is to be said that there is actually another exclusion, which is trivial—that is, the exclusion of all the 3D points projected outside the image plane.

4.2. First Type of Exclusion

In the first type of exclusion, the image plane is firstly divided into a grid, whose cell dimension is defined by the parameter ξ (its measurement unit is pixel, and each cell is composed of $\xi \times \xi$ pixels). All the 3D points are projected onto the image plane, and then clustered depending on which cell of the grid they occupy. For each cluster, the points P_i , such that Equation (2) holds true, are excluded:

$$d_i > d_{min} + \chi \tag{2}$$

where d_i is the distance between the point P_i and the optical center O , d_{min} the minimum distance among the distances d_i and χ a fixed parameter (in millimeter, as for d_i and d_{min}). For higher values of ξ , higher values of χ are needed to take into account the fact that the surface region corresponding to the points not to be excluded is larger and so the difference in distances can be higher.

Theoretically, an optimal choice of the parameters ξ and χ would involve a non-trivial geometrical analysis, considering both the density in space of the 3D points (in general variable from

region to region) and the curvature of the object surface with respect to the viewing angle. In practice, they can simply be adjusted by graphically evaluating the results. In our case, in order to obtain proper experimental results, they were computed according to the experimental relations Equations (3) and (4), suggesting a possible choice of the two parameters:

$$\xi = \text{ceil}(Res/N) + 1 \quad [\text{pixel}] \quad (3)$$

$$\chi = 5/(1 - \xi/20) \quad [\text{mm}] \quad (4)$$

where N is the number of points projected on the image plane and Res the resolution of the image. In Equation (3), “ceil” is a function that rounds a number x to the nearest integer $\geq x$. Since χ is positive, according to Equation (4) ξ must be less than 20, a condition always verified for common resolutions and points cloud densities. It has to be pointed out that, for the vast majority of object geometries, the results are stable with respect to a variation of ξ and χ of some units, and so in any case their resulting values are acceptable as long as they fall into a proper range.

4.3. Second Type of Exclusion

The second type of exclusion exploits the knowledge of the normal vectors \hat{n}_i in each 3D point P_i and requires that all the points for which Equation (5) holds true are excluded:

$$(O - P_i) \cdot \hat{n}_i < 0 \quad (5)$$

where O is the optical center.

This condition derives from the fact that, if the dot product in Equation (5) is negative, the angle between the vector $(O - P_i)$ and the normal vector \hat{n}_i is greater than 90 degrees, and so the point is not seen by the camera. This condition on its own is theoretically suitable for excluding all the inappropriate points in convex surfaces, but it fails when dealing with surfaces that present concavities. Conversely, even if the first exclusion appears to be able to handle a generic shape of the point cloud, it fails in some particular scenarios in which the point cloud lacks some parts, because in these cases some temperatures might be wrongly assigned to the opposite part of the point cloud. For these reasons, the two types of exclusion were combined, in accordance with the experimental results as well. Figure 5 shows two examples of applications of the two exclusions in succession with two different values of ξ , on a region of a hypothetical surface, which is represented in the points cloud by 11 points. The zones between each pair of dotted lines (converging to the optical center) represent the band corresponding to each pixel. Inside each band, the part at the right of the purple segments are excluded according to the first type of exclusion. The points in green are the points to be maintained, the points in blue are the points correctly excluded with the first type of exclusion, the points in yellow the ones correctly excluded only with the second type of exclusion and the points in red the ones not correctly excluded after applying both the exclusions, and to which consequently a wrong temperature is assigned. In Figure 5b, better results are obtained if the value of ξ is doubled, because in this way all the points belonging to a part of the object surface not seen by the camera are correctly excluded. Indeed, the value of the parameter ξ has to be chosen so that in each cell the number of reprojected points is higher than one, so that they form an actual cluster to which we can properly apply Equation (2). Finally, it is worth noticing that the order of the two exclusions is important: in fact, by inverting the exclusions order, it is easy to see that the point P_{11} is never excluded.

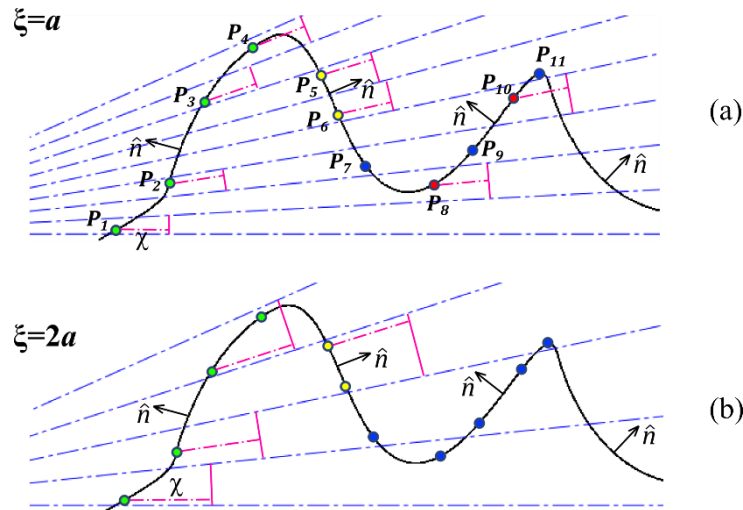


Figure 5. Application of the exclusion methods on a part of a hypothetical surface represented by eleven points, for two values of ξ . The value of ξ in (b) is twice the value of ξ in (a). Green points: to be maintained; blue points: excluded with the first exclusion; yellow points: excluded with the second exclusion; red points: not correctly excluded with the two exclusions.

4.4. Temperature Assignment

Once the subset of the 3D points seen by the thermal camera is defined, the corresponding temperature is assigned at each point, by considering in the temperature matrix the value in correspondence of the pixel in which the 3D point has been projected.

In Figure 6 the overall workflow of the data fusion process is outlined.

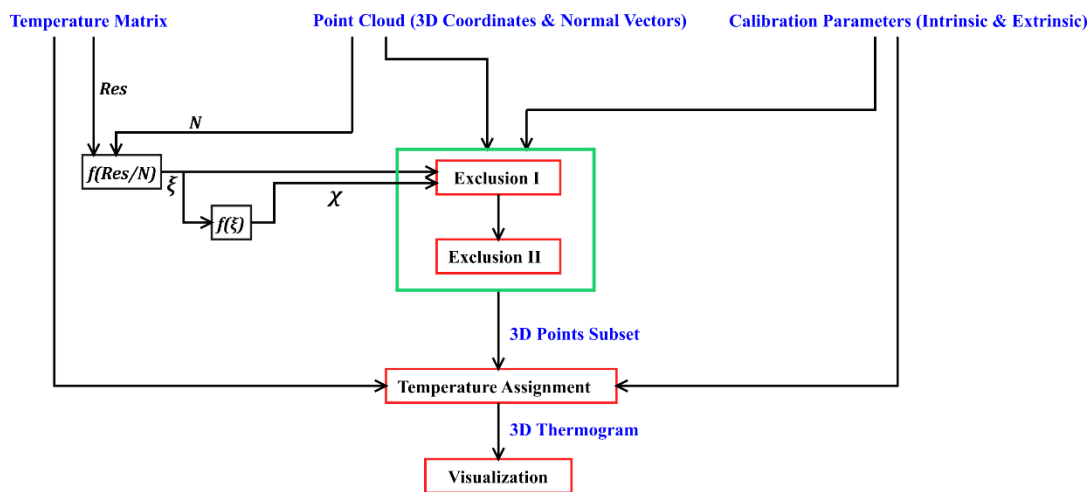


Figure 6. Overall workflow of the data fusion process.

4.5. Multiple Thermograms

In the previous sections, the case of the integration of a single thermogram on a point cloud was considered. There might be, however, the need to integrate more than one thermogram on the same point cloud, in order to have a larger set of points with an associated temperature.

Note that this procedure requires in the first place thermograms acquired in a temporal window in which no measurable changes in the thermal state occur, so as to guarantee that the integration is not carried out on discordant data. Since each thermogram can be individually integrated on the point cloud by the methodology described previously, the problem comes down to handling the overlapping zones, namely the points to which more than one temperature value is assigned. To give these points a single final temperature, the method utilized in [9] was followed. The method relies on the fact that the relative orientation between the surface of the inspected object and the camera affects

the measurement accuracy, and, more specifically, the lower the viewing angle is, the higher the accuracy is. As shown in Equation (6), the temperature T_i assigned to the point P_i is computed as a weighted average of the temperatures T_{ij} :

$$T_i = \frac{\sum_{j=1}^{N_i} c_{ij} T_{ij}}{\sum_{j=1}^{N_i} c_{ij}} \quad (6)$$

where the weight is the confidence factor c_{ij} , the index i refers to the point of the points cloud and ranges from 1 to N , whereas the index j refers to the thermograms that overlap in the point i and ranges from 1 to N_i . The confidence factor c_{ij} is computed as a function of the viewing angle θ_{ij} as shown in Equation (7):

$$c_{ij} = e^{-\kappa\theta_{ij}} \quad (7)$$

where the viewing angle θ_{ij} is the angle from which the thermal camera sees the point P_i , considering the thermogram j , and can be computed as shown in Equation (8):

$$\theta_{ij} = \arccos \frac{(O_j - P_i) \cdot \hat{n}_i}{\|O_j - P_i\|} \quad (8)$$

with point O_j identifying the optical center for the thermogram j .

In this way, a greater weight is assigned to the rays with smaller viewing angles, which allows more accurate measures. More precisely, the weight decreases with an exponential law, depending on a parameter κ , that was set equal to 2 according to experimental evaluations.

4.6. Visualization

The mathematical result of the integration is a N -by-4 matrix, where each line contains the 3D coordinates of a 3D point plus the associated temperature (or, if no temperature is associated, a NaN value). This matrix is visualized in Matlab by assigning to the point subset with an associated temperature a particular colormap, whereas a different color is assigned to the points with no temperature associated.

Figure 7 shows an outline of the workflow in the case of the integration of multiple thermograms on the same point cloud.

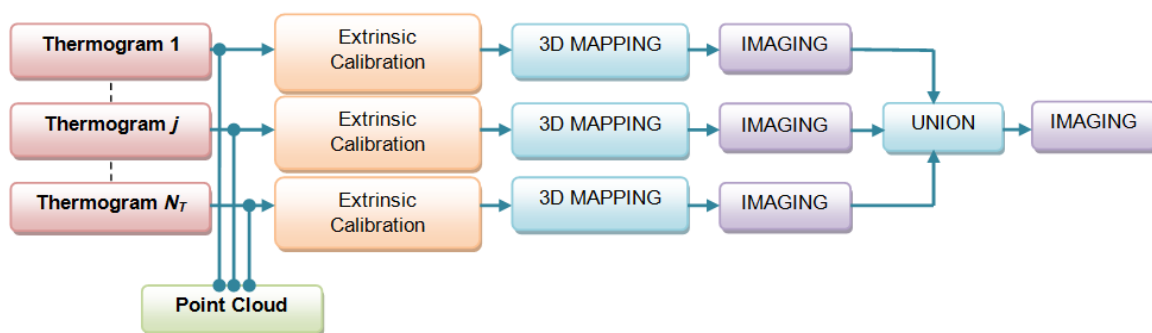


Figure 7. Workflow of the integration of multiple thermograms on the same points cloud. N_T is the total number of thermograms integrated.

5. Results

5.1. Test Object

The experiments were carried out on a particular object, designed so that the two aforementioned extrinsic calibration methods were both equally suitable and easy comparable. In addition, the particular design of the object shown in Figure 8 allowed us to evaluate the effectiveness

of the ray-casting technique on a geometry with concavities. It has an internal conic hole and three different radial parts (shifted by angles of 120 degrees), characterized by many edges.




Figure 8. (a) 3D printed object; (b) Its CAD model.

The object was designed with CAD, 3D-printed with acrylonitrile butadiene styrene (ABS) material and then scanned with the Konica Minolta Vivid 9i laser scanner. The point cloud could be obtained directly from the CAD model, but the object was scanned to take into account the acquisition bias and test the effectiveness of the method on laser scanner data.

Here we present a first set of simulations, carried out to assess the effectiveness of the automatic extrinsic calibration method. This was made by comparing the extrinsic parameter values obtained by means of this automatic procedure with the ones resulting from the method based on the manual selection of homologous points. Because of the specially designed geometry features of the printed object, it was possible to choose for the first method a set of about 20–30 homologous points for each analyzed thermogram, with a mean reprojection error (MRE) of about 1–2 pixels. In this way, the proper reliability of the first method parameters was assured, and so it was possible to evaluate the ones obtained with the second method. Equation (9) shows the expression for computing the MRE. The units of MRE are pixels.

$$MRE = \frac{\sum_{i=1}^n \|P_i - Q_i\|}{n} \quad (9)$$

In Equation (9), n is the number of homologous points selected, P_i the 2D coordinates of the i -th point selected on the thermogram, Q_i the 2D coordinates of the reprojection of the i -th homologous 3D point.

For each thermogram, the convergence from different initial positions was analyzed, within a range of ± 25 degrees for the Euler angles and ± 20 mm for the translation vector from the parameters obtained with the first method.

In Table 1, some of the results of the first calibration method are shown for four different thermograms, whereas Table 2 gives the differences (considering the same thermograms) between the extrinsic parameters obtained with the second method and with the first one respectively.

Figure 9 shows a set of 2D points and the reprojection of 3D homologous points, once the extrinsic parameters were computed.

Table 1. Extrinsic parameters for four different thermograms, in the case of the method based on the manual selection of the homologous points (first method).

Thermogram	N° of Points	MRE	t_1 [mm]	t_2 [mm]	t_3 [mm]	α [°]	β [°]	γ [°]
A	21	1.51	−18.64	19.18	544.17	126.73	3.25	−114.01
B	26	1.51	−16.34	23.23	497.88	−73.48	−1.33	−120.38
C	21	1.61	−3.19	21.62	375.58	25.16	−0.14	−86.90
D	30	1.52	−1.11	12.48	453.16	−3.35	−1.68	−125.87

Table 2. Differences between the extrinsic parameters obtained with the automatic method (second method) and the ones obtained with the first method.

Thermogram	Δt_1 [mm]	Δt_2 [mm]	Δt_3 [mm]	$\Delta\alpha$ [°]	$\Delta\beta$ [°]	$\Delta\gamma$ [°]
A	1.13	0.01	-11.53	0.76	0.83	-0.74
B	0.78	0.96	-4.75	0.29	-0.08	1.67
C	-0.11	1.15	-7.36	0.45	-0.90	0.35
D	-0.04	0.90	-9.66	-0.07	0.41	0.43




Figure 9. Grayscale thermogram with highlighted the set of 2D points (purple circles) and the reprojection of 3D homologous points (yellow cross) once the extrinsic parameters were computed, in the case of thermogram B of Table 1.

With reference to Table 2, thermograms A and B were obtained starting from a first guess of the parameters with the following differences with respect to the ones of Table 1: $\Delta t_1 = 20$ mm, $\Delta t_2 = -20$ mm, $\Delta t_3 = 20$ mm, $\Delta\alpha = -20^\circ$, $\Delta\beta = 0^\circ$, $\Delta\gamma = 20^\circ$. For thermograms C and D, the initial differences were set as follows: $\Delta t_1 = -10$ mm, $\Delta t_2 = 20$ mm, $\Delta t_3 = 10$ mm, $\Delta\alpha = -20^\circ$, $\Delta\beta = 15^\circ$, $\Delta\gamma = 15^\circ$.

As can be seen in Table 2, the final differences are relatively low, except for the Δt_3 , which is of several millimeters. However, this can be acceptable, considering that the t_3 parameter is greater by an order of magnitude compared to the parameters t_1 and t_2 (relative error of about 2%).

In Figure 10, the initial (a) and final (b) projection in the case of automatic extrinsic calibration of the thermogram B are shown. This type of visualization allows for a qualitative consideration on the effectiveness of the method. As can be easily seen, in the case (b), the 3D projected points (in red) fill the thermogram contour (yellow line) well.

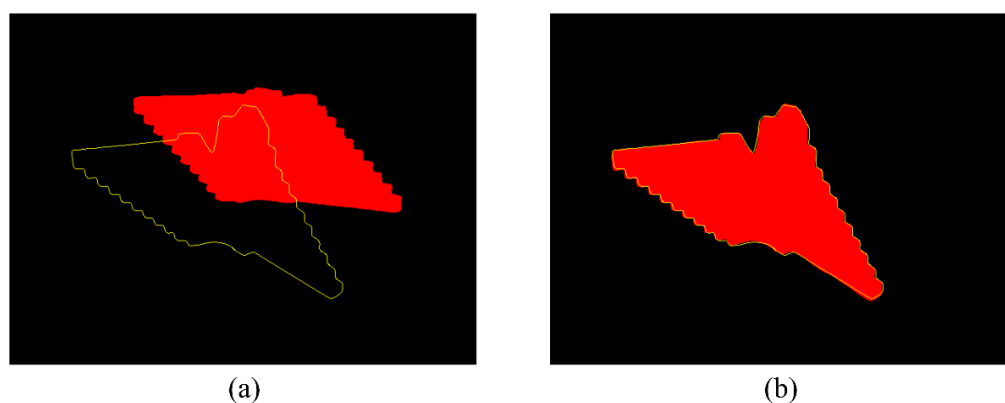


Figure 10. Case of automatic calibration of thermogram B: (a) Initial projection of 3D points, using a set of initial extrinsic parameters; (b) Projection utilizing the parameters for which the objective function of Equation (1) presents a global minimum.

Figure 11 shows a representation of the matrix $\bar{\bar{M}}_1$ (see Section 3, Figures 3–4) converted for visual purposes into a color image.




Figure 11. Visual representation of the matrix $\bar{\bar{M}}_1$ created from the yellow contour of Figure 10. In the scale, y refers to the function graphed in Figure 4, which is here multiplied by a factor of 10^4 for the sake of clarity.

The second set of experiments was aimed at evaluating the general reliability of the technique, by qualitatively assessing the final results obtained with the integration of different thermograms.

In Figure 12, the result in the case of the integration of thermogram B is shown, whereas Figure 13 shows an example of the integration of multiple thermograms (17 in total) acquired from different poses. The shots were taken while the object was being heated with a drier, fixed so that the heat flux came in contact with one side in particular (the one with higher temperature, colored yellow).

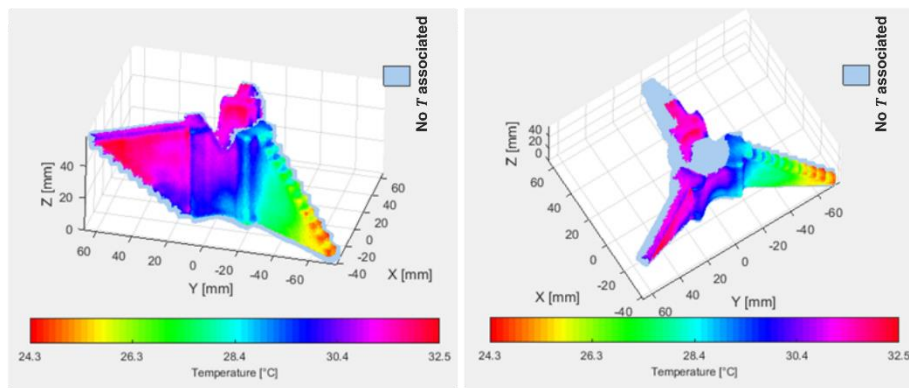


Figure 12. Two different views of the result of the integration of thermogram B on the point cloud.

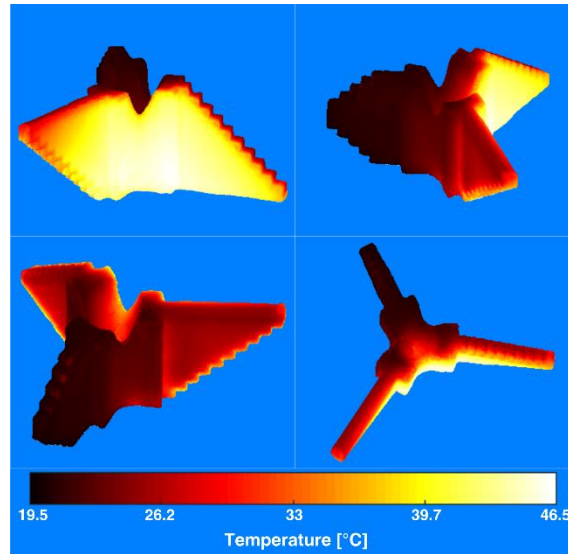


Figure 13. Four different views of the result of the integration of seventeen thermograms.

Figure 14 shows how the union of several thermograms (acquired from different poses) by the method explained in Section 4.5 can compensate for the systematic error in the temperature caused by the dependence of the emissivity on the viewing angle. In Figure 14c, the temperatures of a single thermogram and of the union of several thermograms (the temperature superimpositions vary from 4 to 8 depending on the point) are compared along a key zone, shown in Figure 14a,b, in which the normal vectors of the surface have a significant variation, which implies a significant variation on the viewing angles (referring to the case of the single thermogram, since after the union the concept of viewing angle loses meaning). In this zone, the temperature of the single thermogram is appreciably lower than the temperature of the union (with a maximum difference of roughly 1 °C). This behavior can be explained by correlating the difference ΔT of the temperatures ($\Delta T = T_{UNION} - T_{SINGLE}$) with the variation of the viewing angle, as shown in Figure 14d. The camera view during the acquisition of the single thermogram can be approximately assumed to be the view of Figure 14a, which explains the angle of about 40 degrees for Y in the range 5–10 and 16–25 mm (the camera was tilted downwards with respect to the normal vectors in these points by about 40 degrees). For non-conductor materials, emissivity is nearly constant from 0 to 40–45 degrees, whereas at larger angles it has a significant drop [1] (pp. 39–40). This justifies the fact that the temperature in the zone of high viewing angles is underestimated in the case of the single thermogram. This error can be effectively compensated by the adopted method, as long as, for the same points, temperatures with lower viewing angles associated are available.

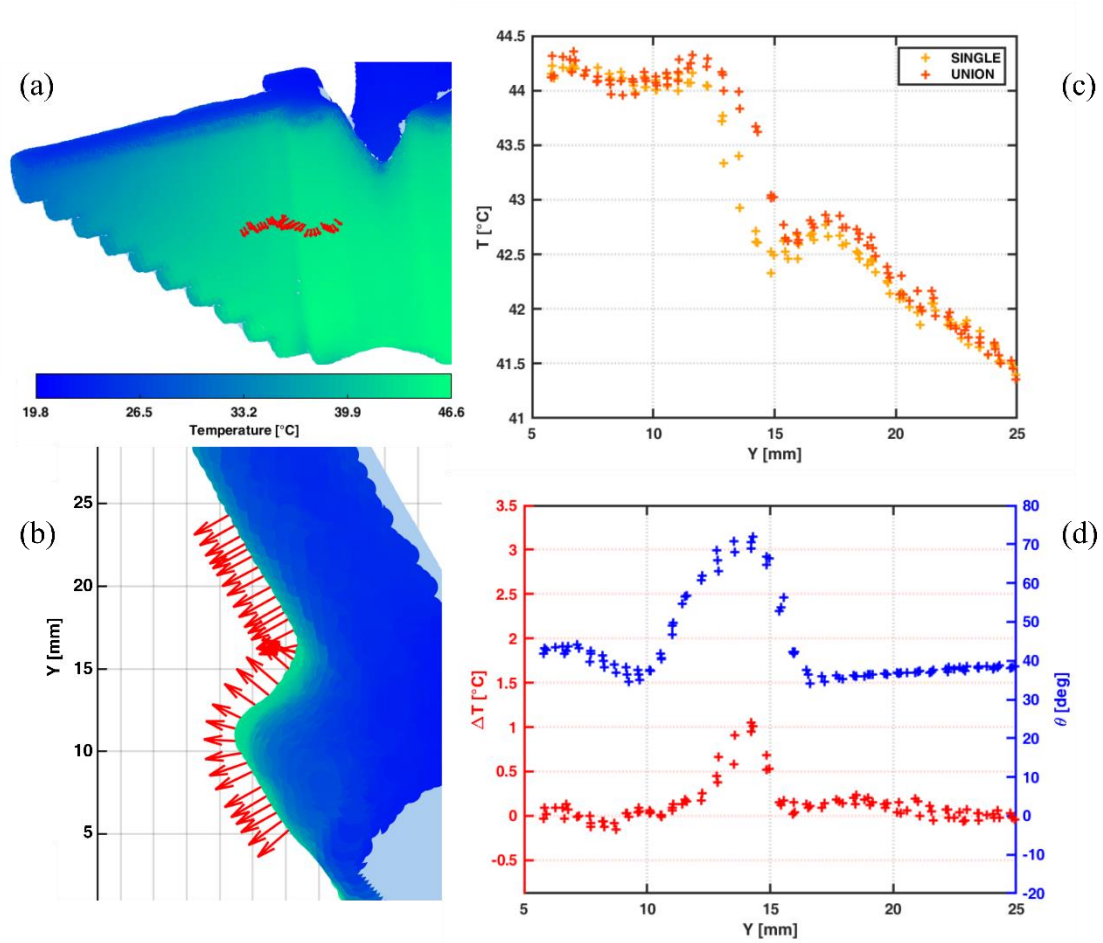


Figure 14. Effect of the temperature correction achieved by the union of different thermograms. (a) Zone in which the temperatures are examined, with the vectors normal to the surface in these points colored red; temperatures are shown with a “winter” colormap; (b) Close-up of the normal vectors, from a top view; (c) Comparison between the temperature of a single thermogram and the one resulting from the union of several thermograms; (d) Temperature difference (T resulting from the union minus T of the single thermogram) and viewing angle for each point examined.

5.2. Statue “Madonna with the Child”

After testing the method on the object previously described, in this section, an application in the cultural heritage field is considered. The integration was carried out on the statue Madonna with the Child, by the Florentine sculptor Agostino di Duccio (1418–1481), with an eye to monitoring its condition and gather additional information about its state. The point cloud of the statue was already available from previous works. Figure 15 shows the statue and the laser scanner in the Bishop’s palace in Forlì. For the sake of brevity here, we present only some of the most significant results.



Figure 15. Statue Madonna with the Child, with the laser scanner in position.

In Figure 16, three thermograms of the statue acquired from different poses are shown.




Figure 16. Three thermograms of the statue, acquired from different poses.

Figure 17 shows several outcomes of the integration. There is a little spot on the top of the head to which no temperature is assigned, and to which a specific colour not belonging to the colormap is assigned. The survey shows that there is one side of the statue that is slightly hotter, whereas the base and the head present lower temperatures. It is probable that the particular lighting system and the statue arrangement near a window with a non-insulated frame give rise to these changes, which is not favourable for optimal preservation. Further investigations are needed to better clarify the cause.

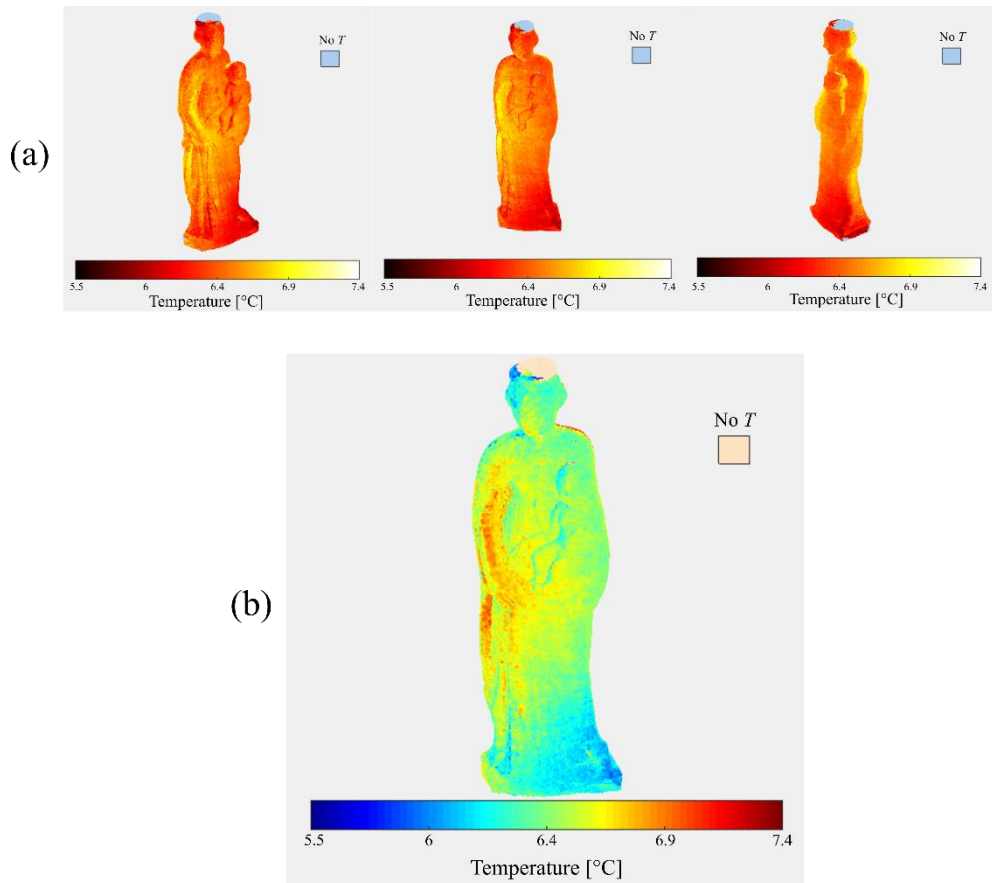


Figure 17. Results of the integration of twelve thermograms. (a) Three different views coloured with the colormap “hot”. (b) Frontal view coloured with the colormap “jet”.

6. Discussion

The work presented includes methodologies to address each integration step, with the final aim of achieving 3D thermography by maintaining the decoupling of the devices. Concerning the extrinsic calibration, we proposed an automatic method that exploits the object silhouette. An evaluation of the accuracy of the method is presented in Table 2. Since our automatic method does not exploit the concept of homologous points (and thus, MRE cannot be computed), the evaluation of its accuracy was made by a comparison with the extrinsic parameters (considered the ground truth) obtained by the manual selection of homologous points, executed on a test object purposely designed. A comparison can be made with the automatic calibration method based on a silhouette developed by J.T. Lussier and S. Thurn [11]. The error concerning the x and y translations (Δt_1 and Δt_2 in Table 2) in our work is about one order of magnitude lower (the max obtained is around 1 mm); concerning the z translation (Δt_3 in Table 2), our maximum error stays within 12 mm, against the over 50 mm reported in [11]. Regarding the error on the rotations ($\Delta\alpha$, $\Delta\beta$ and $\Delta\gamma$ in Table 2), our errors are higher, with a mean of 0.58 degrees against the 0.3 degrees reported in [11]. We want to point out that this comparison is made between integration procedures that suit different integration modalities. In [11], indeed, the authors carried out a real-time integration between thermograms and depth-maps (2.5D), whereas we integrate (offline) thermograms with 3D point clouds. The difference in the type of the range data integrated entails, among other things, the following fact—the accuracy of the method used in [11], as the authors said, depends on the scene coverage, that is to say the percentage of the area covered by the object of interest with respect to the background. The higher the coverage is, the higher the edge variability is, and the lower the error in the extrinsic parameters. In our case, since we did not work with an edge map but with 3D point cloud, the concept of edge variability is not similarly defined and the error did not depend on this parameter.

Unfortunately, regarding the accuracy, a direct comparison with the decoupled method presented by A. G. Krefer *et al.* [22] is not possible because the automatic calibration procedure they used was based on the automatic matching of interest points and was evaluated by the classical MRE. Apart from the calibration method, other differences from the work [22] include the fact that they used the 3D data in the form of a polygonal mesh, whereas we keep the 3D data as a point cloud during the whole integration process. For visual purposes, if need be, the results may be converted into a coloured mesh at a later time. Concerning the handling of the points in which more temperatures are superimposed, we computed the temperature to be assigned as a weighted average, where the weight decreases in an exponential way if the viewing angle increases, which was the method exploited by S. Vidas and P. Moghadam in [9]. In [22], conversely, the weight of each point temperature was computed as a function of the position of the point inside the view frustum of the camera (the weight increases if nearer to the optical axis or to the optical centre). A flaw of this latter approach is that it did not take into account the object geometry (i.e., the normal vectors in each point) and so it was prone to fail to compensate for the variation of the emissivity at high viewing angles. Figure 14 clearly shows that at high viewing angles the temperature can be underestimated (in that specific case up to one degree). The methodology followed, which takes the object geometry into account, allows us to overcome this issue, assuming that thermograms of the interested area can be acquired from different orientations. However, if this is not possible (for instance, because of the position and of the limited mobility of the object or if the acquisition time is limited), an improvement of this method could be to apply a temperature correction to each thermogram singularly, exploiting, for instance, the correction formula proposed in [4], which relies on a theoretical model for the directional emissivity.

The whole integration methodology was first tested on a purposely designed 3D-printed object and then on a historical marble statue, and the results demonstrate the general feasibility of the approach. We are planning, however, further tests, in particular aimed at improving the robustness of the automatic extrinsic calibration method, which is affected, to a certain degree, by the object geometry (especially in terms of level of detail of the geometrical features and of the presence of symmetries). For objects which present a sufficient number of points of interest clearly identifiable both in infrared and in their 3D geometry (e.g., very sharp edges), the manual selection of the homologous points can still be a better option to compute the extrinsic parameters, and could be improved, for instance by applying to the thermogram the intensity transformation proposed in [22], which is able to highlight certain points which are otherwise not clear enough.

7. Conclusions

In this work, a decoupled acquisition method for generating 3D thermal models is proposed. The integration is carried out using the triangulation laser scanner Konica Minolta Vivid 9i and the thermal camera Testo 882, but it can be exploited for generating 3D thermal models with a generic range sensor and a generic thermal camera. The two devices are kept independent during the acquisition phase, allowing the integration of 3D data and thermal data acquired at different times.

With regard to the extrinsic calibration, two methods are used, a more standard one relying on a manual selection of homologous points, and an “automatic” one, the latter based on finding the optimum of a particular function which evaluates the degree of filling of the reprojection of the 3D points inside the object silhouette in the thermogram. The former method is used to assess the effectiveness of the latter, which is proven to work well in the case study, but has room for improvements, especially in terms of robustness. Concerning the data fusion, we propose an easy to implement algorithm which is able to deal with complex object shapes, handle occlusions and cases of incomplete data from the range finder. Furthermore, the viewing angle is computed, and it is used to calculate a weight for each ray, in order to assign a proper temperature value in the zones in which, when integrating multiple thermograms, overlaps occur. It was shown how this can effectively reduce the error in the temperature due to the dependence of the emissivity on the viewing angle.

The integration methodology was first tested on a 3D-printed object and was then applied to a cultural heritage case study, and the results suggest that this approach can be effective and useful

with an eye to integration and restoration. We are planning further tests to better investigate the effectiveness of the method.

Author Contributions: The authors have given an equal contribution to all the different parts of this article. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Acknowledgments: The authors would like to thank Ing. Massimiliano Fantini (Romagna Tech S.C.P.A.) for his important contribution to 3D laser scanner survey about the statue of Agostino di Duccio.

Conflicts of Interest: The authors declare no conflict of interest.

References

- Vollmer, M.; Möllmann, K.-P. *Infrared Thermal Imaging*; WILEY-VCH Verlag GmbH & Co. KGaA: Weinheim, Germany, 2010.
- Soldofieri, F.; Dumoulin, J. Infrared thermography: From sensing principle to nondestructive testing considerations. In *Sensing the Past. Geotechnologies and the Environment*; Masini, N., Soldovieri, F., Eds.; Springer, Cham, Switzerland: 2017; Chapter 11, pp. 233–255.
- Historic England 2018. *3D Laser Scanning for Heritage: Advice and Guidance on the Use of Laser Scanning in Archaeology and Architecture*; Historic England: Swindon, UK, 2018. Available online: <https://historicengland.org.uk/images-books/publications/3d-laser-scanning-heritage/heag155-3d-laser-scanning/> (accessed on 23 January 2020).
- Chernov, G.; Chernov, V.; Barboza Flores, M. 3D Dynamic Thermography System for Biomedical Applications. In *Application of Infrared to Biomedical Sciences*; Ng, E.Y., Etehadtavakol, M., Eds.; Springer: Singapore, 2017; pp. 517–545, doi:10.1007/978-981-10-3147-2_28.
- Ordóñez Muller, A.; Kroll, A. Generating High Fidelity 3-D Thermograms with a Handheld Real-Time Thermal Imaging System. *IEEE Sens. J.* **2017**, *17*, 774–783, doi:10.1109/JSEN.2016.2621166.
- Cardone, G.; Ianiro, A.; dello Ioio, G.; Passaro, A. Temperature Maps Measurements on 3D Surfaces with Infrared Thermography. *Exp. Fluids* **2012**, *52*, 375–385, doi:10.1007/s00348-011-1225-9.
- Ju, X.; Nebel, J.-C.; Siebert, J.P. 3D Thermography Imaging Standardization Technique for Inflammation Diagnosis. In Proceedings of the Photonics Asia 2004, Beijing, China, 8–12 November 2004; pp. 266–273, doi:10.1117/12.577055.
- Skala, K.; Lipić, T.; Sović, I.; Grubišić, I. Dynamic Thermal Models for Human Body Dissipation. *Period. Biol.* **2015**, *117*, 167–176.
- Vidas, S.; Moghadam, P. HeatWave: A Handheld 3D Thermography System for Energy Auditing. *Energy Build.* **2013**, *66*, 445–460, doi:10.1016/j.enbuild.2013.07.030.
- Rangel, J.; Soldan, S. 3D Thermal Imaging: Fusion of Thermography and Depth Cameras. In Proceedings of the 2014 International Conference on Quantitative InfraRed Thermography, Bordeaux, France, 7–11 July 2014; doi:10.21611/qirt.2014.035.
- Lussier, J.T.; Thrun, S. Automatic Calibration of RGBD and Thermal Cameras. In Proceedings of the 2014 IEEE/RSJ International Conference on Intelligent Robots and Systems, Chicago, IL, USA, 14–18 September 2014; pp. 451–458, doi:10.1109/IROS.2014.6942598.
- Cao, Y.; Xu, B.; Ye, Z.; Yang, J.; Cao, Y.; Tisse, C.-L.; Li, X. Depth and Thermal Sensor Fusion to Enhance 3D Thermographic Reconstruction. *Opt. Express* **2018**, *26*, 8179, doi:10.1364/OE.26.008179.
- Yang, M.-D.; Su, T.-C.; Lin, H.-Y. Fusion of Infrared Thermal Image and Visible Image for 3D Thermal Model Reconstruction Using Smartphone Sensors. *Sensors* **2018**, *18*, 2003, doi:10.3390/s18072003.
- Borrmann, D.; Nüchter, A.; Đakulović, M.; Maurović, I.; Petrović, I.; Osmanković, D.; Velagić, J. A Mobile Robot Based System for Fully Automated Thermal 3D Mapping. *Adv. Eng. Inf.* **2014**, *28*, 425–440, doi:10.1016/j.aei.2014.06.002.
- Seccia, L.; Iannucci, A. M.; Santopuoli, M.; Vernia, B. Early Pixels. *Image Process.* **1996**, *8*, 10–12.
- Rizzi, A.; Voltolini, F.; Girardi, S.; Gonzo, L.; Remondino, F. Digital Preservation, Documentation and Analysis of Paintings, Monuments and Large Cultural Heritage with Infrared Technology, Digital Cameras and Range Sensors. In Proceedings of the XXI International CIPA Symposium, Athens, Greece, 1–6 October 2007.

17. Gianinetto, M.; Giussani, A.; Roncoroni, F.; Scaioni, M.; Di Milano, P.; Rilevamento, D.; Di Lecco, P.R. Integration of Multi-Source Close-Range Data. In Proceedings of the CIPA XX International Symposium, Torino, Italy, 26 September–1 October 2005.
18. Scaioni, M.; Rosina, E.; L'Erario, A.; Díaz-Vilariño, L. Integration of Infrared Thermography and Photogrammetric Surveying of Built Landscape. *ISPRS Int. Arch. Photogramm. Remote. Sens. Spat. Inf. Sci.* **2017**, *153*–160, doi:10.5194/isprs-archives-XLII-5-W1-153-2017.
19. Stanco, F.; Battiato, S.; Gallo, G. *Digital Imaging for Cultural Heritage Preservation*; CRC Press: Boca Raton, Florida, USA, 2011.
20. Kawasue, K.; Win, K.D.; Yoshida, K.; Tokunaga, T. Black Cattle Body Shape and Temperature Measurement Using Thermography and KINECT Sensor. *Artif. Life Robot.* **2017**, *22*, 464–470, doi:10.1007/s10015-017-0373-2.
21. Robinson, D.W.; Simpson, R.; Parian, J.A.; Cozzani, A.; Casarosa, G.; Sablerolle, S.; Ertel, H. 3D Thermography for Improving Temperature Measurements in Thermal Vacuum Testing. *CEAS Space J.* **2017**, *9*, 333–350, doi:10.1007/s12567-017-0167-3.
22. Krefer, A.G.; Lie, M.M.I.; Borba, G.B.; Gamba, H.R.; Lavarda, M.D.; De Souza, M.A. A method for generating 3D thermal models with decoupled acquisition. *Comput. Methods Programs Biomed.* **2017**, *151*, 79–90, doi:10.1016/j.cmpb.2017.08.009.
23. Jähne, B. *Practical Handbook on Image Processing for Scientific Applications*; CRC Press: Boca Raton, Florida, USA, 1997.
24. Zhang, Z. A flexible new technique for camera calibration. *IEEE Trans. Pattern Anal. Mach. Intell.* **2000**, *22*, 1330–1334, doi:10.1109/34.888718.
25. Moghadam, P.; Bosse, M.; Zlot, R. Line-Based Extrinsic Calibration of Range and Image Sensors. In Proceedings of the 2013 IEEE International Conference on Robotics and Automation, Karlsruhe, Germany, 6–10 May 2013; pp. 3685–3691, doi:10.1109/ICRA.2013.6631095.
26. Pandey, G.; McBride, J.R.; Savarese, S.; Eustice, R.M. Automatic Targetless Extrinsic Calibration of a 3D Lidar and Camera by Maximizing Mutual Information. In Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence, Toronto, Canada, 22–26 July 2012; pp. 2053–2059.
27. Scaramuzza, D.; Harati, A.; Siegwart, R. Extrinsic Self Calibration of a Camera and a 3D Laser Range Finder from Natural Scenes. In Proceedings of the 2007 IEEE/RSJ International Conference on Intelligent Robots and Systems, San Diego, CA, USA, 29 October–2 November 2007; pp. 4164–4169, doi:10.1109/IROS.2007.4399276.
28. Vidas, S.; Lakemond, R.; Denman, S.; Fookes, C.; Sridharan, S.; Wark, T. A Mask-Based Approach for the Geometric Calibration of Thermal-Infrared Cameras. *IEEE Trans. Instrum. Meas.* **2012**, *61*, 1625–1635, doi:10.1109/TIM.2012.2182851.
29. Gao, X.-S.; Hou, X.-R.; Tang, J.; Cheng, H.-F. Complete Solution Classification for the Perspective-Three-Point Problem. *IEEE Trans. Pattern Anal. Mach. Intell.* **2003**, *25*, 930–943, doi:10.1109/TPAMI.2003.1217599.
30. Reinbacher, C.; Ruther, M.; Bischof, H. Pose Estimation of Known Objects by Efficient Silhouette Matching. In Proceedings of the 2010 20th International Conference on Pattern Recognition, Istanbul, Turkey, 23–26 August 2010; pp. 1080–1083, doi:10.1109/ICPR.2010.270.
31. Boyer, E. On Using Silhouettes for Camera Calibration. In *Computer Vision – ACCV 2006*; Narayanan, P.J., Nayar, S.K., Shum, H.-Y., Eds.; Springer: Berlin/Heidelberg, Germany, 2006; Volume 3851, pp. 1–10, doi:10.1007/11612032_1.
32. Tateno, K.; Tombari, F.; Navab, N. When 2.5D Is Not Enough: Simultaneous Reconstruction, Segmentation and Recognition on Dense SLAM. In Proceedings of the 2016 IEEE International Conference on Robotics and Automation (ICRA), Stockholm, Sweden, 16–21 May 2016; pp. 2295–2302, doi:10.1109/ICRA.2016.7487378.
33. Vidas, S.; Moghadam, P.; Sridharan, S. Real-Time Mobile 3D Temperature Mapping. *IEEE Sens. J.* **2015**, *15*, 1145–1152, doi:10.1109/JSEN.2014.2360709.

