Global Block Adjustment for Mosaicked Stereoscopic Satellite Imagery

Michaël Erblang [1], Emelyne Saulnier [1], Guillaume Laurent [1], Nicolas Delaygue [2], Fabrice Buffe [2], Alice Latourte [2], Mathilde Jaussaud [3], Noémie Bricout [3]

[1] Thales Services Numériques (TSN), 290 Allée du Lac, Labège, France
(michael.erblang, emelyne.saulnier, guillaume.laurent)@thalesgroup.com
[2] Centre National d'Etudes Spatiales (CNES), 18 avenue Edouard Belin, Toulouse, France
(nicolas.delaygue, fabrice.buffe, alice.latourte)@cnes.fr
[3] Institut national de l'information géographique et forestière (IGN), 18 avenue Edouard Belin, Toulouse, France
(mathilde.jaussaud, noemie.bricout)@ign.fr

Keywords: Satellite CO3D, block adjustment, tie point, SRPs, georeferencing, geometric refinement

Abstract

Satellite imagery acquired over large areas from multiple viewpoints introduces subtle geometric misalignments that degrade the quality of derived products such as Digital Surface Models (DSMs). This paper presents a global block adjustment workflow designed to correct these errors across overlapping stereo acquisitions from the "Constellation Optique 3D" (CO3D) constellation, which captures Earth's surface at 50 cm resolution. The proposed pipeline operates in three stages: individual acquisition refinement using Space Reference Points (SRPs) as Ground Control Points; tie point extraction between overlapping scenes through two-pass image correlation; and a weighted global spatio-triangulation simultaneously optimizing attitude biases, attitude drifts, and per-satellite magnification parameters. Applied to a large stereo acquisition dataset over the Aorounga crater, Chad, the method demonstrates strong geometric performance. The results highlight that careful parameterization — combining observation weighting, n-tuple point filtering, and per-satellite sensor refinement — is key to producing accurate, geometrically consistent large-scale mosaics from bi-satellite stereo imagery. This paper does not include the in-orbit performances due to confidentiality agreement.

1. Introduction

The increasing availability of high-resolution satellite constellations opens new opportunities for large-scale Earth observation, enabling the production of detailed 3D surface models and orthophotos over vast territories. However, mosaicking multiple stereo acquisitions remains a geometric challenge: each image is captured from a slightly different viewpoint, and sensor attitude imperfections introduce subtle but cumulative localization errors across the scene. Without correction, these misalignments degrade the consistency of derived products.
This paper presents a global block adjustment workflow developed in the context of the CO3D constellation, addressing this challenge through a three-stage pipeline combining individual image refinement, tie point extraction, and weighted global optimization.

2. Image Dataset

The dataset consists of multiple synchronous stereo acquisitions in the Red-Green-Blue (RGB) bands from the pair of satellites CO3D1 and CO3D2. Each image is covering approximately 5×7 km² with a spatial resolution of 50 cm.
CO3D mission is based on a constellation of four optical satellites, developed by Airbus Defence and Space (ADS) in partnership with CNES, and launched on July 25, 2025, at the Guiana Space Center.
The CO3D constellation is divided into two pairs of satellites positioned on either side of a near-polar orbit at an altitude of 502 kilometers. Within each pair, the satellites are separated by 100 kilometers along-track to produce a stereoscopic view of the Earth's surface and enable its 3D reconstruction. The two pairs increase revisit frequency and coverage over large areas of interest, while the synchronous stereo acquisition mode is particularly suited to the 3D reconstruction of moving objects such as vehicles or ocean waves (cf. Lebègue, L. et al., 2026).
CO3D employs a matrix sensor imager, capturing a complete image of a ground rectangle in a single acquisition. The primary instrument onboard each CO3D satellite is its imaging sensor, based on a RGB pixel array (Bayer filter) and a separate array sensitive to the near-infrared spectrum (cf. Buffe, F. et al., 2026).

Each acquisition is calibrated using Ground Control Points (GCPs) to ensure local alignment. These points are a global reference layer called Space Reference Points (SRPs), consisting of a set of image chips spaced every 2km², with a 3D center coordinate and an accuracy at 3m CE90. The SRP database is built from "Satellites Pour l'Observation de la Terre" (SPOT6/7) and is archived by the French mapping agency (IGN) and Airbus (cf. Chandelier, L. et al., 2020).

Finally, the Copernicus Digital Surface Model (DSM), providing global coverage at a 10-meter altimetric accuracy, serves as our reference DSM for calibrating all observation points and especially the altitude of our Tie Points (TPs).
Copernicus is the Earth observation component of the European Union's Space programme. Copernicus is served by a set of dedicated satellites (the Sentinel families) and contributing missions (existing commercial and public satellites).

3. Methodology

3.1 Main pipeline

The refinement process comprises three main stages:

- Individual refinement of each acquisition, correlating with GCPs and refining through spatio-triangulation.

https://doi.org/10.5194/isprs-archives-XLIX-B2-2026-25-2026

- Correlation between stereo and non-stereo images to link acquisitions together.
- Global refinement through uncertainties settings, weighting and attitude and sensor refinement via spatio-triangulation.

3.2 Individual refining with GCPs

Each pixel in the image has an associated ground location. A geometric model (physical or simplified) is used to project each pixel (i, j) onto the ground in a geographic reference system. The geometric model for each acquisition is then refined using only SRP through image correlation. Refining the geometric model consists in optimizing its parameters by matching image pixels to ground locations provided by SRP, resulting in precise local alignment. Each acquisition is refined using only SRP data. A good distribution of SRPs in the image improves the accuracy of per-acquisition refinement.

For this purpose, the acquisitions are projected into the geometry defined by the SRP, and a filter is applied to equalize the blur levels across images. Each SRP is defined by a stack of thumbnails, each of which is correlated with the acquisition.

Correlations are performed using a sub-pixel two-dimensional stereo matching algorithm applied to an image pair, based on a local matching approach (cf. Cournet, M et al., 2016). For each SRP thumbnail, a unique correspondence is estimated at the central pixel of the thumbnail. This correlation is carried out using a sliding window approach: for which a small patch is extracted and compared to all candidate patches within a search area in the target image, using a similarity score - typically the Normalized Cross-Correlation (NCC) (cf. Figure 1).

- *Correlation window size*: size in pixels of the patch extracted around reference point; a larger window is more robust to noise but less sensitive to fine spatial details.
- *Search area size*: extent of the region in pixels in the target image within which the best match is sought; it must be large enough to encompass the expected geometric offset between the two images.
- *Correlation score threshold:* minimum similarity score below which a match is rejected, filtering out unreliable correspondences. NCC score is a scalar value between -1 and 1.

For each stack of thumbnails, the median correlation is selected, yielding one correspondence per SRP.

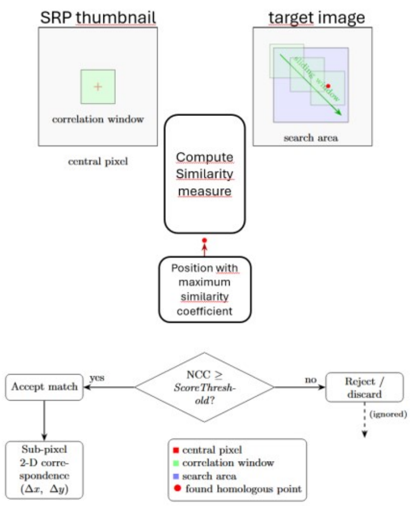

Figure 1: Principle of correlation

The correlation results are saved in two files: the first contains the image values (including the pixel coordinates of the image measurements with their associated uncertainties and weights), and the second contains the ground values (including the planimetric coordinates of the ground measurements, the planimetric and altimetric coordinates of the ground estimates, and the planimetric precision and weight of the ground measurements). This step is called observation context generation.

The geometric modeling error considered here corresponds to a satellite attitude error, represented as a bias. The previously generated observation context serves as input to a refinement procedure that estimates a constant roll, pitch, and yaw rotation, minimizing the residuals between the acquisitions and the SRP patches.

And so, for each acquisition:

$$\sigma_{SRPmodelsAttitude} = (\sigma_{SRProll\ x0}, \sigma_{SRPpitch\ y0}, \sigma_{SRPyaw\ z0})$$

3.3 Image Correlation between stereo pairs and adjacent scenes

3.3.1 Geometric preprocessing for correlation: Once each image has been individually refined through correlation with SRP chips – and so it is now tied to its ground-truth coordinates - they must be geometrically linked to one another to ensure the overall consistency of the mosaic, while preserving their individual anchorage to the ground. Each image must be connected to:

- its stereo pair,
- neighboring acquisitions.

In both cases, common points are extracted from overlapping acquisitions using linear image correlation, providing relative geometric constraints between stereo pairs and/or adjacent scenes.

The extraction process follows three steps. First, all pairwise intersections are computed from the footprints of each acquisition. Then, for each intersection, the images involved are resampled into the geometry of a common reference image, ensuring that the correlation is performed within a consistent coordinate frame.

3.3.2 Correlation settings: As mentioned in Section 3.2, the correlation itself is carried out using a sliding window approach: for each point on the reference image grid, a small patch is extracted and compared to all candidate patches within a search area in the target image, using a similarity score. One additional key parameter is introduced here, as the size of our observation context depends strongly on it.

- *Correlation grid spacing* is the spacing in pixels between two consecutive grid points; a finer spacing yields more tie points but increases computational cost.

3.3.3 N-tuple points in observation context: Once the observation framework contains both ground GCPs and TPs, each point must be characterized by the number of images in which it is observed. A point seen in exactly two images forms a pair, or tuple, while a point observed in more than two images is referred to as an n-tuple point. A point is assigned a multiplicity N if it belongs to an N-tuple point.

This characterization serves two purposes:

- it identifies the most geometrically constraining points within the observation framework.
- it distinguishes between stereo linking points — with a multiplicity of at most 2 — and junction points between adjacent scenes, which exhibit a multiplicity greater than 2.

https://doi.org/10.5194/isprs-archives-XLIX-B2-2026-25-2026

Together, these points constitute the observation framework. SRPs serve as ground anchors for the images, as they carry precise terrain measurements and are directly correlated to the imagery. Tie points, on the other hand, link images to one another but provide different levels of geometric constraint depending on their multiplicity (cf. Figure 2). In the context of a global adjustment, points with a multiplicity of 4 or more form the backbone of the block and are all equally critical. By contrast, multiplicity-2 points derived from stereo pairs are far more numerous but contribute primarily to ensuring internal stereo consistency.

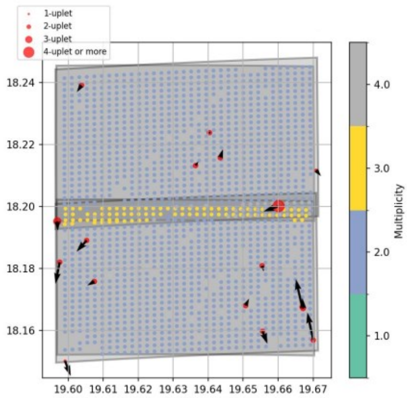

Figure 2: multiplicity of GCP (red dot) and TP (colored)

It is therefore necessary to balance the observation framework by reducing the number of tie points. This can be achieved through several strategies: increasing the correlation grid spacing during the extraction step, filtering the observation framework after correlation, applying observation weighting, or combining several of these approaches (cf. Figure 3). A good balance would be 1 GCP for every 10 TPs.

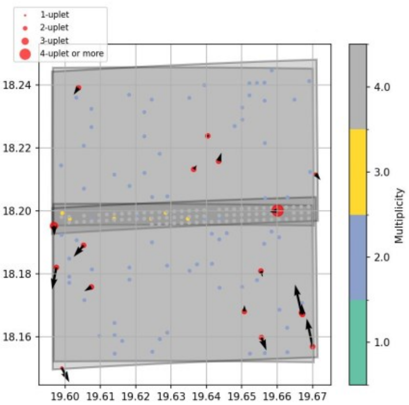

Figure 3: Filtering TP by multiplicity

3.4 Global refining

3.4.1 Pipeline overview: Global refinement consists of optimizing model parameters to improve both relative precision (between models) and absolute accuracy (with respect to the ground).

Refining process requires:

- Image measurements with their associated uncertainties σ_{img} and their weights ω_{img}.
- Ground measurements with their associated uncertainties σ_{gnd} and their weight ω_{gnd}
- Models with deviations characterized by an uncertainty σ_{models}
- Digital Elevation Model (DEM) deformation with associated unknowns

The resulting outputs are:

- Modeling unknowns (deviations) stored for each state.
- Image positions estimated by the corrected model.
- Ground positions estimated by the corrected model.
- Updated DEM deformation unknowns.

3.4.2 Image and ground accuracy: The initial step involves defining the optimal parameters for our image and ground measurements. Since our dataset consists of both Ground Control Points (GCPs) and tie points (TPs), a clear distinction between these two classes is mandatory to assign appropriate uncertainties. For image measurements, uncertainties depend on the accuracy of our correlation maps (correlation noises included) of our GCPs and our TPs. For our ground measurements, uncertainties depend on the native ground accuracy of our GCP and of the referenced MNS. And so for every measurements i:

- Image measurements with their associated uncertainties:
  $\sigma_{imgGCP_i} = (\sigma_{colGCP_i}, \sigma_{rowGCP_i}) \approx$ GCP correlation accuracy
  $\sigma_{imgTP_i} = (\sigma_{colTP_i}, \sigma_{rowTP_i}) \approx$ TP correlation accuracy
- Ground measurements with their associated uncertainties
  $\sigma_{gndGCP_i} = (\sigma_{xyGCP_i}, \sigma_{zGCP_i}) \approx$ GCP accuracy
  $\sigma_{gndTP_i} = \sigma_{zGCP_i} \approx$ Referenced MNS accuracy

(noting that TPs do not have planimetric ground measurements).

3.4.3 Image and ground weighting: Weighting the points within our observation context serves as a balancing method. While all points are weighted at one by default, it can be advantageous to assign higher weights to points with high multiplicity or to GCPs to emphasize their contribution relative to other points. It is worth noting that weighting should be applied only to ground measurements, as over-constraining image measurements may lead to poor refinement performance. For example, a normalized weighting scheme can be implemented for each point i in the observation context:

$$\omega_{imgGCP_i} = 1$$
$$\omega_{imgTP_i} = 1$$
$$\omega_{gndGCP_i} = 1 / \sqrt{(\text{number of GCP})}$$
$$\omega_{gndTP_i} = 1 / \sqrt{(\text{number of TP})}$$

Consequently, each GCP or TP is reinforced relatively to other points in its class, ensuring a balance between point classes. However, manipulating weights increases the complexity of the context and may hinder the analysis of the global refinement. Too large a difference between image and ground weighting can introduce other issues. Therefore, weighting must be handled with caution.

3.4.4 Attitude and sensor refining: Finally, a global refinement of the geometric models is performed through spatio-triangulation. To compensate for localization and orientation errors, the attitude of each acquisition is adjusted by optimizing roll, pitch, and yaw parameters. Constant attitude biases, modeled as zero-order polynomials, are refined according to their respective uncertainties σ_{x0}, σ_{y0} and σ_{z0}. Furthermore, attitude drifts (first-order polynomials) associated with uncertainties σ_{x1}, σ_{y1} and σ_{z1} can also be estimated to account for potential residual temporal instabilities.
Given the dual-satellite stereo configuration, a per-satellite scale refinement (magnification) can also be implemented. This

specific adjustment focuses on correcting a sensor-specific bias in the Z-direction.
And so for each acquisition:

$$\sigma_{modelsAttitude} = ((\sigma_{roll\ x0}, \sigma_{roll\ x1}), (\sigma_{pitch\ y0}, \sigma_{pitch\ y1}), (\sigma_{yaw\ z0}, \sigma_{yaw\ z1}))$$

And for each satellite:

$$\sigma_{modelsSensor} = \sigma_{magnification\ z0}$$

A mathematical adjustment model minimizing planimetric and altimetric residuals across the block - the least squares method - simultaneously optimizes the position and orientation of each image. During this process we also filter out outliers running a median absolute deviation (MAD) at k=3σ. We choose to keep 99.7% of the data points in both image and ground observation contexts.

Combining optimized points accuracy, weighting and models refining is the key to obtain a successful global refining.

4. Study Site and Results

4.1 Choice of dataset and targeted accuracy

4.1.1 Study site: The study site is in Chad, centered over the over Aorounga crater. The dataset consists of a 4x3 grid of stereo pair acquisitions covering approximately 20x21 km², resulting in twenty-four acquisitions (cf. Figure 4).
This site was selected for several reasons. First, its desert and relatively flat terrain ensures that altimetric references remain consistent between the SRPs and the Copernicus reference DSM. This is a key consideration since SRPs are accurate to approximately 3 meters and can inadvertently lock onto building rooftops in urban areas, whereas the Copernicus DSM, accurate to 10 meters, models bare ground more reliably — making a flat, open environment preferable to avoid such discrepancies. Second, the site exhibits sufficient surface texture to enable robust image correlation across acquisitions, yielding high-quality correlation maps. Finally, we possess a 3D reference derived from Pléiades Neo — accurate to better than 1.5 meters in altimetry over this site — allowing us to validate the geometric refinement performance by comparing the CO3D DSM against the Pléiades Neo DSM. The present work focuses on the geometric refinement stage while the full 3D reconstruction pipeline is detailed in a companion paper (Youssefi, D. et al., 2026).

Figure 4: Study site and its extend in Google

4.1.2 Targeted accuracy: The primary benefit of global block adjustment is to ensure geometric consistency across all acquisitions while maintaining high planimetric and altimetric geolocation accuracy. To evaluate this, two distinct metrics are considered:

- Geolocation error of the mosaic acquisitions, defined as the planimetric and altimetric discrepancy between the refined ground measurements and the true ground reference.
- Altimetric shift between adjacent acquisitions, defined as the elevation discontinuity at the junction zone between two neighbouring scenes.

The geolocation error, expressed as Circular Error at 90th percentile (CE90) for planimetry and Linear Error at 90th percentile (LE90) for altimetry, shall not exceed 5 meters when measured against the SRPs (this budget includes the intrinsic CE90 and LE90 uncertainty of the SRPs themselves, estimated at 3 meters).
In this paper, planimetric refinement performance will not be disclosed due to a confidentiality agreement. Instead, we rely on the altimetric performance of the SRPs to assess the quality of the global adjustment. While the altimetric performance of tie points could also be analysed, their accuracy is tied to that of the Copernicus DSM, which is of lower quality than the SRPs and therefore less suitable as a validation reference.

Regarding the altimetric shift, the elevation discontinuity between two adjacent 5×7 km² scenes shall remain below 0.6 m (LE90). Although this specification can only be fully verified by comparing two 3D products, we can attempt to assess whether we approach the requirements at our 2D processing stage and whether the refined geometric models are of sufficient quality to support the subsequent generation of 3D DSMs.
To measure this altimetric bias at a junction zone, for every 4-tuple point — i.e., a point observed simultaneously by four acquisitions — we compute the difference between the pseudo-intersection derived from the two lines of sight of stereo pair A and that derived from the two lines of sight of stereo pair B.
It should also be noted that one production tile encompasses one distribution tile plus a half-scene margin on each side. In our specific case, the distribution tile corresponds to the central 2×1 scene plus half scene of the margin.

4.2 Individual Refining with GCP

Each of the 24 acquisitions in the dataset must be processed individually to achieve precise local alignment
The implemented individual refinement method, together with the parameter settings described below, is expected to comply with the mission specifications described above.
To begin, let’s visualize the SRP data in Google’s georeferenced image dataset (cf. Figure 5).

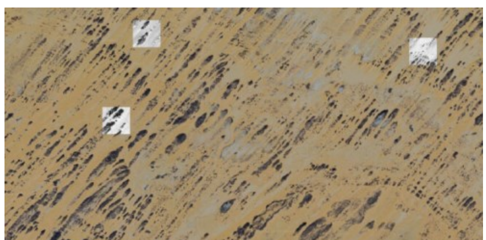

Figure 5: Four SRPs (in grey) footprints in Google

4.2.1 Correlation Parameterization: The steps presented in the methodology use parameter settings that comply with the technical specifications, set at the maximum acceptable values given the computational constraints.
First, the CO3D products are projected into the geometry defined by the SRP, and a filter is applied on CO3D to harmonize the blur.
The correlation step is then performed using the following parameter settings:

Correlation window size	45x45 pixels srp geometry
Search area size	80x80 pixels srp geometry
Correlation score threshold	0,45

Table 1: correlation settings used for SRP correlation

The following parameters (cf. Table 1) are generally applicable to a variety of landscapes (urban, vegetated, desert).
The correlation coefficient may appear low and, consequently, may not allow for the elimination of outliers at this stage. This is intentional, as outlier filtering will be performed later during the refinement of the geometric model.
Each SRP thumbnail is therefore aligned with image CO3D. The figure below (cf. Figure 6) illustrates this visually: at the center is a section of image CO3D, surrounded by the corresponding SRP thumbnail.

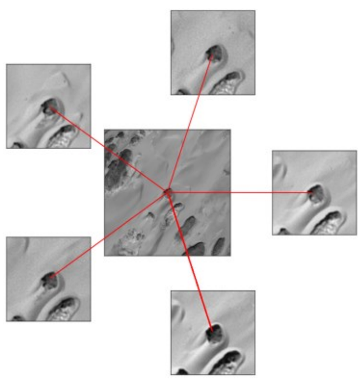

Figure 6: CO3D–SRP Correlation

In this case, it works well; all SRP thumbnails converge to the same point. The median value obtained will be retained.

4.2.2 Unit Refinement Parameterization:
To correct geolocation and orientation errors, the attitude of each acquisition is refined by optimizing the roll, pitch, and yaw angles (cf. Table 2). We assign bias uncertainties to each acquisition’s attitude model, defining

Settings	Roll (in rad)	Pitch (in rad)	Yaw (in rad)
Attitude bias	$\sigma_{SRP\ roll\ x0} = 10^{-3}$	$\sigma_{SRP\ pitch\ y0} = 10^{-3}$	$\sigma_{SRP\ yaw\ z0} = 10^{-2}$

Table 2: Attitude refining accuracies - individual refining

which corresponds to 1000 µrad, 1000 µrad, and 10 000 µrad respectively.
We intentionally set a large bias of 1000 µrad for roll and pitch so that unit-level refinement can operate with maximum flexibility.
For yaw, the bias is set even higher, at 10 000 µrad. Preliminary studies have demonstrated that this choice yields results consistent with the observations, although a residual yaw component can remain after the unit-level refinement based on the SRPs. That residual can subsequently be removed during the global refinement stage.

All bias values are deliberately chosen to exceed the mission-specified limits, thereby providing a generous safety margin for the optimization process while ensuring coherence of the results at every processing step.
During the unit-level adjustment, the points used for the refinement are first passed through a k-MAD (k-median absolute deviation) filter with a threshold of 2.5, to reject outliers and retain only the most reliable observations.

In certain specific cases, we do not have enough SRPs to make the refinement robust (either there are not enough SRPs within the CO3D domain, or the SRPs did not correlate well). To address this, the refined bias can be obtained by refining data with lower resolution, such as Sentinel-2.

Note that, in addition, we correct some residual temporal defects associated with the CO3D matrix sensor through additional refinement operations (cf. Laurent, G. et al., 2026).

4.3 Image correlation between stereo pairs and adjacent scenes

4.3.1 Identifying blocks intersections and distribution tile: Since we refined properly each acquisition, we can now compute our mosaicked ground footprint with its associated models. We also identify the distribution tile (white box in Figure 7) where we must target accuracy requirements defined in 4.1.2).

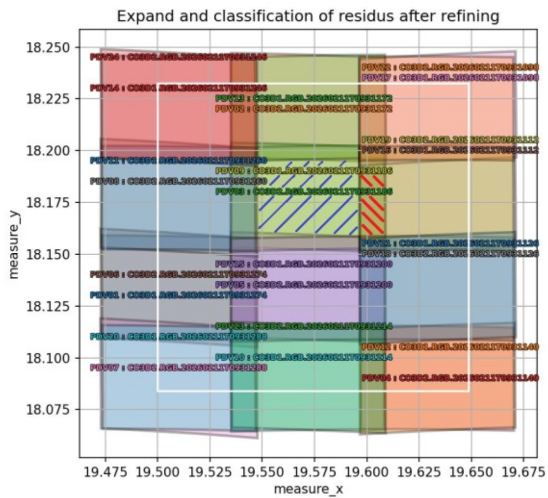

Figure 7: Study site footprints and models

The dataset comprises twelve synchronous stereoscopic acquisition pairs and seventeen intersection areas between stereo pairs. In the following analysis, we focus primarily on the two stereo pairs and the seventeen intersections contained within the distribution tile. The stereo pair hatched in blue in the figure corresponds to the stereo formed between tiles PDV09 (CO3D1) and PDV03 (CO3D2), while the intersection hatched in red corresponds to the overlap between PDV09 (CO3D1) and PDV19 (CO3D2). These will be referred as *stereo_PDV03-09* and *intersection_PDV09-19* respectively and will be examined in detail in the following section (cf. Figure 7).
For each intersection area, the data are preprocessed by resampling all acquisitions into a common geometry prior to the next processing step.

4.3.2 Correlation between acquisitions: In our study, we aim to define correlation parameters that can be uniformly applied across the entire dataset, covering all intersection types — bi-satellite, mono-satellite, stereo, and junction. We therefore adopt a two-pass correlation strategy: a first pass with a large search area and restrictive correlation thresholds, retaining only the most reliable matches; then, following an initial refinement, a second pass with a narrower search area and relaxed thresholds to further populate the correlation map (cf. Table 3).

Correlation settings	1st pass	2nd pass
Correlation window size	51x51 pix²	51x51 pix²
Search area size	20x20 pix²	5x5 pix²
Correlation score threshold	0.9	0.7
Correlation grid spacing	400x400 pix²	400x400 pix²

Table 3: correlation settings for TPs

Correlation grid spacing: A relatively coarse grid spacing is chosen to avoid an excessive density of tie points.
Correlation window size: This represents a good trade-off between capturing fine surface patterns and avoiding noise.
Search area size: Since each acquisition has already been independently refined at the individual stage, large search areas are not required.
Correlation score threshold: The final thresholds remain demanding to minimize erroneous matches.

With these parameters, we obtain for our stereo zone *stereo_PDV03-09* the following results:

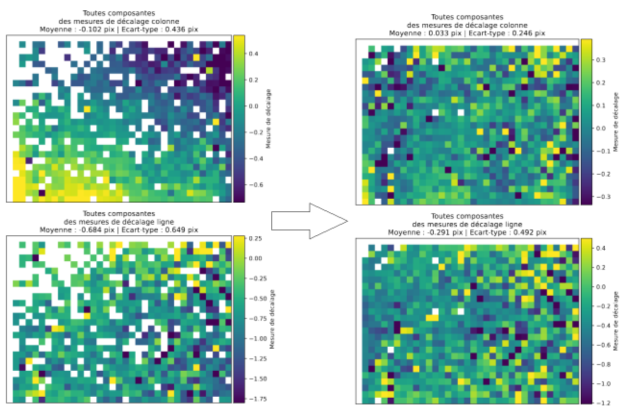

Figure 8: two-pass correlation map for stereo_PDV03-09

And we obtain for the overlapping area *intersection_PDV09-19* the following correlation maps:

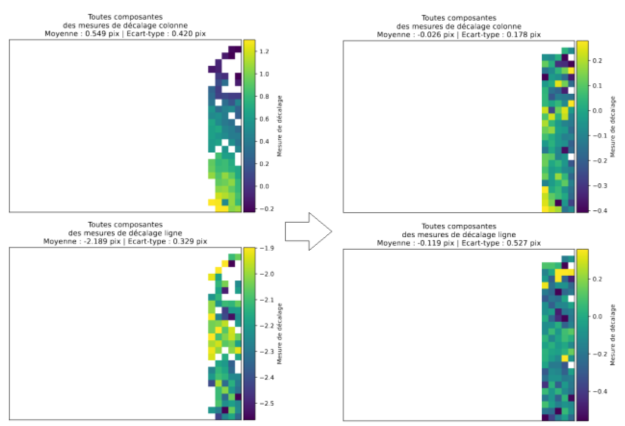

Figure 9: two-pass correlation map for intersection_PDV09-19

4.3.3 Identifying n-tuple point: The acquisitions successfully correlated across nearly all points within the overlapping zones, yielding a dense set of matched points (cf. Figure 8 and Figure 9). We can now present our final observation context and identify the multiplicity of each point.

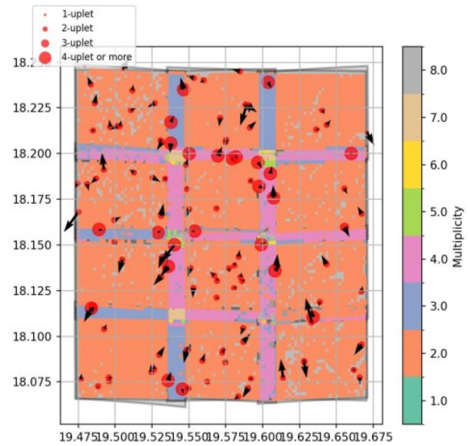

Figure 10: Ground observation context and points multiplicity

Our observation context is particularly rich. We primarily seek to retain 4-tuple points to link the junction zones between adjacent scenes. However we can see some junction zones contain no 4-tuple or higher-tuple point (pink point in Figure 10). After investigation, this was identified as a limitation of our correlation tool. We therefore choose to retain all points of multiplicity 3 and above, and to keep only 5% of multiplicity-2 points, selected through random sampling.

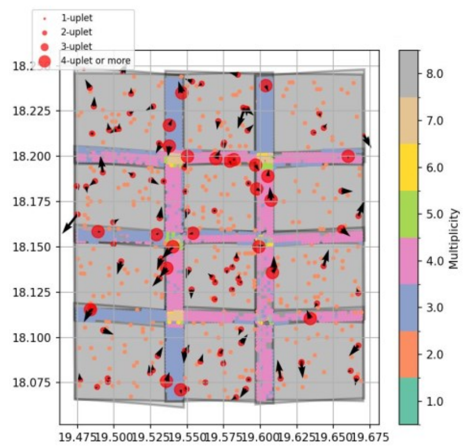

Figure 11: Filtered ground observation context and points multiplicity

Our observation context looks ready, we can start the global refining process.

4.4 Global Refining on case study

4.4.1 Image and ground precision: as explained in the 3.4.2, image uncertainties of our GCPs match the correlation accuracy of our SRPs, and the image uncertainties of our TP matches the correlation accuracy of our TPs. In both cases our correlation accuracy is comparable to the correlation noise, which is around 0.2 pixel.

Ground uncertainties of our GCPs match the ground accuracy of our SRPs. It is defined as 3m CE90/ LE90 but we rely on the

accuracy of our SRPs’ geolocation, especially in reference sites such as Aorounga crater area. So we evaluate SRPs’ planimetric accuracy to 0.5m. For SRPs’ altimetric accuracy we are less confident and we prefer sticking to higher values. We choose 2m accuracy in altimetry. For the ground uncertainties of our TPs we chose the accuracy of our referenced MNS Copernicus, which is defined as 10m accuracy.
And so we define, for each point i (cf. Table 4),

Settings	Image obs. Context	Ground obs. context
SRPs	$\sigma_{imgGCP_i} = (0.2, 0.2)$ pix²	$\sigma_{gndGCP_i} = (0.5, 2)$ m²
TPs	$\sigma_{imgTP_i} = (0.2, 0.2)$ pix².	$\sigma_{gndTP_i} = 10$ m

Table 4: Image and ground uncertainties

4.4.2 Image and ground weighing: Our observation context appears well-balanced between SRPs and tie points; therefore, no specific weighting will be applied to either the image measurements or the ground measurements in this study case.
For each point i we leave $\omega_{imgGCP_i} = \omega_{imgTP_i} = \omega_{gndGCP_i} = \omega_{gndTP_i} = 1$.

4.4.3 Attitude and sensor refining: To compensate for geolocation and orientation errors, the attitude of each acquisition is adjusted by optimizing roll, pitch, and yaw parameters. We set uncertainties for bias and a drift in attitude, and so for each acquisition attitude model we define (cf. Table 5),

Settings	Roll (in rad)	Pitch (in rad)	Yaw (in rad)
Atttitude bias	$\sigma_{roll\ x0} = 10^{-4}$	$\sigma_{pitch\ y0} = 10^{-4}$	$\sigma_{yaw\ z0} = 10^{-2}$
Attitude drift	$\sigma_{roll\ x1} = 10^{-5}$	$\sigma_{pitch\ y1} = 10^{-5}$	$\sigma_{yaw\ z1} = 10^{-10}$ (frozen)

Table 5: Attitude refining accuracies – global refining

A large attitude bias of 100 µrad is set in roll and pitch, allowing the global refinement sufficient freedom to correct any residual errors, even though a localization error of that magnitude is highly unlikely at this stage. The yaw bias is set to an even larger value of 10^4 µrad, as individual refinement using SRPs has demonstrated that, in some cases, a residual yaw bias may persist.
An attitude drift of 10 µrad/s is allowed in roll and pitch to account for potential residual drift. Yaw drift is frozen at 10^{-4} µrad/s to prevent the introduction of additional errors that have been observed when this parameter is left unconstrained.

To correct for potential scale errors related to satellite-induced magnification — a recurring issue in bi-satellite stereo mosaics — the refinement is extended to include a sensor Z-bias, which physically models a magnification correction. By sharing this state across all models from the same satellite, we ensure that the magnification parameter is refined on a per-satellite basis (cf. Table 6),

Settings	X (no unit)	Y (no unit)	Z (no unit)
Sensor bias	10^{-10}	10^{-10}	$\sigma_{modelsSensor} = 10^{-5}$

Table 6: Sensor refining accuracies

Following analysis of the magnification effect on CO3D, an uncertainty of 10^{-5} corrects pixel-level deviations on the order of 0.1 pixel at the edge of the field of view.

Once the parameterization is complete, the spatio-triangulation is run to globally refine the observation context.

4.4.4 Analysis of the results: The results can be examined from two complementary perspectives: the analysis of altimetric residuals from the SRPs (and TPs despite their lower accuracy), and the analysis of the altimetric shift at junction zones.

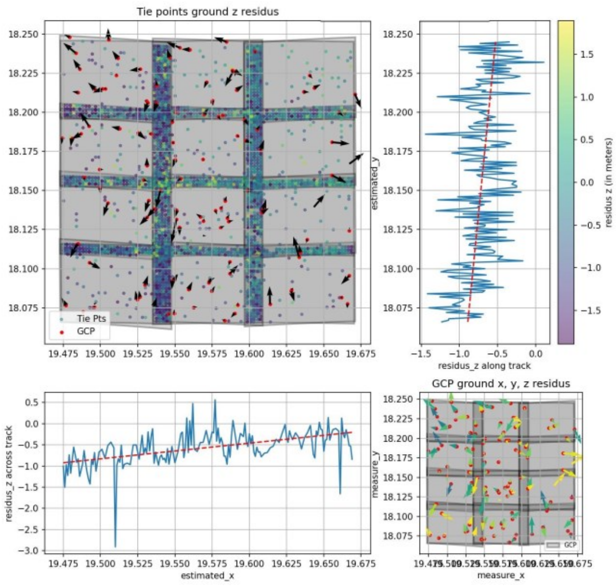

Figure 12: Altimetric residues of ground observation context

The Figure 12 above consists of four plots. The top-left panel displays the altimetric residuals of the tie points, color-coded according to the colorbar on the right. Red dots represent SRPs retained (active) after k-MAD filtering, and the arrows indicate planimetric residuals. The bottom-left and top-right panels show the mean altimetric residual profiles in the across-track and along-track directions, respectively. The red dashed lines provide an indication of the linear trend in these residuals. Finally, the bottom-right panel displays the altimetric residuals of the SRPs, also color-coded according to the colorbar.

The altimetric residual analysis of the SRPs yields a mean of 0.88 m, a standard deviation of 0.66 m, and a maximum absolute value of 1.91 m. For the tie points, the corresponding values are a mean of −0.65 m, a standard deviation of 0.86 m, and a maximum absolute value of 3.19 m.
Several observations can be drawn from these results:

- The altimetric residuals appear well-refined, showing no spatial pattern and appearing to be randomly distributed.
- The altimetric profiles of the tie points exhibit very slight trends in both the across-track and along-track directions, while the SRP altimetric profiles remain close to zero with no discernible trend.

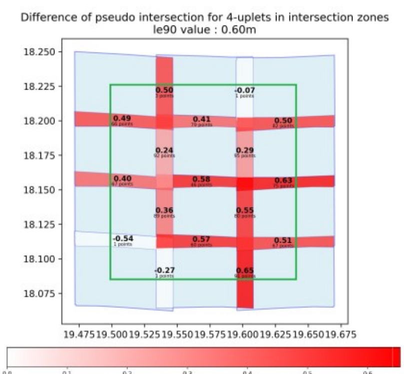

Figure 13: altimetric shift between two stereos at junction zones

Regarding the altimetric shift at each intersection, the Figure 13 above presents the LE90 shift value of every point at each junction zone along with the number of points used in its computation. In some areas, the pseudo-intersection difference was computed from very few or even a single point, corresponding to zones where no 4-tuple points were found. In such cases, the pseudo-intersection was computed from only the 4-tuple SRPs present on this junction (cf. Figure 11).
Within the distribution tile (green square in Figure 13), the LE90 altimetric shift of these LE90 intersection values shifts results in 0.60 m. Excluding junction zones computed from a single point, the highest measured altimetric shift is 0.65 m.

5. Conclusion and future research work

This paper presented a global block adjustment workflow for mosaicked stereoscopic satellite imagery from the CO3D constellation. Through a three-stage pipeline combining individual SRP-based refinement, two-pass tie point extraction, and weighted spatio-triangulation, the method achieves strong geometric consistency across a 4×3 stereo acquisition dataset over the Aorounga crater, Chad.
The results confirm that the global refinement meets the targeted accuracy requirements: the altimetric shift between adjacent scenes within the distribution tile reaches an LE90 of 0.6 m against a 0.6 m threshold, and the SRP altimetric residuals show no systematic spatial trend after adjustment.
Several methodological choices were found to be critical to the success of the pipeline. The two-pass correlation strategy effectively balances tie point density and quality across heterogeneous intersection types. The n-tuple point filtering ensures a geometrically well-conditioned observation framework without over-constraining the adjustment. Finally, the introduction of a per-satellite magnification parameter proved particularly valuable in correcting scale-related errors inherent to bi-satellite stereo configurations.

Future work will focus on improving key stages of the pipeline such as extracting more 4-tuple points to strengthen the observation framework, extending this workflow to larger mosaics, assessing performance under varying landscape conditions, and integrating the geometric refinement pipeline with the full 3D reconstruction chain described in the companion paper (Youssefi, D. et al., 2026).

References

Buffe, F., Lucas, S., Latourte, A., Delaygue, N., Accadia, T., Youssefi, D., Dick A., Luitot, C., Bruno, M., Coggiola, C., Tanguy, Y., Lebègue, L., Thomas, C., Schillinger, M., Chambrelan, A.: CO3D IMAGE QUALITY CALIBRATION, ISPRS 2026 (Submitted)

Chandelier, L., Coeurdevey, L., Bosch, S., Favé, P., Gachet, R., Orsoni, A., Tilak, T., and Barot, A.: A WORLDWIDE 3D GCP DATABASE INHERITED FROM 20 YEARS OF MASSIVE MULTI-SATELLITE OBSERVATIONS, ISPRS Ann. Photogramm. Remote Sens. Spatial Inf. Sci., V-2-2020, 15–23, https://doi.org/10.5194/isprs-annals-V-2-2020-15-2020, 2020

Cournet, M., Giros, A., Dumas, L., Delvit, J. M., Greslou, D., Languille, F., Blanchet, G., May, S., and Michel, J.: 2DSub-Pixel Disparity Measurement Using QPEC / Medicis,Int. Arch. Photogramm. Remote Sens. Spatial Inf. Sci., XLI-B1,291-298, http://dx.doi.org/10.5194/isprs-archives-XLI-B1-291-2016 .

Laurent, G., Latourte, A., Buffe, F., Vidal, B.: LINE OF SIGHT CALIBRATION FOR SATELLITE IMAGERY BASED ON MATRIX DETECTOR, ISPRS 2026 (Submitted)

Lebègue, L., Buffe, F., Melet, O., Youssefi, D., Henry, J.: The CO3D mission, a worldwide one-meter accuracy Digital Surface Model, ISPRS 2026 (Submitted)

Youssefi, D., Accadia, T., Colleu, E., Godet, A., Bellet, V., Lallement, D., Tanguy, Y., Lebègue, L., 2026: 3D Product Quality Control in the CO3D Mission: A Critical Role, ISPRS 2026 (Submitted)