A STATISTICAL ANALYSIS FOR THE ASSESSMENT OF CLOSE-RANGE PHOTOGRAMMETRY GEOMETRICAL FEATURES

A. di Filippo[1], S. Antinozzi[1], A. Dell'Amico[2], A. Sanseverino[2]

[1] Dept. of Civil Engineering (DICIV), Università degli Studi di Salerno, Italy – (anddifilippo1, santinozzi)@unisa.it
[2] Dept. of Civil Engineering and Architecture (DICAR), Università degli Studi di Pavia, Italy – (anna.dellamico, anna.sanseverino)@unipv.it

Technical Commission II

KEY WORDS: Photogrammetric processing, Accuracy evaluation, Agisoft Metashape, Tie Points Filtering, Covariance Matrix.

ABSTRACT:

An examination of the traceability and dependability of the virtualisation properties is prompted by the widespread use of three-dimensional models. The challenge of obtaining accuracy indicators directly from the photogrammetric method when a reference model is missing is widely acknowledged. In this study, a robust method based on a statistical analysis of the uncertainty associated with Tie Points (TPs) is presented to provide a strict framework for the informed processing of photogrammetric survey data. In the phases of Structure estimation, Structure optimisation, and Dense Cloud generation, the key steps and variables affecting data processing are described. The workflow is then applied to a specific bronze museum finding smaller than 20 cm in size. All tie points that overcome the filtering phase are included in the procedure and for their coordinates the covariance matrix is examined. The error ellipsoid is calculated and the distribution of the lengths of the major semi-axes is analysed to calculate an appropriate tolerance interval which can be used as an indicator of the accuracy of the entire photogrammetric process. Indeed, using the tolerance intervals tool allows for the derivation of a representative indicator that can be compared with the outcomes of other photogrammetric processes while overcoming the ambiguity of statistical indicators that are not representative in the case of a non-normal distribution.

1. INTRODUCTION

The demand for 3D reality-based models is growing in popularity and ambition due to the level of realism and detail expected today. Although hardware and software solutions are getting increasingly innovative and outstanding, somehow fulfilling the challenging task, a question arises regarding the traceability and reliability of virtualisation properties (Brusaporci, 2017). Particularly nowadays, the simplicity of generating 3D geometries due to the spread of completely automated image-based technologies makes it incredibly more complex to trace the quality of outcomes. Thus, leading to the false belief that these 3D models are accurate. This assumption is most likely a result of the common idea that the widespread use of image-based computer programmes automatically guarantees consistency and reproducibility (Remondino et al., 2012). It is therefore expected that modelling processes, and in particular algorithms for handling raw data, can improve significantly (Morena et al., 2019). Hence, despite the extensive use of three-dimensional models, automation and innovation have otherwise made it more challenging to come up with effective approaches for determining the accuracy of a photogrammetric model in relation to its geometric features, resulting in the absence of a common criterion for formalising errors (Puerto et al., 2022).

Moreover, a systematic method to describe the quality of virtualisation has not yet been well developed, despite the extensive research in the literature on evaluating the accuracy of photogrammetry in relation to the acquisition phase's best practices, with a specific focus on systematic errors due to: (i) camera factors (i.e., type, principal point, principal distance, and camera lens distortion coefficients); (ii) imaging settings (i.e., shooting distances, baselines, percentage of photo overlaps, number of overlapping photos, camera intersection angles, and angles of incidence) (Dai et al., 2014).

The widespread way of dealing with the quality task is based on assessing the accuracy of geometric descriptions of an object by comparing homologous models produced through different techniques. The model considered most reliable can be adopted as a reference, calculating the deviations of other digital descriptions from it.

The need for data integration and/or for producing multi-resolution models results in the redundancy of data useful for such distance-based comparative analyses. However, several issues do not make the choice of reference entity trivial: (i) how to unambiguously define the accuracy and reliability of processes and tools that make the reference appropriate; (ii) how to register the models before comparison; (iii) which algorithm is most suitable for quantifying distances. The instrumental resolution, followed by the rigour of the acquisition campaign and the technology employed, are the factors that mainly contribute to making the reference model unambiguous. For the definition of a common reference system, two main types of approaches can be employed: the former based on the identification of specific points, such as targets distributed over the survey scene, and the latter involving the entire object, such as ICP-derived algorithms (Chetverikov et al., 2002). Regarding the choice of algorithm, it seems appropriate to adopt the cloud-to-model (C2M) distance in the case of a polygonal mesh reference model; while for other scenarios, the multiscale model-to-model cloud comparison (M3C2) technique is preferable because it employs parameters that allow for better control over the sources of uncertainty (Lague et al., 2013). When a reference model is not available, the evaluation of the accuracy of the final model is more challenging, as accuracy indicators need to be derived directly from the photogrammetric process. Even when Ground Control Points (GCPs) or Check Points (CPs) are evenly distributed in the scene, basing the entire analysis on a limited number of points is not a robust approach (Luhmann, 2010). Thus, the direct accuracy assessment of the photogrammetric model should be performed by studying the uncertainties related to the Tie Points (TPs). Nevertheless, it must also be taken into account that the quality of the orientation process and associated reconstruction outputs is inevitably affected by the feature extraction and matching techniques (Nocerino et al., 2013), for which not only the number of TP must be considered representative, but also their correctness (Barazzetti, 2017).

Figure 1. "La Sirena di Murgie", National Archaeological Museum of Crotone.

Under these assumptions, a statistical analysis of the covariance matrix associated with the estimated TP coordinates is proposed in this study. The primary objective is to present a potential strategy for the conscious treatment of survey data.
The experimentation is conducted on the close-range photogrammetric process (Luhmann et al., 2014) applied to a bronze ointment vase (Askòs) called "La Sirena di Murgie" (Figure 1). The finding, dated 5th century BC, is preserved at the National Archaeological Museum of Crotone, Calabria Region, and depicts a Mermaid: the hybrid mythological being that, according to Greek iconographic tradition, is a half-human and half-bird (Marino, 2010). The total height of the vessel is 15.3 cm, while the length and width measure 18.7 cm and 8.4 cm.
Although generating models of small objects in the field of Cultural Heritage is often only aimed at AR or Web visualisation applications, in this case, it seemed appropriate to test the validity of the acquisitions that were conducted in a difficult context, such as the museum environment, thus lacking controlled settings such as a well-equipped laboratory (Parrinello et al., 2019). The objective was to guarantee high accuracy with portable and affordable equipment, even in extreme conditions.

2. MATERIALS AND METHODS

The influence of factors governing the orientation process on the quality of a three-dimensional model obtained from photogrammetry is well-known. The essential steps and parameters governing data processing are outlined below. As explained in the following section, they were employed in a specific workflow, highlighting the necessity for a case-by-case approach to the procedures (section 3).

2.1 Structure estimation

During the orientation phase, within the piece of software used – Agisoft Metashape Professional v.1.8.3 (AM) – the pixels in the input photos that are stable under changes in lighting and viewpoint are identified. Then each is assigned a descriptor that places it on the basis of its immediate neighbours.
Similar to the well-known scale-invariant feature transform (SIFT) algorithm, these entities will be later utilised to locate correspondences between the images (feature detection).
After approximately identifying the camera locations (feature matching), the calculation is refined via Self-Calibration Bundle Block Adjustment (structure estimation), solving both the problem of internal and relative external orientation (Nex et al., 2014). The parameters involved in this process are presented below:

- Accuracy, which adjusts the downsampling of the input data according to the original pixel size.
- The key Points limit concerns the maximum number of feature points searched for in each image. The value 0 corresponds to "no limit"; 40,000 is the recommended default value (more points would be obtained by setting higher values, but their reliability would gradually decrease). However, an upper limit still ensures reliable points if high-quality images are employed in the procedure.
- The TPs limit concerns the maximum number of matching points found in each image. The value 0 corresponds "no limit"; 4,000 is a recommended default value. It would be preferable not to limit TPs if one intends to reduce them after alignment.

The recently added option in the software, called Adaptive Camera Model Fitting, if selected, introduces unpredictable solutions in order to find the combination of camera model coefficients that best fits the minimum error model.
Allowing AM control over the decision-making process may lead to lesser error, but it also causes overfitting of the data or the generation of a complex camera model that doesn't exactly reflect the equipment utilised.

2.2 Structure optimisation

The subsequent filtering phase is aimed at obtaining only high-quality TPs and, although subjective in terms of the number of points to be eliminated, is governed by the following parameters:

- Reconstruction Uncertainty, which reflects poor geometric relations between cameras. Reconstruction Uncertainty is calculated as follows:

$$\sqrt{\lambda_1/\lambda_3} \quad (1)$$

  Where λ_1 is the largest eigenvalue of the tie-point covariance matrix, and λ_3 is the smallest. Basically, it is the ratio between the largest and the smallest semi-axis of the error ellipsoid for 3D point coordinates (k = 1). The latter region corresponds to the uncertainty of the point triangulation alone, without considering the propagation of ambiguities arising from interior and exterior orientation parameters.
- Projection Accuracy, which identifies low match accuracy internally assigned by the software, is defined as the average image scale at which picture coordinates of the TP are measured, computed as:

$$\sum_i s_i/n \quad (2)$$

  Where s_i is the image scale at which corresponding projections are identified on the i^{th} image, and n is the number of photos where the TP is detected. This criterion enables the removal of points whose larger size causes their projections to be relatively poorly localised. Projection Accuracy is essentially a representation of the fidelity in the identification of the TP location, given the size of the key points that intersect to create it.
- Reprojection Error, which highlights false correspondences and is defined as the distance between the point on the image where a 3D point can be projected and the original projection of that 3D point detected on the photo and used as a basis for the 3D point reconstruction procedure. The filter evaluates the maximum Reprojection Error in normalised units across all pictures where TP is identified:

$$Max_i \left| x_i' - x_i \right| / s_i \quad (3)$$

Where x_i' is the point projection according to adjusted orientation parameters on the i^{th} image in pixels, x_i represents the measured point projection coordinates on the i^{th} image in pixels, and s_i is the image scale at which corresponding projection is evaluated on the i^{th} image.

Iterative selection and removal procedures are used in the application, as specifically described in the next paragraph. The aim is to improve the estimated internal and external orientation parameters by removing poor TPs. However, each time TPs are deleted, the accuracy of the remaining TPs changes, requiring re-optimisation of the project before going on.
Alongside these considerations on the robustness of TPs, there are those relating to marker accuracy (regarding the image space) and measured control data accuracy (regarding the object space). Precisely, the proportion between these two parameters – the quality of the coordinates identified in the image and the quality of the coordinates measured in the scene – allocates the weight assigned to the markers and TPs throughout the process. In this regard, in AM, the Reference Settings panel allows the inputs of Images Coordinates Accuracy and Measurement Accuracy to be correctly balanced. This setting prevents misleading statistics, as their incorrect estimation generates unrepresentative error models since lens coefficients are very sensitive to these parameters (Over et al., 2021).
The section Measurement Accuracy, which refers to external real world (m), deals with capturing geotags and calculating the accuracy of targets used as CGPs and the overall surveying methods.
Focusing on the Image Coordinates Accuracy (pixel) section, the two parameters to be refined are:

- TPs Accuracy, which corresponds to the normalised accuracy of TPs projections detected at the scale equal to 1, considering a pyramid built applying Gaussian blur. TPs identified on other scales will have accuracy proportional to their ranks.
- Marker Accuracy, which depends to a large extent on how they are positioned in the frames when pinning a marker on a feature, typically on a target, with an indirect correlation with the capture resolution.

In both cases, establishing a priori the proper values is difficult. Therefore, the default values (1 pixel for the TPs Accuracy and 0.5 pixels for the Marker Accuracy) can be used for the first optimisation step and then gradually adjusted during the optimisation iterations. It is worth mentioning that an indicator of the goodness of the realistic estimation related to these metrics is the discrepancy between the error (m) and accuracy (m) values for GCPs and CPs. In the case of a correct estimate, they will converge. For monitoring the refinement of the camera model, several strategies can be implemented:

- Checking the generated reports in the Console Panel during the optimisation process. The number of iterations required to calculate the lens coefficients will be displayed via a sequence of "x". If the string is long and does not get shorter during the procedure, the solution might be divergent or the modelled coefficients may not be sufficient to reach the internal trigger that would otherwise put an end to the optimisation process. In these scenarios, starting from the post-alignment phase is preferable to overfitting the solutions. The reports also show the values of Sigma0, the AM equivalent of the photogrammetric adjustment quality indicator sigma naught (σ_0), which is the Standard Error of Unit Weight (SEUW). The farther the SEUW is from 1, the poorer the estimated TP accuracy. It corresponds to the degree of deviation from an assumed quality or how closely the RMS Reprojection Errors match the predefined error values. However, since the weighting it is based on is not fully documented, using it as the main indication for process monitoring is impractical. Instead, its convergence towards unity becomes valuable in the advanced optimisation stages.
- Observing the number of projections for each image, defined as the number of valid TPs found on a given capture. A specific image will not be used to generate the final outputs if the number of projections for that image is less than 100. To achieve a robust orientation, it would be better to adopt a cautious limit of projections.
- Monitoring the Root Mean Square (RMS) Reprojection Errors in the chunk point cloud property. To increase the robustness of the process, especially when there are blurry photographs or photos with few distinguishing features or textures, the TPs are placed at several map scales. Then, the map scale information is used to weigh the TPs Reprojection Errors by the software. Unfortunately, since the scaling parameters and weighting procedures are not disclosed in AM, the meaning of the mentioned metric becomes convoluted: of the two reported RMS Reprojection Errors related to orientation, the weighted value is the first one in the units of key point scale, and the unweighted one, written in pixels, is the second one.
  The optimal unweighted value is generally less than 0.3 pixels. If too many TPs are deleted during the subsequent filtering steps, the camera model may deviate from a particular solution, increasing the unweighted RMS Reprojection Error.

Estimating the TPs covariance matrix, which is related to the execution of the Bundle Block Adjustment, is a way to visualise the uncertainty in the camera models after each optimisation cycle. The evaluation outcomes can be examined by switching the TPs cloud display mode to a specific one (Figure 2).
The vector associated with each TP indicates the direction and magnitude of the error for the TP estimated position (its three components correspond to the semi-axes of the error ellipsoid with k = 1 determined by the covariance matrix).

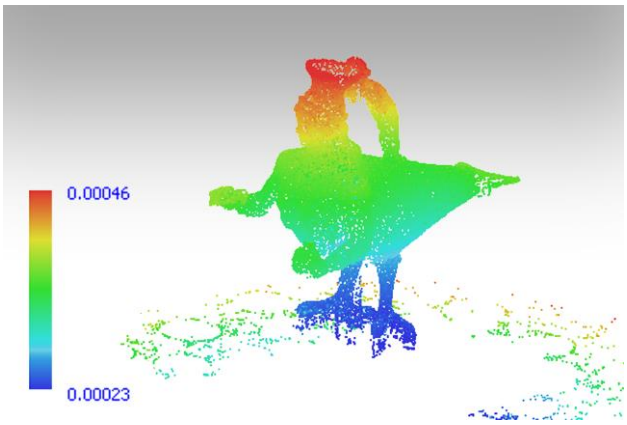

Figure 2. TPs covariance view mode, after Gradual Selection filtering, in Agisoft Metashape SfM software.

2.3 Dense Cloud generation

Dense Clouds are reconstructed from depth maps computed through pairs of overlapping images, identified by their internal and relative external orientation parameters, which are estimated using Bundle Block Adjustment. Multiple pairwise depth maps generated for each camera are merged into a combined depth map, employing excessive information in the overlapping regions to discard wrong depth measurements. The products are transformed into partial dense point clouds, which are subsequently merged into a final model, thanks to an additional noise filtering step applied in the overlapping regions. The normals in the partial dense point clouds are calculated using plane fitting to the pixel neighbourhood in the combined depth maps, and the colours are sampled from the images. The number of contributing maps is recorded and stored as a confidence value for every point in the final dense point cloud. This number will be used later to perform additional filtering. A list of parameters governing the procedure is explored below:

- Quality, which affects the detail and accuracy of the reconstructed geometries as well as the processing time, depending on the number of depth maps calculated.
- Depth filtering, which removes outliers from the cloud due to aspects such as noisy or badly focused images by checking them within the raw depth maps. This step is performed using a connected component filter that operates on segmented depth maps based on the pixel values.
- Whether and how the process associates the RGB information to the individual points.
- Point confidence calculation, which enumerates how many depth maps are generated for each point so that further filtering can be performed.

Once the Dense Cloud has been reconstructed, noise removal can be performed. In addition to the interactive intervention, which is useful for deleting badly located elements, a filter can be applied according to the number of depth maps per point. The results of this operation depend primarily on the pattern chosen to capture the frames. The algorithm is based on a non-linear selection scale ranging from 0 to 255, where lower values indicate fewer depth maps involved in the reconstruction of the 3D position of the point (Figure 3).
This last operation ends the photogrammetric process.

2.4 Uncertainty assessment

To provide a rigorous framework for the informed handling of photogrammetric survey data in the absence of a reference model, a statistical analysis of the uncertainty associated with Tie Points (TPs) in object space can be performed. Sources of uncertainty related to instrumentation, techniques, their integration, and the statistical description of their distributions can be identified. In this way, their propagation law and a dispersion indicator or an interval that effectively summarises the accuracy of the survey can be obtained. However, finding sources of uncertainty and combining them is not easy. There is no single unambiguous solution, which depends heavily on the detection methods and the transformation of the raw data into a reality-based model. Given that it is impossible to list and describe every possible combination, an inductive approach is preferable, starting from a practical case and highlighting the distinctive aspects that affect the decisions. Nonetheless, evaluating the sources of uncertainty in photogrammetric operations cannot rely just on the coordinates of the codified targets.

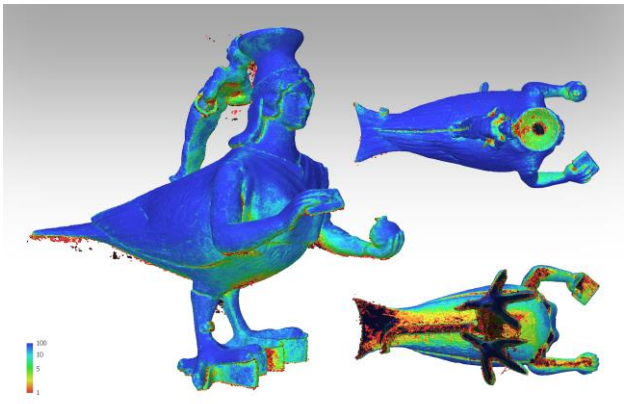

Figure 3. Point Confidence view mode, before Dense Cloud filtering, in Agisoft Metashape SfM software.

In fact, as a hypothetical solution, the uncertainty related to the Check Points (CPs), expressed by the difference between input and output data, and the uncertainty connected to the support survey may be combined. However, the two sources of uncertainty have a correlation, which would need further elaboration to be adequately described. Albeit the points are evenly distributed around the picture, basing the entire analysis on a small number of points is unquestionably not a robust strategy. The same reasoning may be extended to GCPs, with the aggravating factor that their coordinates, utilised to resolve the structure optimisation phase, generate a reduction in the error associated with them, making them inappropriate for this kind of evaluation.
The previous paragraphs clarified how crucial it is for a rigorous photogrammetric procedure to trace the accuracy of all input data (related to both the object space and the picture space). After the orientation optimisation phase, the covariance matrix connected to the predicted coordinates for the TPs in the object space is exported using an *ad-hoc* Python script we developed. From the matrix, the standard error ellipsoid and the configuration corresponding to a probability of containing the theoretical mean value of the coordinates of 97.75% (k =3) are derived. Of this last ellipsoid, the major semi-axis is considered, and an appropriate tolerance interval is calculated by studying its distribution for all the Tie Points. This tool allows estimating, from a sample, the extremes that contain a certain percentage of a population associated with a specific level of confidence. The approach is therefore considered reasonable, noting that the TPs will constitute only a part of the final photogrammetric cloud. To be fair, the dense image-matching phase and its algorithms should also be involved, but this would become too complicated according to the current state-of-the-art. We will then merely employ the results of the Structure from Motion step here. Section 3.2 elaborates on the procedure for determining these ranges.

3. APPLICATIVE WORKFLOW AND RESULTS

3.1 The SfM process

The workflow is applied to data from Nikon D800E SRL camera, equipped with the AF-S Micro NIKKOR 60mm f/2.8G ED, mounted on a tripod in landscape orientation and operated by remote control.
The frames captured have a size of 7360 × 4912 pixels. Considering the focused distance from the object to be 0.6 m, in order to frame the object entirely in a single shot and limit the time and number of acquisitions, a GSD of about 0.04 mm/px

was determined. In addition, a Depth Of Field of 12 cm was achieved with an f/32 aperture, using a Circle Of Confusion of 0.015 mm (about half the conventional values for such equipment, to be on the safe side). The ISO sensitivity values were kept low (100) to avoid digital noise. The shutter speed was estimated to obtain the correct exposure, requiring a remote trigger to avoid vibration. Manual focusing was employed since the main distance would not change during the acquisition campaign, as the focus would only be adjusted just once before starting to capture the photos. The object, placed on a turntable, was rotated in front of the camera approximately every 15 degrees at three different camera angles. The complete dataset consists of about 73 images, ready to be oriented. Files imported into AM for processing were converted from the proprietary .NEF format to lossy JPG sRGB 24-bit type, which is a good compromise between file size and quality. An initial estimate of the quality of such data was based on sharpness. The method consists of comparing the contrast gradients in the most peculiar regions of the source images, taking into account both the originals and the images with the Gaussian blur filter applied. All frames with a rate below 0.5 were excluded from the process.

Then, the calibration table was examined. The software produces a single calibration group since the loaded photos for each project are identical in terms of size, focal length, and other factors derived from the exchangeable image file (EXIF). Therefore, the camera type (frame), pixel size, and focal length were checked. No pre-calibration data were available for the camera. No position data were known, and therefore a local system was set for the cameras.

The orientation phase was governed by the parameters shown in Table 1, so as to use a higher limit for Key Points due to the high image quality and not limiting the number of TPs as they will be reduced after alignment (section 2.1).

Parameter	**Setting**
Accuracy	Highest
Generic preselection	Yes
Reference preselection	Sequential
Key point limit	100.000
TP limit	0
Exclude stationary TPs	No
Guided image matching	No
Adaptive camera model fitting	No
Matching time	6 minutes 12 seconds
Alignment time	55 seconds

Table 1. Orientation parameters.

The import of external references is the subsequent operation to optimise the configuration and address the issue of absolute external orientation. A metric reference – that employs a non-repetitive pattern and 56 targets whose coordinates are known by design – was placed on the rotary table and used as an external reference. The targets were visible in the photos, so AM could locate them automatically by simply choosing the type of artificial object placed in the scene (16-bit circular encoded targets). After the recognition, a visual check was performed, followed by an eventual optimisation. Not all markers were used as GCPs to solve the absolute external orientation; some would indeed be discarded from this process, to be later employed as CPs to validate the optimisation results, ensuring a homogeneous distribution of the two groups in the detected scene. The optimisation phase started here (section 2.2).

At this stage, activating every coefficient might reduce error. Still, it also could cause overfitting of the data or the generation of an overly complex camera model that does not accurately reflect the equipment used, leading to underestimating or increasing actual error. Therefore, the affinity and nonorthogonality, also known as skew coefficients (B_1 and B_2), should be initially suppressed and only included in case of an inflated RMS Reprojection Error value indicating the possible presence of distortion related to these phenomena. The same applies to the radial distortion coefficient K_4.

After duplicating the chunk containing the data processed so far to store a backup copy and checking the correctness of the input reference data, the first optimisation cycle was performed by selecting the appropriate coefficients and estimating the covariance of the TPs. Along the procedure, the number of projections per frame and the unweighted RMS Reprojection Error were monitored. For calibrating the accuracy of the image coordinate, the latter is helpful mainly when associated with RMS Reprojection Error on markers and mean Key Point size.

The error reduction procedure first consists in removing points resulting from poor camera geometry through the Reconstruction Uncertainty filter. High Reconstruction Uncertainty is typical for points identified through nearby photos with a small baseline. They can noticeably deviate from the object's surface, causing noise in the cloud. While removing such elements should not affect the optimisation accuracy, it may be useful to delete them to improve the visual appearance of the reality-based model. A Reconstruction Uncertainty of 10 is roughly equivalent to a good base-to-height (or base-to-distance) ratio of 1:2.3 (parallax angle of about 23°), whereas 15 is almost equal to a marginally acceptable ratio of 1:5.5 (about 10°) (Gujski et al., 2022). Our workflow reiterated the filtering of TPs with this parameter twice, taking care not to remove more than 50 per cent of the TPs. The optimisation process used the same lens coefficients as the preliminary stage. After removing the TPs, a single iteration is sufficient if a Reconstruction Uncertainty of 10 is achieved in the first attempt and less than 50% of the TPs are selected. Repeated filtrations have diminishing returns and may overfit the camera model before more poor-quality TPs can be removed. Once again, we monitored the status indicators used during the first cycle.

The second part of the error reduction procedure aims to remove points based on Projection Accuracy. AM saves an internal accuracy and scale value for each TP as part of the correlation process. The highest reliability points are assigned to level 1 and are weighted based on the relative size of the pixels. A TP allocated to level 2 has twice as much inaccuracy as level 1. Not all projects can tolerate removing points at a level of 2 to 3, particularly if the images have been compressed or their quality is compromised due to noise or blur. A gradual selection level of 5 or 6 may be the best that can be obtained. The threshold limit can be defined by remembering that this filtering phase should not eliminate more than 50 per cent of the TPs inherited from the previous one, targeting two cycles. Lens coefficients used in the process are the same as in the second step. Repeated applications show reduced returns and may overfit the camera model. If the project can support an initial filtering of level 3 without selecting more than 50 per cent of the TPs, a single optimisation after deleting the points is sufficient.

In our applications, the threshold value was 2. Once again, we monitored the process quality indicators before going any further.

The third step of the error reduction procedure concerns the Reprojection Error, which, unlike the previous ones, is directly related to the parameters governing internal and external orientation. It is, therefore, necessary to correctly define the accuracies for the measurements and image coordinates before proceeding. As far as measurements are concerned, these are known by design, so their accuracy is only related to the printing process for their physical realisation. Printer accuracy, linked to critical dot size, shape and position accuracy, is not evaluated

here. However, for inkjet printers, the error is 10 µm on average (Creagh et al., 2003), then this value can be taken as a lower limit for the marker accuracy parameter (m) in the measurement accuracy section.
Concerning the image coordinates parameters, the marker accuracy can assume values below 0.5 pixels if the identification procedure is rigorous, as in our case, and supported by automatic extraction algorithms. Considering the high resolution and the quality of the captures, we assumed the marker accuracy value equal to 0.05 pixels, verifying that the discrepancy between the error (m) and the accuracy (m) associated with the GCPs and CPs does not grow uncontrollably during the next optimisation phase.
As for the accuracy of TPs we gradually reduced it by checking the convergence of Sigma0 to 1. As a rule, we always give markers more weight than tie points. After the calibration of these parameters, we performed a new optimisation cycle and tested if the value of Sigma0 tended towards 1. If this did not happen, it was necessary to pay attention and, eventually, repeat the previous phases. Finally, we achieved a value of 0.1 pixels for the accuracy of TPs.
High reprojection error usually indicates poor localisation accuracy of the corresponding projections at the point matching step. It is also typical of false correlations. The threshold set for filtering in our applications is 0.3 pixels, not reached directly but by operating in successive cycles and selecting only 10 per cent of the tie points at a time. After the first cycle of this phase, we checked that the error of the markers (m) did not exceed the accuracy (m). In this case, it is convenient to stop the process, revaluate, and correct the a priori reliability of the image coordinates in light of new data, such as the RMS reprojection error. We then checked whether Sigma0 converged to 1. The selected lens coefficients were again those chosen for the previous steps. If confluence does not occur, it is possible to consider involving other coefficients and use additional corrections.
At the end of all the cycles of the task, we check the number of projections for each frame, the effective reduction of the iterations within the succession of the cycles, the survival of at least 15-20% of the original tie points and the possible overfitting of the camera model. If no inconsistencies or anomalies emerge, the orientation process (so-called Structure from Motion – SfM) can be considered concluded (Table 2).

Parameter	Setting
Reconstruction Uncertainty filter	10
Projection Accuracy filter	6
Marker Accuracy (m)	0.0005
Marker Accuracy (px)	0.05
Tie Points Accuracy (px)	0.1
Reprojection Error filter	0.3
Control Points error (m)	0.0003
Check Points error (m)	0.0004
Sigma0	0.97

Table 2. Error reduction and reference system parameters consistent with input data.

In order to be thorough, the relationship between the calibration coefficients of the camera was examined (Remondino et al., 2006; Tang, 2016). This is an effective approach to check for overfitting and may show that some of these parameters are really not relevant to represent the equipment.
The set used for the application is an 8-term set derived from the one originally formulated by Brown (1971), including internal orientation parameters of main distance and principal point offset, as well as the three coefficients of radial and two of tangential distortion.

	Value	Error	F	Cx	Cy	K1	K2	K3	P1	P2
F	12598.7	0.672439	1.00	0.04	-0.22	-0.16	0.22	-0.18	0.03	-0.06
Cx	90.3494	0.986818		1.00	-0.03	-0.07	0.06	-0.08	0.97	-0.06
Cy	-19.782	1.30691			1.00	-0.28	0.37	-0.28	-0.02	0.55
K1	-0.0405119	0.000623855				1.00	-0.92	0.83	-0.09	-0.11
K2	0.731	0.0138645					1.00	-0.97	0.08	0.08
K3	-1.60927	0.0861545						1.00	-0.09	-0.07
P1	0.00267437	2.95114e-05							1.00	-0.05
P2	-0.000827401	2.39623e-05								1.00

Figure 4. Correlation analysis between camera calibration coefficients.

The correlation between the radial coefficients (Figure 4) is physiological and depends on the structure of the model itself; for this reason, we can neglect it. The decentring distortion is also strongly projectively coupled with the principal point offset, and, in general, this relationship increases with focal length or when a poor convergence of the image axes occurs. These reasons confirm that the correlation is negligible.
Based on the exterior and interior orientation parameters, the dense image matching phase (section 2.3) was then performed.
A list of the parameters governing the procedure is summarised in Table 3.

Parameter	Setting
Quality	High
Depth filtering	Moderate
RGB data	Yes
Points Confidence	Yes

Table 3. Error reduction and reference system parameters consistent with input data.

The high-quality setting used in our applications implies preliminary image size downscaling by a factor of 4 (2 by each side), a good compromise between model density and computational load. The moderate depth filter was selected to preserve the small, discernible details in the scene to be reconstructed.
The noise remotion was carried out after the generation of the Dense Cloud. Unfortunately, the confidence-based filter working process is blinded, therefore we rely on the indications provided by the developers, according to which the relevant part of the noise belongs to the range 1-5. After filtering, the SfM process can be considered terminated.

3.2 Towards tolerance limit formation

After deriving the error ellipsoid with k = 3 from the covariance matrix and selecting the major semi-axis, its distribution was studied to build an appropriate tolerance interval. This tool allowed us to estimate, from a sample, the extremes that contain a certain percentage of a population (p) with a specific level of confidence (α). The appropriate construction of tolerance intervals takes into account the type of distribution, making it possible to compare different statistical models, involving several steps: (i) test for normality, which, if satisfied, leads to the calculation of normal tolerance limits; (ii) searching for normalising transformation, to be implemented when the distribution is not normal. If it deviates only slightly, an

acceptable transformation is found, and normal tolerance limits can be calculated from transformed data; (iii) alternative distributions to be performed when the transformation approach fails by searching for a good fit so that tolerance limits from that distribution can be calculated; (iv) non-parametric tolerance limits when the others approach fails.
In any case, it must be ensured that the sample size is adequate for the statistical treatment. All intervals are computed for a population percentage of p = 95% and a confidence value of α = 95%. It should be borne in mind that, in the case of non-parametric tolerance limit calculations, only one of these two features can be defined, the other being estimated downstream of the procedure. In that case, preventive treatment of the distributions may be necessary to eliminate possible outliers, e.g., by constructing a box-plot diagram. Precisely the latter approach was applied to the case study fixing α = 95%, and Figure 5 shows the actual distribution of the major semi-axes of the ellipsoids with k = 3. The approach, therefore, seems reasonable, considering that the tie points will constitute only a part of the final photogrammetric cloud. To be fair, the dense image-matching phase and its algorithms should also be involved, but this would become too complicated. Thus we chose to employ the Structure from Motion classic workflow for the subsequent steps. Since the semi-axis length is a positive definite quantity, we outlined a one-sided interval, obtaining an upper tolerance limit of 1.35 mm, which can be used as an indicator of the accuracy of the entire photogrammetric process.

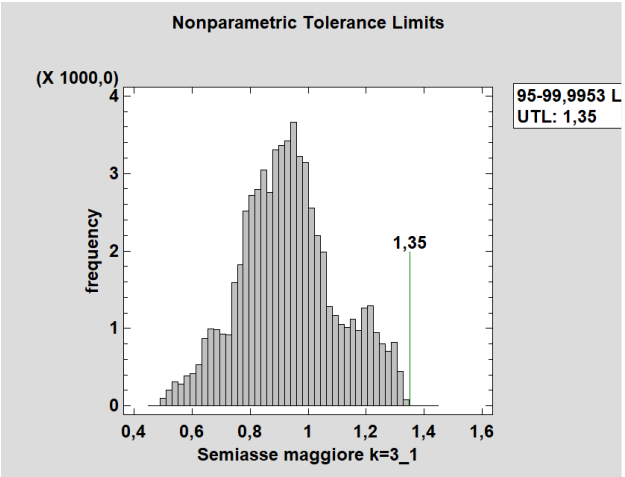

Figure 5. Distribution of lengths of major semi-axes with k = 3 and upper tolerance limit.

4. CONCLUSIONS

The question of photogrammetric survey data traceability is becoming more relevant than ever. There is still ambiguity about how to define the error, given the numerous parameters that control the processing workflow and the countless factors that can affect it. In the absence of more reliable analyses, the error on GCPs and CPs is not representative and does not justify any defects in the modelling, leaving the reasons for deficiencies hidden and not allowing for process improvement.
This study aimed to address this demand by attempting to achieve accuracy traceability while also providing maximum adaptability and flexibility to the requirements of the potential case studies and proposing possible solutions. In fact, to define the project requirements and assess the accuracy of geometric features, a procedure has been proposed. This is based on the reliability of the data processing that was observed to create source-based virtualisation.
One advantage of the suggested methodology is that it ensures a level of accuracy in the modelling phase can be easily traced. Thus, making the framework repeatable and compatible with data from various sensors.
In detail, the procedure is robust, as it involves all tie points that survive the filtering phase. For them, it examines the covariance matrix, deriving the error ellipsoid and analysing the distribution of the lengths of the major semi-axes. This study was conducted in a rigorous manner. The distribution hardly ever follows the normal curve; therefore, position and dispersion indicators, such as mean and standard deviation, do not always describe the phenomenon effectively. In contrast, statistical tolerance intervals make it possible to overcome this ambiguity and derive a representative indicator that can be compared with the results of other photogrammetric processes.
Future developments will concentrate on a more detailed description of the survey-related uncertainty components and the opportunity to specify accuracy levels during a preliminary design phase in accordance with the client's desired tolerances.

REFERENCES

Brusaporci, S., 2017. The Importance of Being Honest: Issues of Transparency in Digital Visualization of Architectural Heritage. In Ippolito, A. (ed), *Handbook of Research on Emerging Technologies for Architectural and Archaeological Heritage*. IGI Global: Hershey, Pennsylvania, USA, pp. 66-93.

Remondino, F., Del Pizzo, S., Kersten, T., Troisi, S., 2012. Low Cost and open-source solutions for automated image orientation – A critical overview. *Lecture Notes in Computer Science*, 7616, 40-54.

Morena, S., Barba, S., Álvaro-Tordesillas, A., 2019. Shining 3D Einscan-Pro, application and validation in the field of Cultural Heritage, from the Chillida-Leku Museum to the Archaeological Museum Of Sarno. *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci.*, XLII-2/W18, 135-142.

Puerto, P., Heißelmann, D., Müller, S., Mendikute, A., 2022. Methodology to Evaluate the Performance of Portable Photogrammetry for Large-Volume Metrology. *Metrology* 2, pp. 320-334.

Dai, F., Feng, Y., Hough, R., 2014. Photogrammetric error sources and impacts on modeling and surveying in construction engineering applications. *Visualisation in Engineering* 2(2), pp. 1-14.

Chetverikov, D., Svirko, D., Stepanov, D., Krsek, P., 2002. The Trimmed Iterative Closest Point algorithm. *Object recognition supported by user interaction for service robots* 3, pp. 545-548.

Lague, D., Brodu, N., Leroux, J., 2013. Accurate 3D comparison of complex topography with terrestrial laser scanner: Application to the Rangitikei canyon (N-Z). *ISPRS Journal of Photogrammetry and Remote Sensing* 82, pp. 10-26.

Luhmann, T., 2010. Close range photogrammetry for industrial applications. *ISPRS Journal of Photogrammetry and Remote Sensing* 65, pp. 558-569.

Nocerino, E., Menna, F., Remondino, F., 2014. Accuracy ff typical photogrammetric networks in Cultural Heritage 3D modeling projects. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, Volume XL-5, pp. 465-472.

Barazzetti, L., 2017. Network Design in Close-Range Photogrammetry with Short Baseline Images. *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci., IV-2/W2*, 17-23.

Luhmann, T., Robson, S., Kyle, S., Boehm, J., 2014. *Close-Range photogrammetry and 3D Imaging*. De Gruyter: Boston, Massachusetts, USA.

Marino, M., 2010. *Le Sirene di Kroton*. Museo Archeologico Nazionale di Crotone: Crotone, Italy, p. 23.

Parrinello, S., La Placa, S., 2019. Vectorialization practices of the image drawing of the floor mosaics of the Basilica of Nativity in Bethlehem, *SCIRES-IT, SCIentific RESearch and Information Technology*, Volume IX (II), 95-104.

Nex, F., Remondino, F., 2014. UAV for 3D mapping applications: a review. *Applied Geomatics, 6(1)*, 1-15.

Over, J.R., Ritchie, A.C., Kranenburg, C.J., Brown, J.A., Buscombe, D., Noble, T., Sherwood, C.R., Warrick, J.A., and Wernette, P.A., 2021. *Processing Coastal Imagery With Agisoft Metashape Professional Edition, Version 1.6 - Structure From Motion Workflow Documentation.* U.S. Geological Survey: Reston, Virginia, pp. 23-25.

Gujski, L.M., di Filippo, A., Limongiello, M., 2022. Machine Learning Clustering for Point Clouds Optimisation via Feature Analysis in Cultural Heritage. *Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci., XLVI-2/W1-2022*, 245-251.

Creagh, L.T., McDonald, M., 2003. Design and Performance of Inkjet Print Heads for Non-Graphic-Arts Applications. *MRS Bulletin, 28*, 807-811 (2003).

Remondino, F., Fraser, C., 2006. Digital camera calibration methods: considerations and comparisons. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, 36(5)*, 266-272.

Tang, R., 2013. *Mathematical Methods for Camera Self-Calibration in Photogrammetry and Computer Vision*. Deutsche Geodätische Kommission: München, Germany.