Comprehensive Evaluation of Small-Format Multi-Head Camera Systems for 3D Topographic Mapping

Thirawat Bannakulpiphat[1,2], Wilfried Karel[1], Ana-Maria Loghin[3,4], Valeria-Ersilia Oniga[3], Norbert Pfeifer[1]

[1] Department of Geodesy and Geoinformation, Technische Universität Wien, Vienna, Austria
(thirawat.bannakulpiphat, wilfried.karel, norbert.pfeifer)@geo.tuwien.ac.at
[2] Mapping and Positioning from Space (MAPS) Technology Research Center, Department of Survey Engineering, Chulalongkorn University, Bangkok, Thailand
[3] Department of Terrestrial Measurements and Cadastre, Technical University „Gheorghe Asachi" of Iași, Iași, Romania
(ana-maria.loghin, valeria-ersilia.oniga)@academic.tuiasi.ro
[4] Federal Office of Metrology and Surveying (BEV), Vienna, Austria

Keywords: Small-Format, Multi-Head Camera Systems, Oblique Camera, UAV, Topographic Mapping, Façade Reconstruction

Abstract

Small-format multi-head camera systems are increasingly deployed for UAV-based 3D mapping; however, systematic evaluations that explicitly link acquisition geometry to network robustness and façade reconstruction completeness remain scarce. This study establishes a geometry-driven, multi-dataset evaluation framework to assess such systems in topographic mapping. Six UAV datasets with varying flying altitudes, overlaps, platforms, cameras, and urbanization are analyzed using a rigorous bundle block adjustment workflow. Internal network strength is quantified using image connectivity, image ray intersection angles, and tie point multiplicity, while object-space accuracy is assessed using ground control point residuals. Reconstruction completeness is evaluated using façade-specific point density metrics, enabling a controlled comparison between nadir-only and multi-head (nadir + oblique) configurations. The results show that oblique imagery strengthens multi-directional image network geometry and substantially improves façade reconstruction. This study establishes empirical relationships between acquisition design, network stability, and object-space completeness.

1. Introduction

Small-format multi-head camera systems represent a recent advancement in photogrammetric sensing. Their compact and lightweight design makes them particularly suitable for unmanned aerial vehicles (UAVs), where payload constraints directly influence operational efficiency. By combining UAV mobility with oblique imaging capabilities, these systems enable high-quality data acquisition in complex and difficult-to-access environments. The simultaneous capture of multiple viewing directions increases spatial coverage and geometric redundancy, thereby supporting the generation of high-resolution cadastral products, 3D building models, and comprehensive urban documentation (Remondino and Gerke, 2015, Gerke et al., 2016, Toschi et al., 2017).

Recent improvements in UAV platforms and small-format sensors have further enhanced the capabilities of aerial photogrammetry. Compared with traditional single-camera configurations, multi-head systems employ synchronized cameras oriented in complementary directions, thereby strengthening image overlap and the geometric network structure. These advantages are particularly relevant in urban environments, where vertical façades and occlusions limit the effectiveness of nadir-only acquisition. However, the integration of multiple sensors and oblique viewing geometries increases the complexity of image orientation, tie point extraction, and bundle block adjustment.

Despite their growing use in operational mapping, most existing studies investigate multi-head UAV systems through isolated case studies or focus primarily on adjustment formulation. A systematic, geometry-driven evaluation that quantitatively links acquisition parameters, internal network strength indicators, bundle adjustment behavior, and façade reconstruction completeness across multiple datasets remains largely lacking. In particular, the interaction between directional connectivity, ray intersection geometry, and dense reconstruction performance under varying flight strategies has not been sufficiently analyzed for small-format UAV-based multi-head systems. To address this gap, this study integrates network geometry diagnostics, statistical adjustment analysis, and façade-specific completeness metrics within a multi-dataset framework.

1.1 Research Aim

The primary objective of this study is to investigate how acquisition geometry and computational processing influence the accuracy and completeness of UAV-based 3D mapping using small-format multi-head camera systems. A geometry-driven evaluation framework is applied across multiple datasets to (i) quantify internal network strength using directional connectivity, ray intersection geometry, and tie point multiplicity; (ii) relate these indicators to bundle adjustment precision; and (iii) assess façade reconstruction completeness. The framework establishes explicit links between acquisition design parameters, internal network strength, adjustment stability, and object-space completeness within a unified comparative setting, thereby enabling geometry-aware performance diagnosis of multi-head UAV systems.

2. Related Work

2.1 Modern Oblique SfM

Structure-from-Motion (SfM) constitutes the standard framework for orienting oblique aerial image blocks. The integration of nadir and oblique views enhances façade visibility and reconstruction completeness but also increases geometric complexity

This contribution has been peer-reviewed.
https://doi.org/10.5194/isprs-archives-XLIX-B1-2026-239-2026

in feature matching, image connectivity, and network adjustment. For large-format airborne systems, multi-head configurations demonstrate improved vertical accuracy and façade modeling compared with nadir-only blocks, provided that sufficient cross-directional tie points establish strong inter-camera connections (Nex et al., 2015, Rupnik et al., 2015). These studies highlight that block stability depends not only on image overlap but also on the spatial distribution of viewing directions.

With the transition to UAV platforms, SfM processing must cope with higher image redundancy, shorter baselines, and reduced direct georeferencing accuracy. Topology-aware matching strategies aim to reduce computational costs in high-overlap datasets (Jiang and Jiang, 2017); however, these approaches primarily address efficiency rather than analyzing the intrinsic geometric structure of the image network. Recent investigations of small-format penta-camera UAV systems show that overlap of nadir images does not necessarily ensure symmetric overlap in oblique directions, thereby directly influencing tie point multiplicity and ray intersection geometry (Bannakulpiphat et al., 2024). Despite these advances, systematic geometry-driven evaluations of high-overlap small-format multi-head UAV blocks remain limited, particularly with respect to the interaction between image connectivity, ray intersection angles, and adjustment robustness.

2.2 Multi-Head Camera Bundle Adjustment

Bundle block adjustment (BBA) for multi-head systems requires explicit modeling of the geometric relationships between individual camera heads. Rig-constrained formulations preserve system rigidity by parameterizing oblique cameras relative to a reference camera, thereby reducing parameter redundancy and improving adjustment stability (Rupnik et al., 2015). Reduced normal equation strategies further enhance scalability for large image blocks (Sun et al., 2016). Recent studies indicate that treating camera heads as independent sensors increases the effective degrees of freedom and may lead to noise fitting, thereby degrading object-space accuracy (Maset et al., 2024). Two principal strategies are commonly applied in multi-camera orientation: (i) constraining cameras via externally calibrated relative poses (Schneider and Förstner, 2013), and (ii) estimating relative orientations directly from image correspondences within a Structure-from-Motion (SfM) framework, such as *OrientAL*, *MicMac*, and *COLMAP* (Karel et al., 2013, Pierrot-Deseilligny et al., 2014, Schonberger and Frahm, 2016).

For UAV-based small-format systems, rigorous relative orientation modeling is particularly critical due to compact rig geometry and limited external positioning accuracy. Under such conditions, internal network characteristics—tie point multiplicity, image connectivity, and ray intersection geometry—become decisive for BBA precision and reliability. Although constrained BBA formulations are mathematically well established, comparatively few studies systematically examine how acquisition design parameters, such as overlap configuration and viewing direction distribution, influence geometric stability in high-overlap small-format UAV blocks. This gap motivates the geometry-driven, multi-dataset evaluation presented in this study.

3. Materials and Methods

3.1 Datasets

Six UAV datasets acquired in Thailand (THA) and Romania (ROU) were used to evaluate the geometric accuracy and reconstruction completeness of small-format multi-head camera systems for 3D mapping. Two camera systems were employed: the JOUAV CA502R and the SHARE PSDK 102S V3. Both systems comprise five synchronized cameras—one nadir (Na) and four oblique (Ob, 45°)—mounted on either vertical take-off and landing (VTOL) or multi-rotor UAV platforms. Each camera captures 24 MP images (6000 × 4000 or 6144 × 4096 pixels), with focal lengths ranging from 25 to 28 mm for the nadir camera and from 35 to 43 mm for the oblique cameras. All camera heads operate with an electronic rolling shutter. Among the six datasets, two pairs cover similar geographic areas. One pair was acquired using different camera systems, enabling a direct comparison of system performance under comparable conditions. Although these pairs share similar naming conventions, they may represent adjacent rather than identical areas. Table 1 summarizes the main acquisition parameters, direct georeferencing (DG) accuracy, and ground control point (GCP) statistics. Figure 1 illustrates the spatial distribution of image centers and GCPs for all datasets.

3.2 Data Processing

All datasets were processed using the *OrientAL* software developed at TU Wien. The software implements a complete Structure-from-Motion (SfM) workflow combined with rigorous bundle block adjustment (BBA). Feature extraction and matching are performed for all overlapping image pairs, followed by essential matrix estimation using RANSAC (Fischler and Bolles, 1981) for gross error suppression. Robust parameter estimation within the bundle adjustment is carried out using an M-estimator to mitigate the influence of remaining outliers. To ensure a homogeneous spatial distribution of tie points while favoring high-multiplicity tracks, the feature reduction strategy proposed by (Karel et al., 2016) was applied. This approach reduces redundant correspondences in high-overlap blocks while preserving network stability and orientation accuracy.

OrientAL performs incremental BBA to iteratively refine exterior orientation and camera calibration parameters. It supports the modeling of multi-head systems and integrates direct georeferencing (DG) observations together with ground control point (GCP) observations in object space. All observations are weighted according to their respective stochastic models, with DG and GCP measurements treated as independent object-space observations.

The oriented image blocks were subsequently processed using *PMVS* (Furukawa and Ponce, 2010) to generate dense point clouds for further analysis. As a patch-based multi-view stereo method, *PMVS* reconstructs 3D points only in areas where sufficient texture and favorable viewing geometry provide reliable support, without performing surface interpolation. This conservative behavior makes *PMVS* particularly suitable for evaluating the contribution of oblique imagery, as variations in reconstruction completeness directly reflect the underlying image geometry rather than artifacts introduced by regularization.

In addition, selected datasets were processed using a modern dense image matching approach to enable a comparative assessment of both reconstruction strategies.

3.3 Geometry and Network Strength

The design of acquisition parameters for small-format multi-head UAV systems requires a systematic understanding of how flight altitude, image overlap, and platform configuration influence geometric network robustness. The geometric characteristics of the image block are quantified using three complementary indicators: (i) the image connectivity matrix, (ii) image ray intersection angles, and (iii) tie point multiplicity combined with image residual statistics.

Together, these indicators characterize the internal reliability and spatial strength of the photogrammetric network. They directly

Dataset	Location	Platform + Camera	# Images (Na + Ob)	Nadir GSD	Nadir Overlap (along/across)	Altitude	DG Accuracy (Horiz/Vert)	# GCPs	GCP Accuracy (Horiz/Vert)
1	Bangkok, THA	VTOL + JOUAV CA502R	16,595	5.7 cm	80/80 %	400 m	5.00 / 10.00 m	48	0.02 / 0.05 m
2	Nonthaburi, THA (N-S)	VTOL + JOUAV CA502R	5,660	3.6 cm	80/80 %	250 m	0.02 / 0.05 m	16	0.02 / 0.05 m
3	Nonthaburi, THA (E-W)	VTOL + JOUAV CA502R	5,635	3.6 cm	80/80 %	250 m	0.02 / 0.05 m	16	0.02 / 0.05 m
4	Saraburi, THA	VTOL + JOUAV CA502R	6,065	2.6 cm	80/80 %	150 m	5.00 / 10.00 m	43	0.02 / 0.05 m
5	Saraburi, THA	Multi-rotor + SHARE PSDK 102S V3	5,080	2.9 cm	80/80 %	180 m	5.00 / 10.00 m	39	0.02 / 0.05 m
6	Iași, ROU	Multi-rotor + SHARE PSDK 102S V3	5,150	1.3 cm	80/70 %	60 m	0.03 / 0.20 m	28	0.02 / 0.03 m

Table 1. Summary of acquisition and georeferencing parameters for the six UAV-based multi-head camera datasets.

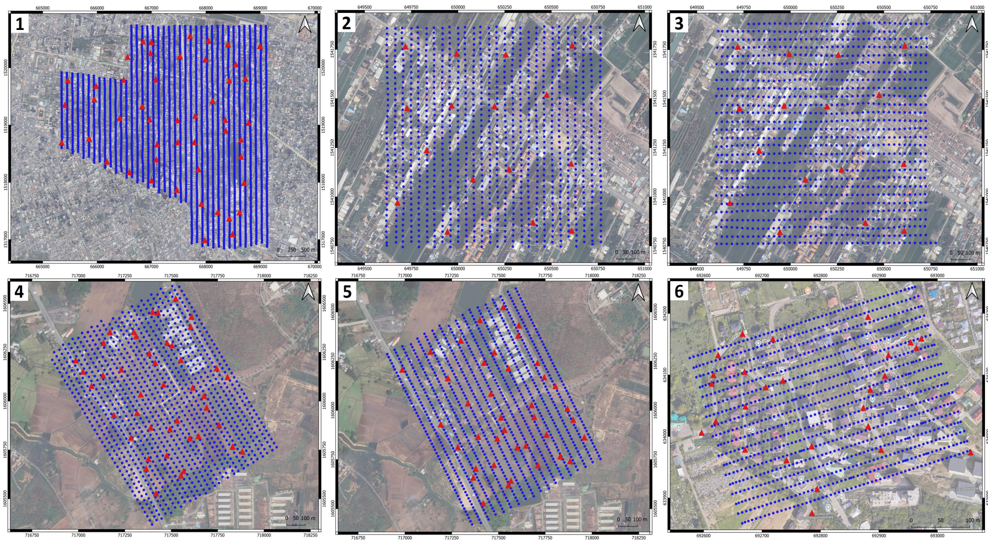

Figure 1. Overview of the six UAV-based multi-head camera datasets. Each subfigure illustrates the spatial distribution of image centers (blue) and ground control points (red), highlighting variations in flight geometry and survey extent.

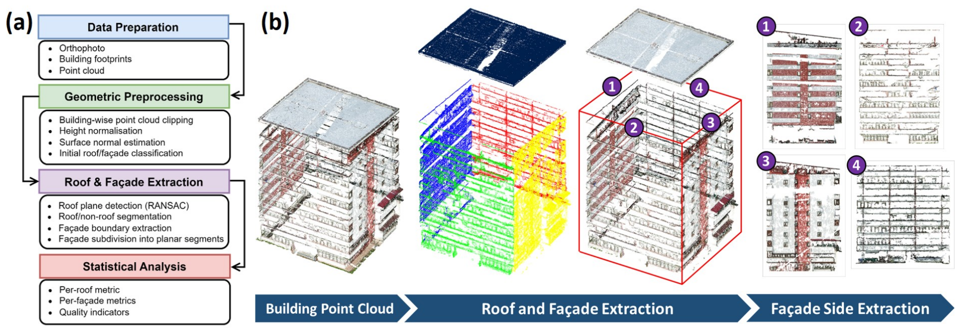

Figure 2. Façade completeness evaluation workflow. (a) Pipeline for extracting façade sides from dense point clouds. (b) Example building showing successive processing stages, from clipped point cloud to roof–façade separation and final façade-side segmentation for density analysis.

influence the stability of the bundle block adjustment (BBA), the precision of the estimated parameters, and ultimately the geometric quality of the reconstructed 3D model.

The methodology for deriving and interpreting these indicators was previously established (Bannakulpiphat et al., 2025) using Dataset 4 (Saraburi), where the effects of viewing geometry and flight configuration were analyzed in detail. In this study, the validated framework is extended to all six datasets, enabling a comparative evaluation across different acquisition geometries, camera systems, and flight strategies. This multi-dataset approach enables a geometry-driven assessment of network robustness beyond isolated case studies.

3.4 Façade Completeness

Façade completeness was quantified using façade-specific point cloud density metrics. For each building, the number of reconstructed 3D points per façade side was evaluated to enable a direct comparison between nadir-only (1CAM) and multi-head (5CAM) reconstructions. Both dense point clouds were derived from identical orientation results in order to isolate the influence of acquisition geometry and camera configuration.

As illustrated in Figure 2, dense point clouds were first clipped per building using manually digitized footprints derived from the orthophoto. The same spatial footprints across the study area were applied to all processing configurations to ensure geometric consistency and to prevent neighboring structures from influencing façade statistics. Geometric preprocessing was then performed. Surface normals were estimated using local plane fitting, and height normalization was applied based on the minimum and maximum elevations of each building. Roof and façade points were separated using a rule-based classification that combined (i) the vertical component of the surface normal and (ii) the normalized height attribute. Roof points were characterized by near-horizontal normals at higher elevations, whereas

Dataset	Camera configuration	A posteriori variance factor (σ_0) [px]	GCP residuals																	
			BBA with GCP observations									BBA with DG + GCP observations								
			Mean residual (μ) [m]			Standard deviation (σ) [m]			σ / Nadir GSD			Mean residual (μ) [m]			Standard deviation (σ) [m]			σ / Nadir GSD		
			X	Y	Z	X	Y	Z	X	Y	Z	X	Y	Z	X	Y	Z	X	Y	Z
1	1CAM	0.805	0.003	0.003	0.003	0.062	0.080	0.097	1.1	1.4	1.7	0.003	0.003	0.002	0.062	0.080	0.097	1.1	1.4	1.7
	5CAM	0.933	-0.002	0.016	0.026	0.048	0.053	0.133	0.8	0.9	2.3	0.000	0.018	0.010	0.048	0.051	0.129	0.8	0.9	2.3
2	1CAM	0.889	0.001	-0.007	-0.015	0.035	0.053	0.068	1.0	1.5	1.9	0.018	0.007	0.040	0.029	0.028	0.049	0.8	0.8	1.4
	5CAM	0.970	0.002	0.001	0.017	0.053	0.039	0.044	1.5	1.1	1.2	0.019	-0.028	0.029	0.016	0.011	0.021	0.5	0.3	0.6
3	1CAM	0.949	-0.008	0.011	-0.002	0.058	0.056	0.088	1.6	1.6	2.5	0.056	0.002	0.020	0.062	0.040	0.054	1.7	1.1	1.5
	5CAM	1.014	-0.010	0.007	-0.012	0.023	0.041	0.081	0.6	1.2	2.3	0.078	0.012	-0.038	0.011	0.024	0.030	0.3	0.7	0.8
4	1CAM	0.973	-0.008	0.000	-0.004	0.028	0.036	0.049	1.1	1.4	1.9	-0.008	0.000	-0.009	0.028	0.035	0.049	1.1	1.4	1.9
	5CAM	1.043	-0.004	-0.006	0.011	0.023	0.023	0.102	0.9	0.9	4.0	-0.003	-0.002	-0.014	0.024	0.021	0.099	0.9	0.8	3.9
5	1CAM	0.678	0.008	0.004	0.000	0.022	0.016	0.031	0.8	0.6	1.1	0.008	0.004	0.000	0.022	0.016	0.029	0.8	0.6	1.0
	5CAM	0.932	0.003	-0.003	0.041	0.073	0.054	0.310	2.5	1.9	10.8	0.003	-0.002	0.037	0.072	0.054	0.306	2.5	1.9	10.6
6	1CAM	0.773	-0.006	-0.016	-0.006	0.053	0.075	0.043	4.1	5.9	3.4	-0.016	-0.027	0.083	0.020	0.039	0.115	1.6	3.0	9.0
	5CAM	1.002	-0.002	-0.014	0.036	0.035	0.063	0.089	2.7	4.9	7.0	-0.032	-0.030	-0.105	0.057	0.044	0.160	4.5	3.4	12.5

Table 2. Bundle block adjustment (BBA) statistics for single-camera (1CAM) and multi-head (5CAM) configurations under GCP-constrained and DG-supported adjustment scenarios. The a posteriori variance factor reported in the table refers to the value obtained after relative orientation and is therefore independent of whether GCP constraints were applied.

façade points exhibited near-vertical normals within a defined height range. All preprocessing parameters were kept identical across configurations to ensure a fair comparison.

The extracted façade points were projected onto the horizontal plane to derive a building footprint using a convex hull approach. For improved stability and comparability, the footprint was optionally regularized using a rotated bounding box aligned with the dominant façade orientation. Façade sides were then defined using a centroid-based angular sector method to ensure consistent subdivision of the building envelope. For each façade side (f), the planar width (w_f) was derived from the corresponding footprint edge, and the vertical extent (h_f) from the Z-range of the assigned points. The façade reconstruction density is computed as

$$\rho_f = \frac{N_f}{w_f \cdot h_f}, \tag{1}$$

where N_f denotes the number of assigned points. The metric ρ_f (points/m^2) provides a geometry-controlled indicator of façade reconstruction completeness and enables a consistent comparison across acquisition geometries and camera configurations. Since *PMVS* performs conservative patch-based reconstruction without surface interpolation, the resulting density primarily reflects geometric support and multi-view consistency.

4. Results and Discussion

4.1 Accuracy

The accuracy analysis links bundle block adjustment (BBA) performance to object-space georeferencing quality. Internal precision is represented by the a posteriori variance factor (σ_0), while accuracy is assessed using the mean (μ) and standard deviation (σ) of ground control point (GCP) residuals in object space. The distortion modeling strategy validated for Dataset 4, including affine and tangential parameters, was consistently applied to all datasets to ensure comparability across configurations and to isolate the influence of image geometry. Given the use of electronic rolling shutter sensors, rolling shutter effects were compensated within the bundle adjustment to prevent systematic temporal distortions from affecting the accuracy assessment.

Two camera configurations were evaluated: (i) single-camera blocks using nadir imagery only (1CAM) and (ii) multi-head blocks including nadir and oblique images (5CAM). For each configuration, two adjustment scenarios were analyzed: (i) BBA with GCP observations only and (ii) BBA with combined DG and GCP observations. The complete statistics are summarized in Table 2.

The inclusion of oblique imagery is expected to strengthen the geometric network by increasing ray intersection diversity and tie point redundancy, thereby potentially improving adjustment stability and vertical precision. Likewise, incorporating DG observations as weighted constraints should enhance global block stability. However, the results reveal notable dataset-dependent variations in adjustment performance and object-space accuracy.

To obtain an overall view of the trends, the median values of the GSD-normalized standard deviation of object point coordinates ($\sigma/Nadir\ GSD$) were analyzed. Compared with the nadir-only configuration, the combined nadir–oblique system reduces the GSD-normalized planimetric standard deviation (X and Y) from a median of 1.2 to 0.9. This improvement can be attributed to increased redundancy and enhanced network strength. Furthermore, incorporating direct georeferencing observations reduces the planimetric GSD-normalized standard deviation from 1.3 to 1.0, likely due to improved gross error detection and stabilization of the block geometry. In the vertical component (Z), the results are more heterogeneous. The inclusion of oblique views leads to an increase in the GSD-normalized standard deviation from 1.8 to 2.3, which is not expected. In contrast, supplementing the GCP configuration with direct georeferencing reduces the vertical GSD-normalized standard deviation from 2.1 to 1.8.

For Datasets 2 and 3, the 5CAM configuration reduces the GSD-normalized standard deviations compared with 1CAM, particularly in the vertical component. This indicates that additional oblique viewing directions improve vertical ray intersection geometry under symmetric acquisition conditions. However, Datasets 2 and 3 were acquired at the same location and contain the smallest number of GCPs.

In contrast, Dataset 5 exhibits increased residual dispersion in the 5CAM configuration, especially in height. This suggests that geometric strengthening through oblique imagery depends strongly on acquisition symmetry, overlap consistency, and the quality of direct georeferencing (DG), rather than being inherently guaranteed. The degradation observed in Dataset 5 demon-

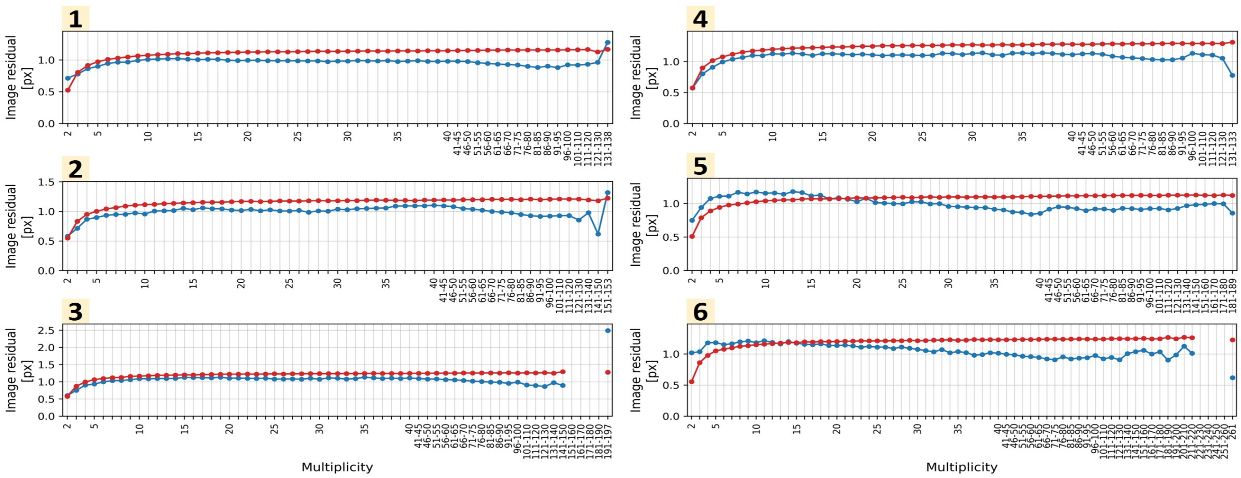

Figure 3. Mean image residual length as a function of tie point multiplicity. Blue lines represent empirical residual statistics derived from the bundle adjustment, while red lines indicate simulated residuals resulting from synthetic noise applied to the tie point projections used as image observations. The noise was sampled from a normal distribution with standard deviation equal to the estimated σ_0.

strates that adding camera heads does not automatically improve accuracy. Increased parameter dimensionality may amplify instabilities when geometric symmetry or external constraints are suboptimal. The geometric benefit of oblique imagery is therefore conditional rather than universal. Systematic differences at the GCPs are generally within the centimeter domain or better, particularly when direct georeferencing is not incorporated. Only in Datasets 1, 4, and 5 are these values maintained after incorporating DG; however, in these cases, the reported DG accuracy introduced into the adjustment is relatively low. This indicates unresolved challenges in reliably realizing the geodetic datum at centimeter-level accuracy.

Extended camera modeling effectively reduces systematic effects, with mean image residual magnitudes generally remaining below ±0.1 pixels. Figure 3 presents the mean image residual length as a function of tie point multiplicity. The blue curves represent empirical residual statistics, while the red curves correspond to simulated residuals generated by perturbing the image observations with normally distributed noise, with a standard deviation equal to the estimated σ_0 value, following the simulation approach described in (Bannakulpiphat et al., 2025).

Across all datasets, the empirical residual distributions deviate systematically from the simulated normal behavior. Whereas the simulated curves follow smooth theoretical trends, the observed residuals exhibit multiplicity-dependent structure and non-random deviations. This indicates that the image observations do not fully satisfy the classical least-squares assumptions of independence, homoscedasticity, and a correctly specified stochastic model. These findings suggest that stochastic and functional modeling strategies for high-overlap multi-head UAV blocks may require refinement. In particular, multiplicity-dependent weighting schemes or direction-dependent variance modeling could better represent the structured residual behavior observed in the empirical data.

Overall, the results demonstrate that the geometric advantage of multi-head imagery is scenario-dependent, with consistent vertical improvements observed only under geometrically balanced overlap configurations across viewing directions. Agreement between simulated and empirical results is substantially better for Datasets 1–4, whereas Datasets 5 and 6 exhibit similar systematic deviations. Two possible explanations are that (i) the geometric and stochastic modeling of the SHARE PSDK 102S V3 multi-head system may not yet be fully represented within the current bundle block adjustment formulation, particularly given the strong contribution of oblique views, and (ii) the multi-rotor acquisition platform may introduce additional dynamic effects that are not fully compensated in the current adjustment model.

4.2 Image Connectivity

Image connectivity was analyzed using the compass-direction framework introduced in (Bannakulpiphat et al., 2024). In multi-head systems, the effective viewing direction of oblique cameras changes with flight direction. For each flight strip, two pairs of opposite oblique cameras switch their viewing orientation, while the compass reference frame remains fixed. Aligning the connectivity analysis with this fixed compass frame enables consistent interpretation across different flight configurations.

Given the large number of images per dataset, connectivity was evaluated at the block level using a super-image connectivity approach. Each image block was subdivided into a 5×5 grid, and connectivity statistics were aggregated within each cell prior to global evaluation, following (Bannakulpiphat et al., 2025). This approach provides a scalable and robust characterization of directional connectivity for high-overlap UAV blocks.

Across all datasets, a consistent ranking of connectivity strength was observed: same-direction image pairs exhibit the highest connectivity, followed by nadir–compass-direction pairs, 90°-difference pairs, and finally 180°-difference pairs. This ordering agrees with the findings for Dataset 4 in (Bannakulpiphat et al., 2025) and confirms that viewing direction governs tie point establishment. Connectivity variations are influenced by flight geometry, overlap configuration, and side-lap percentage; however, the results demonstrate that connectivity does not scale linearly with nominal overlap. Instead, it is primarily controlled by viewing-direction diversity and scene structure. In particular, establishing correspondences between images acquired from opposing directions (180° difference) remains challenging even under high-overlap conditions.

The applied feature selection and thinning strategy contribute to a more homogeneous spatial distribution of tie points, reducing local oversampling while preserving directional connections. This homogenization improves the interpretability of connectivity statistics and strengthens geometric robustness in subsequent bundle adjustment. Overall, these findings demonstrate that nom-

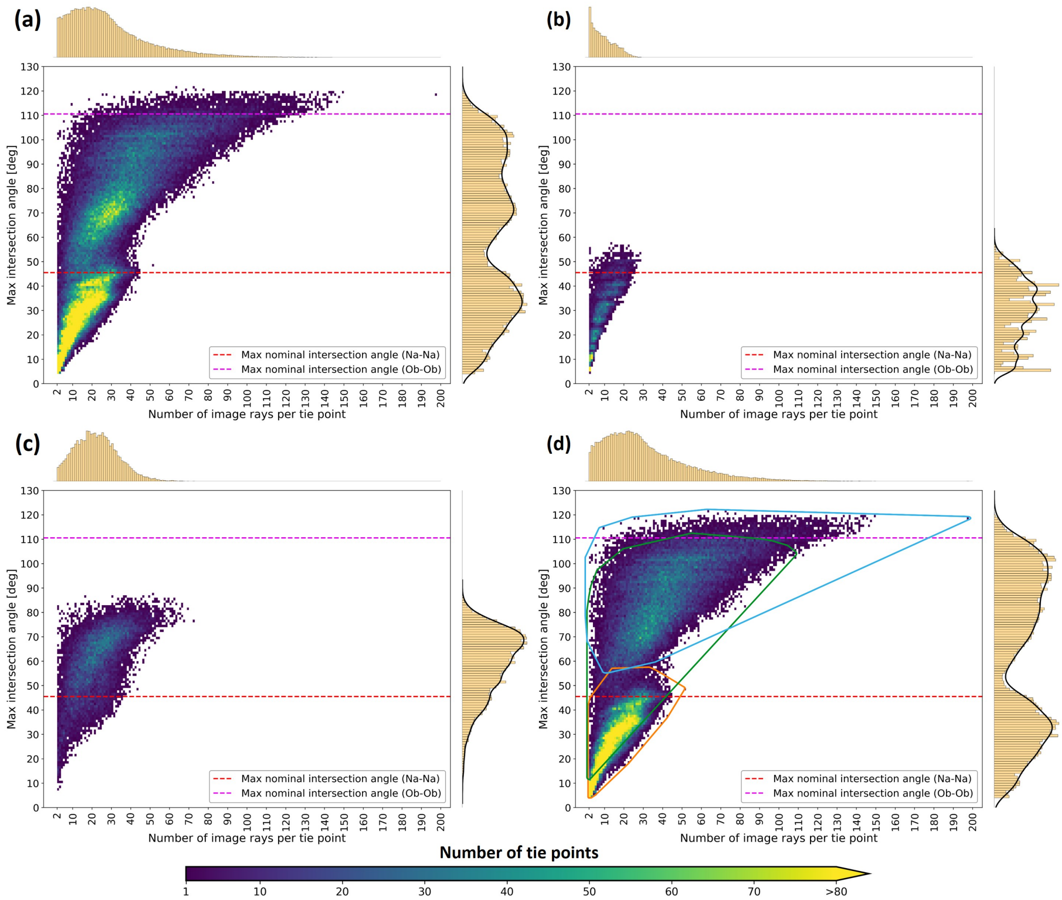

Figure 4. Maximum image ray intersection angle as a function of tie point multiplicity for Dataset 3. (a) All directional combinations; (b) nadir–nadir (Na–Na); (c) nadir–oblique (Na–Ob); (d) oblique–oblique (Ob–Ob). Dashed lines indicate nominal maximum intersection angles. Colored envelopes (orange, green, and blue) in (d) represent sub-configurations with directional differences of 0°, 90°, and 180°. The color scale indicates the number of tie points. Marginal histograms illustrate the distributions of tie point multiplicity and maximum intersection angle.

inal overlap percentage alone cannot serve as a reliable proxy for geometric network strength in multi-head UAV systems.

4.3 Image Ray Intersection Angle

The image ray intersection angle quantifies the angular separation between viewing rays from different camera positions to the same 3D point. It serves as a key geometric indicator, as larger intersection angles generally improve spatial precision, particularly in depth estimation. For each tie point, all viewing rays were considered pairwise, and the intersection angle was computed for every combination. The maximum pairwise intersection angle per tie point was used for statistical evaluation, as it represents the strongest geometric constraint available.

Figure 4 presents 2D histograms of the maximum intersection angle versus tie point multiplicity. The nominal maximum angles between nadir-to-nadir (Na–Na), nadir-to-oblique (Na–Ob), and oblique-to-oblique (Ob–Ob) pairs are 46°, 78°, and 110°, respectively. These values are derived under idealized assumptions of flat terrain, constant flying height, fixed oblique viewing angles of 45° relative to nadir (i.e. off-nadir viewing direction), and strictly nadir-oriented cameras; that is, all cameras have viewing directions of either 0° (nadir) or 45° (oblique).

All subplots are based on the same reconstructed tie point set, meaning that no tie points are excluded based on viewing direction. Tie point multiplicity is defined as the total number of image observations per point, including both nadir and oblique images, and is therefore identical across all subplots. In contrast, for each subplot, the maximum intersection angle is computed only from the corresponding subset of image ray pairings (Na–Na, Na–Ob, or Ob–Ob). Thus, while the intersection angle depends on the selected ray pairing, the multiplicity reflects the overall visibility of a point within the image network.

Figure 4a shows the combined result for Dataset 3, including all directional pairings. Figures 4b–4d present the individual configurations: (b) Na–Na, (c) Na–Ob, and (d) Ob–Ob. For the Ob–Ob configuration, three subcases are distinguished according to the angular difference between camera viewing directions: (i) 0° (same viewing direction), (ii) 90°, and (iii) 180°. These subcases are indicated by colored envelopes representing the geometric bounds of each configuration.

Two principal observations emerge. First, only a relatively small proportion of tie points attain their maximum intersection angle through Na–Na configurations. This can be explained not only by the limited angular separation of nadir viewing rays but also by the comparatively small number of available Na–Na ray pairs per tie point. Since tie points are typically observed by fewer nadir images than oblique images, the number of possible Na–Na ray combinations is inherently limited, reducing the likelihood that a pair approaches the nominal maximum intersection angle. Consequently, purely nadir imagery contributes only limited geometric strength to the bundle block. In contrast, the majority of tie points achieve their strongest intersection geometry through

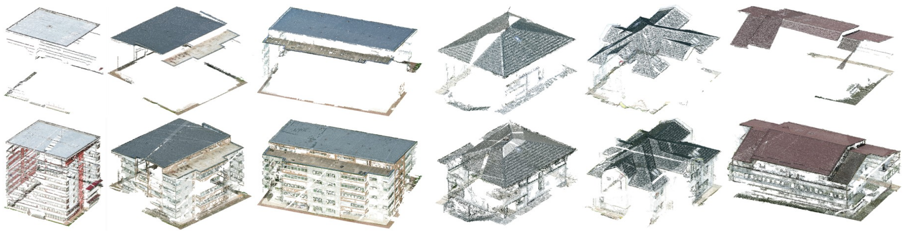

Figure 5. Comparison of dense 3D building reconstructions for Dataset 6 (Iași). Top: nadir-only configuration (1CAM); bottom: multi-head configuration (5CAM).

configurations involving oblique imagery (Na–Ob and Ob–Ob). These configurations provide a substantially larger number of ray pair combinations, particularly in the case of Ob–Ob, where multiple viewing directions are available. The increased number of combinations raises the probability that at least one pair forms a geometrically favorable intersection with a large angular separation. This confirms that oblique views enhance not only angular diversity but also the combinatorial richness of ray geometry, thereby strengthening object-space reconstruction.

Second, nearly all observed maximum intersection angles lie at or slightly below their nominal limits, indicating that the realized acquisition geometry closely follows the planned configuration. The ranking of intersection angle strength mirrors the connectivity hierarchy discussed in Section 4.2. Configurations formed by images with identical viewing directions (Na–Na) yield the smallest angles, whereas increasing directional diversity, from Na–Na to Na–Ob and further to Ob–Ob, leads to progressively larger intersection angles. This pattern demonstrates that directional diversity, rather than overlap magnitude alone, governs the potential for depth precision.

Overall, the empirical distributions align well with nominal expectations. The observed minimum and maximum intersection angles indicate stable flight execution, with only minor deviations in pitch and roll. Occasional values that slightly exceed the nominal limits can be attributed to small discrepancies between the idealized geometric assumptions used to derive the nominal bounds and the estimated exterior orientation parameters obtained from bundle adjustment. These deviations are marginal and lie within the expected uncertainty range. The fact that most tie points do not attain the nominal maximum intersection angle is consistent with the spatial distribution of tie points in a realistic image network. The nominal maximum represents an ideal geometric upper bound under specific viewing constellations, whereas in practice tie points are observed under varying angular configurations. Consequently, the realized intersection angles reflect the overall network geometry rather than indicating deficiencies in flight execution or parameter estimation. These results demonstrate that oblique imagery does not merely increase redundancy; rather, it fundamentally enhances the angular geometry governing depth accuracy.

4.4 Point Cloud and Façade Completeness

The final stage of the analysis evaluates reconstruction completeness using façade-specific point cloud density metrics. Dataset 6 (Iași) was selected due to its diverse urban morphology. Twenty buildings were analyzed, comprising 10 tall (≥4 storeys) and 10 low-rise (≤3 storeys) structures, enabling an assessment of how acquisition geometry affects façade reconstruction across different typologies.

Figure 5 presents representative dense reconstructions generated with *PMVS* from nadir-only (1CAM) and multi-head (5CAM) imagery. Nadir-only reconstructions exhibit façade gaps and limited vertical coverage, particularly for tall buildings. The inclusion of oblique imagery substantially improves façade continuity and roof–wall definition, consistent with the enhanced ray intersection geometry demonstrated in the previous sections.

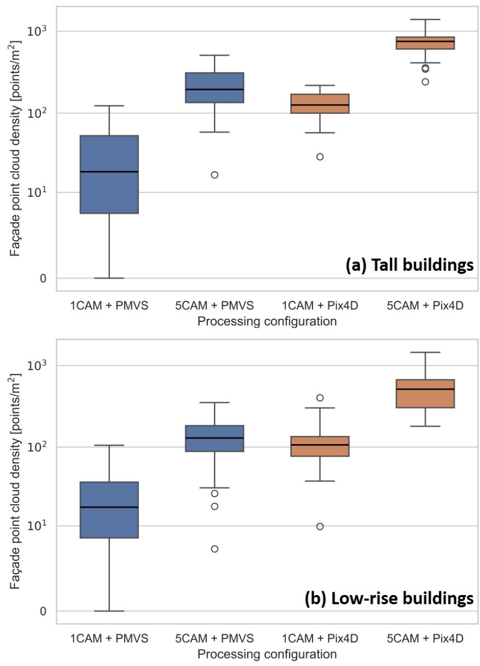

Figure 6. Boxplots indicate the median and IQR of façade point cloud density for (a) tall and (b) low-rise buildings under four processing configurations.

For quantitative evaluation, façade sides were extracted following Section 3.4, and façade point density (ρ_f) was computed for each side. Four processing configurations were analyzed: *1CAM + PMVS*, *5CAM + PMVS*, *1CAM + Pix4D*, and *5CAM + Pix4D*. *Pix4D* was included as a comparative reference representing a commercially established multi-view stereo workflow that incorporates regularization and interpolation during dense reconstruction. Consequently, differences in façade point density reflect not only acquisition geometry but also methodological differences in dense matching and surface regularization.

Figure 6 summarizes the median façade densities and interquartile ranges for tall and low-rise buildings. In both software pipelines, multi-head configurations consistently produce higher façade densities than nadir-only configurations, as reflected by increased median values. The spread of the distributions also differs between the workflows, with *PMVS* exhibiting a noticeably wider interquartile range, an effect further emphasized by the logarithmic scaling of the density axis, indicating greater variability in façade point density. The improvement is markedly stronger for tall buildings, where nadir-only geometry provides limited vertical intersection strength.

The density ratio (ρ_{5CAM}/ρ_{1CAM}) equals 10.503 and 7.349 for tall and low-rise buildings, respectively, in the *PMVS* workflow, and 5.959 and 4.830 in the *Pix4D* workflow. The larger relative gain observed with *PMVS* reflects its conservative patch-based reconstruction strategy, whereby additional viewing geometry directly translates into valid façade points. In contrast, *Pix4D* produces denser baseline reconstructions even under nadir-only conditions, resulting in smaller, yet still substantial, relative improvements.

5. Conclusions and Future Work

This study investigated the influence of acquisition geometry and computational processing on the geometric performance and reconstruction completeness of small-format multi-head UAV camera systems. Six datasets acquired under varying flight altitudes, overlaps, platforms, and urban conditions were analyzed using a consistent Structure-from-Motion and bundle block adjustment workflow. A geometry-driven evaluation framework was applied, integrating image connectivity, ray intersection geometry, tie point multiplicity, adjustment diagnostics, and façade-specific density metrics within a unified comparative setting.

The results demonstrate that acquisition geometry is a decisive factor for both adjustment stability and reconstruction completeness. Connectivity and intersection-angle analyses confirm that oblique imagery strengthens the directional network structure and increases effective intersection angles. Although the inclusion of oblique cameras does not universally reduce the a posteriori variance factor (σ_0), this behavior may reflect limitations in the adopted stochastic or functional model. In particular, potential differences in image quality or noise characteristics between nadir and oblique cameras were not explicitly considered in the weighting scheme. Quantitative façade-density analysis reveals substantial improvements in reconstruction completeness when oblique imagery is included. For tall buildings, façade point density increases by approximately one order of magnitude in the *PMVS* workflow and by a factor of five in the *Pix4D* workflow. These results confirm that façade completeness is primarily governed by acquisition geometry, while the reconstruction algorithm modulates the magnitude, but not the direction, of improvement. Overall, the findings demonstrate that nominal overlap alone is insufficient to characterize geometric network strength in multi-head UAV systems. Instead, directional viewing diversity constitutes the fundamental driver of depth precision and façade reconstruction quality.

Future work will focus on refining stochastic modeling within bundle adjustment, extending façade completeness analysis to additional urban morphologies, and investigating adaptive acquisition strategies that optimize viewing geometry for task-specific mapping objectives.

6. Acknowledgments

The authors express their sincere gratitude to the Department of Survey Engineering, Faculty of Engineering, Chulalongkorn University, and to N.P. Survey Company Limited, Thailand, for providing and supporting the majority of the datasets used in this study. T.B. received scholarship funding from the Royal Thai Government.

References

Bannakulpiphat, T., Karel, W., Ressl, C., Pfeifer, N., 2024. Optimizing Bundle Block Adjustment for High-Overlap Small-Format Multi-Head Camera Systems. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLVIII-2/W7-2024, 17–24. https://isprs-archives.copernicus.org/articles/XLVIII-2-W7-2024/17/2024/.

Bannakulpiphat, T., Karel, W., Ressl, C., Pfeifer, N., 2025. Evaluation of Bundle Block Adjustment for 3D Geospatial Modelling Using UAV-Based Small-Format Multi-Head Camera Systems. *RevCAD Journal of Geodesy and Cadastre*, 39, 13–24.

Fischler, M. A., Bolles, R. C., 1981. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6), 381–395. https://doi.org/10.1145/358669.358692.

Furukawa, Y., Ponce, J., 2010. Accurate, dense, and robust multiview stereopsis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(8), 1362–1376.

Gerke, M., Nex, F., Remondino, F., Jacobsen, K., Kremer, J., Karel, W., Huf, H., Ostrowski, W., 2016. Orientation of oblique airborne image sets-experiences from the ISPRS/EUROSDR benchmark on multi-platform photogrammetry. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences-ISPRS Archives 41 (2016)*, 41, 185–191.

Jiang, S., Jiang, W., 2017. Efficient structure from motion for oblique UAV images based on maximal spanning tree expansion. *ISPRS Journal of Photogrammetry and Remote Sensing*, 132, 140–161.

Karel, W., Doneus, M., Verhoeve, G., Bries, C., Ressl, C., Pfeifer, N., 2013. ORIENTAL - Automatic Geo-Referencing and Ortho-Rectification of Archaeological Aerial Photographs. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, II-5/W1, 175–180. https://isprs-annals.copernicus.org/articles/II-5-W1/175/2013/.

Karel, W., Ressl, C., Pfeifer, N., 2016. Efficient orientation and calibration of large aerial blocks of multi-camera platforms. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 41, 199–204.

Maset, E., Magri, L., Fusiello, A., 2024. Principled bundle block adjustment with multi-head cameras. *ISPRS Open Journal of Photogrammetry and Remote Sensing*, 11, 100051.

Nex, F., Gerke, M., Remondino, F., Przybilla, H.-J., Bäumker, M., Zurhorst, A., 2015. ISPRS benchmark for multi-platform photogrammetry. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 2, 135–142.

Pierrot-Deseilligny, M., Jouin, D., Belvaux, J., Maillet, G., Girod, L., Rupnik, E., Muller, J., Daakir, M., Choqueux, G., Deveau, M., 2014. Micmac, apero, pastis and other beverages in a nutshell. *Institut Géographique National*, 14.

Remondino, F., Gerke, M., 2015. Oblique aerial imagery : a review. D. Frietsch (ed.), *Proceedings of Photogrammetric Week '15, 7-11 September, Stuttgart, Germany*, Wichmann, 75–83.

Rupnik, E., Nex, F., Toschi, I., Remondino, F., 2015. Aerial multi-camera systems: Accuracy and block triangulation issues. *ISPRS Journal of Photogrammetry and Remote Sensing*, 101, 233–246.

Schneider, J., Förstner, W., 2013. Bundle adjustment and system calibration with points at infinity for omnidirectional camera systems. *Photogrammetrie–Fernerkundung–Geoinformation (PFG)*, 4, 309–321.

Schonberger, J. L., Frahm, J.-M., 2016. Structure-from-motion revisited. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4104–4113.

Sun, Y., Sun, H., Yan, L., Fan, S., Chen, R., 2016. RBA: Reduced Bundle Adjustment for oblique aerial photogrammetry. *ISPRS Journal of Photogrammetry and Remote Sensing*, 121, 128–142.

Toschi, I., Ramos, M. M., Nocerino, E., Menna, F., Remondino, F., Moe, K., Poli, D., Legat, K., Fassi, F., 2017. Oblique Photogrammetry Supporting 3D Urban Reconstruction of Complex Scenarios. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, XLII-1/W1, 519–526. https://isprs-archives.copernicus.org/articles/XLII-1-W1/519/2017/.