

Cold Region Building Inspection using UAV-based Three-dimensional Reconstruction

Nicholas Chodura^{1,2}, Derek Boase^{1,3}, Joshua Woods^{1,2}, and Melissa Greeff^{1,3}

¹Ingenuity Labs Research Institute, Queen's University, Canada -
(nicholas.chodura, derek.boase, joshua.woods, melissa.greeff)@queensu.ca

²Department of Civil Engineering, Queen's University, Canada

³Department of Electrical and Computer Engineering, Queen's University, Canada

Keywords: Infrastructure Inspection, Building Envelope, Infrared Thermography, Snow Depth, Uncrewed Aerial Vehicles, Photogrammetry, 3D Reconstruction

Abstract

This paper presents a UAV-based workflow for the inspection of buildings in cold-region climates, focusing specifically on rooftop snow depth estimation and thermal reconstruction of building envelopes. Accurate measurement of rooftop snow depth is critical, as an excessive snow load can lead to structural damage or roof collapse, resulting in substantial economic losses and potential fatalities. Thermal reconstruction can indicate thermal bridging in the building envelope, a phenomenon that can significantly reduce heating efficiency, causing increased energy consumption and higher utility expenses. Case studies conducted at multiple locations demonstrate the efficacy of UAV photogrammetry in accurately measuring rooftop snow depth, validated using manual field measurements and LiDAR scanning. We also propose a method of joint RGB-thermal reconstruction, capable of producing models with a high degree of geometric and radiometric accuracy without the requirement of homography or 3D transformations for image pair alignment. Results validate the approach through field measurements and pixel-wise thermal model evaluation. The proposed methods provide efficient, accurate, and safe alternatives to traditional inspection practices.

1. INTRODUCTION

The ubiquitous application of aerial vehicles in surveillance and monitoring tasks is an artifact of the recent advancements in mobile technologies (Price et al., 2022). The growth in these fields opens the door to more frequent and broad uses of these vehicles in tasks that represent extreme challenges or dangers to workers, such as roof inspection (Graybeal et al., 2002). In such cases, photogrammetry can be leveraged to offer the end user a more global and context rich representation of a physical asset, allowing the user to query the data from a 3D model without exposure to the same dangers and limitations that accompany manual inspection.

Within global cold-climate regions, there are additional challenges that can impose more requirements for the inspection process. One of these challenges is the accumulation of snow and ice on rooftops. Heavy snowfalls have led to numerous roof collapses, with incident frequency and severity varying geographically. North America has documented a high rate of structural failures, highlighted by a 20-year survey that recorded 1,029 snow-induced building collapses in the United States from 1989 to 2009 (Geis et al., 2012). Scandinavia and Central Europe have also experienced winters with particularly severe snowfall, including the winter of 2010-2011, where at least 85 farm buildings in Sweden suffered serious damage or collapse (Nilsson and Friberg, 2012), and the winter of 2005-2006, where over 200 roof failures occurred across Central Europe (Croke et al., 2018). Common causes identified in these failures are extreme snow accumulation beyond design assumptions, often exacerbated by rain-on-snow events or drift piling, combined with structural deficiencies or maintenance issues that undermine capacity (Geis et al., 2012, Nilsson and Friberg, 2012). Impacts can range from substantial economic losses to loss of life. Individual incidents have incurred hundreds of millions in damages (a single blizzard in 1993 caused over \$200 million in U.S. roof damage) (Geis et

al., 2012). The Katowice Trade Hall collapse in Poland claimed 65 lives when the long-span exhibition hall roof collapsed under extreme snow load in January 2006 (Biegus and Rykaluk, 2009). Moscow's Basmanny Market roof collapse in 2006 is another example of snow-based roof collapse, resulting in 68 fatalities and 32 injuries (Lobkina, 2021).

A primary factor contributing to these incidents is the lack of regular roof inspection and maintenance (Holický and Sýkora, 2009). Many jurisdictions do not mandate the routine inspection of rooftops by building managers and, due to the cost of traditional inspections, many buildings do not receive regular inspection. By reducing the cost and time required for inspections through the use of UAVs and 3D Reconstruction, routine winter inspections can become more common and the chances of snow-based roof collapse can be reduced.

Another challenge that buildings face in cold climate regions is thermal bridging, where structural discontinuities within the building envelope result in localized areas of increased heat transfer. Thermal bridges typically occur where materials with higher thermal conductivity, such as concrete or metal structural elements, penetrate or bypass insulation layers. This commonly occurs at studs, joists, balconies, and junctions between external walls, roofs, and floors. This significantly reduces the thermal performance of buildings, causing increased heat loss, reduced interior air temperatures, and heightened risks of condensation and mold growth (Alhawari and Mukhopadhyaya, 2018). In well-insulated residential buildings situated in cold climates, studies indicate that thermal bridges can be responsible for a 30% increase of total heating energy loss (Theodosiou and Papadopoulos, 2008). Consequently, accurate detection, assessment, and mitigation of thermal bridging effects are essential to achieving sustainable energy efficiency, ensuring occupant comfort, and maintaining the long-term integrity of building assets.

Infrared Thermography (IRT) has become a popular method for

identifying this issue, using terrestrial or UAV-based sensors to measure the radiometric properties of the building surface (Tomita and Chew, 2022). Advances in thermal 3D reconstruction have provided opportunities to visualize thermal bridging in structures and better interact with digitized models of these assets.

With data collected from multi-spectral sensor systems, UAV-based 3D reconstruction can provide a cost-effective alternative to traditional inspections, while allowing asset managers and engineers to analyze 3D building data off-site.

2. RELATED WORK

2.1 Rooftop Snow Depth

UAV-based snow depth measurement has emerged over the past decade as a versatile technique applicable across several fields, including alpine hydrology, avalanche forecasting, ecological studies, and infrastructure assessment. Early efforts predominantly targeted alpine catchments and mountain hydrology, where precise snow-depth mapping can provide improved snow water equivalent (SWE) estimation and water resource forecasting (Bühler et al., 2015, Maier et al., 2022).

Despite extensive research in natural environments, fewer studies have explored the application of UAV snow depth measurement in infrastructure contexts, particularly for assessing rooftop snow loads. One study investigated photogrammetry for rooftop snow depth measurement (Chiba and Thiis, 2016). Four roofs were identified for testing, and snow probes were used for field measurement of the snow depth. A combination of Remotely Piloted Aircraft System (RPAS) and handheld cameras was used to capture the data with image overlaps of at least 60%. The method for extracting model snow depths was not specified, but it appears to be discretely sampled across the roof surface. Potentially due to this low overlap, or the method of depth estimation, the mean snow depth error was approximately 9 cm for the roof that relied entirely on RPAS image captures. Due to these higher depth errors, and a lack of clearly visualized, high-resolution rooftop snow depth information, further research is required in this field to better understand this workflow.

Our paper presents a workflow for UAV-based snow depth estimation that uses consistent, replicable flight parameters that can be specialized to specific roof requirements. This method results in a mean error of less than 1 cm, and the photogrammetry workflow provides comparable results to LiDAR scanning.

2.2 Building Envelope Thermal Reconstruction

Three-dimensional thermal reconstruction combines thermal imaging with 3D geometry reconstruction, offering valuable insights for building diagnostics and structural inspections. Traditional photogrammetric methods, such as Structure-from-Motion (SfM) and Multi-View Stereo (MVS), are widely adopted but face challenges when applied directly to thermal imagery due to low resolution, low contrast, and feature sparsity (Ramón-Constantí et al., 2022). Recent research emphasizes multi-modal fusion techniques to address limitations inherent to thermal images. RGB-IR image fusion has become prevalent, leveraging high-resolution RGB imagery for geometry reconstruction and thermal images for surface temperature mapping.

2D image transformations have been applied in prior work to provide efficient methods of image alignment between RGB

and IR image pairs, including overlap transformations (Andras et al., 2019), and Normalized Cross-Correlation (NCC) image registration (Yang et al., 2018). A popular method of image alignment is through a homography transformation, whereby the RGB images are shifted in some way to fit the image contents and size of the thermal images. This can be done using 2D affine transformations (Javadnejad et al., 2020), or using calibration targets with various detection algorithms (Adamopoulos et al., 2020, Daffara et al., 2020, Motayyeb et al., 2023), among other methods. The primary issue with these 2D transformations is they assume a perfectly planar scene and therefore neglect any depth variations, such as parallax, that are present in real-world imagery. When images are captured with oblique views to better represent non-planar building components, relying on homography or other 2D transformations for alignment can lead to inaccuracies in the final registration between thermal and RGB images (Adamopoulos et al., 2020).

3D pose estimation methods have been developed that aim to address these issues, including precise modeling of the 3D lever arm and boresight geometry of the camera rig (Javadnejad et al., 2020), and registration of the thermal images to an RGB 3D model by matching features between modalities, projecting them onto the RGB point cloud, and iteratively refining the thermal pose with Gauss–Newton optimization (Lin et al., 2025). 3D camera pose estimation requires extremely precise multi-camera calibration, which is particularly challenging with consumer-grade thermal cameras due to lower resolution, low-contrast thermal images. Consequently, these calibration uncertainties can lead to significant errors in 3D pose estimation (Javadnejad et al., 2020).

Methods that do not rely on image alignment are less common. A study used manually collected building measurements to create a 3D CAD model that could be textured using thermal imagery. However, the resulting 3D model lacks detailed building features due to the manual model generation method (Lagüela et al., 2014). Iwaszczuk et al. proposed a method with a prior known building model, where thermal camera poses are refined based on line-based model-to-image matching techniques (Iwaszczuk and Stilla, 2017). This method was also limited by the requirement for a manually constructed CAD building model.

2.3 Contributions

Through the study outlined in this paper, the authors present two main contributions to these fields of research:

1. A method for UAV-based rooftop snow depth estimation using low-cost and open-source tools. A field study validated the method using 17 rooftop snow probe measurements and LiDAR scanning.
2. A UAV-based RGB-IR 3D reconstruction workflow that does not use image alignment between the RGB and IR image sets, eliminating the need for complex homography or 3D transformations. Workflow was validated through radiometric control points and pixel-wise temperature mean average error (MAE) from selected camera poses.

3. 3D RECONSTRUCTION

This section introduces the high-level workflow for 3D reconstruction from multi-spectral data gathered using UAVs, detailing the stages of a 3D reconstruction pipeline that is used in this work.




Figure 1. Extracted features using Reality Capture, seen in green.



Figure 2. Tie-points connecting features from the overlapping regions of each constituent image, seen in orange.

The initial phase of the photogrammetry pipeline is the Data Acquisition (DAQ). The quality of the reconstruction is sensitive to factors including lighting conditions and time of day (Burdziakowski and Bobkowska, 2021, Luna Torres et al., 2024), seasonal variations and time of year (Brown et al., 2024, Luna Torres et al., 2024), and user-defined flight parameters including flight altitude and ground sampling distance (GSD) (Chodura et al., 2025). For physical assets such as building envelopes, high-quality reconstructions may require hundreds to thousands of images. Using UAVs with automated or even autonomous flight regimes and multi-spectral sensors may be performed with relatively reduced resources compared to manual collection.

Following the DAQ stage, calibration focuses on performing a preliminary refinement of the sensor data. This step extracts the camera intrinsic parameters (focal length, principal point, distortion coefficients) and performs sensor calibration—both geometric and radiometric—to correct lens distortion or account for environmental factors, including temperature-dependent sensor drifts and noise due to ambient light and temperature. It is also the point at which image enhancements (e.g., brightness, contrast) may be applied, improving the effectiveness of feature extraction during the next stage.

In the feature extraction and matching block, keypoints and descriptors are extracted from the images. Examples of the keypoints are shown in Figure 1 as green dots. This is commonly done using local feature-based matching for traditional images and techniques like L2-SIFT for aerial data (Jiang et al., 2020). The keypoints that are extracted from the image form the initial approximation of the sparse point cloud (López et al., 2021). At the end of this stage, tie-points are found connecting the 2D features from overlapping regions of the constituent images, shown as orange lines in Figure 2.

In structure from motion (SfM) the relative poses of the cameras in the scene are approximated and a sparse point cloud is recovered, representing a rough initial geometry estimate. This phase does not require apriori knowledge of the camera poses, although they can be used if available for faster convergence (Jiang et al., 2020). The camera intrinsics are further refined in this phase.




Figure 3. Photogrammetry pipeline used for 3D reconstruction.

Following SfM, the bundle adjustment stage refines the camera and scene parameters under a single global optimization, minimizing the reprojection error. This optimization is applied across the dataset, leading to improvements in the camera pose estimations and the sparse point cloud reconstruction (Seitz et al., 2006). Although this stage improves the current state of the reconstruction, it tends to be expensive, especially for datasets with large image counts or feature dense images. The sparse point cloud is densified in the penultimate stage of multi-view stereo (MVS). This is done by traversing the constituent image pixels and estimating new points outside of the tie-points found in earlier stages. This geometry augmentation is responsible for a richer geometry approximation (Seitz et al., 2006). Post-processing is an optional step for user-defined manipulation of the 3D model. This can include mesh texturing, hole filling, smoothing, and decimation (reducing the number of polygons in a mesh).

4. METHODS

All testing was conducted using the UAV platforms outlined in Table 1, with the DJI Matrice 350 RTK being used in conjunction with the DJI Zenmuse H20T for thermal data collection and the DJI Zenmuse L2 for LiDAR data collection. Preliminary field testing was conducted at Ellis Hall on Queen’s University campus in Kingston, Ontario as shown in Figure 4. Secondary locations include Mitchell Hall on Queen’s campus, and commercial buildings located in Kingston and Ottawa that have been kept anonymous for privacy reasons. For the photogrammetry flights, double-grid boustrophedon flight patterns were used, with overlap and ground sampling distance parameters based on a previous rooftop 3D reconstruction study (Chodura et al., 2025).

Table 1. Platform, Camera, and Image Specifications

Platform	Camera	Format	Resolution
Matrice 350	H20T, Thermal	RJPEG	640x512
	H20T, Wide	JPEG	4056x3040
Mavic 3E	L2	Point cloud	-
Mini 3 Pro	Wide	JPEG	5280x3956
	-	JPEG	8064x6048

4.1 Rooftop Snow Depth

Snow depth and volume are calculated using vertical axis volumetric differencing between a no-snow building model and the building covered in snow. For the case study on Ellis Hall, photogrammetry and LiDAR were used for comparison, and manual measurements of the snow depth were taken at various points on the roof immediately after the flight for validation. Additional small-scale testing was done on snow piles of known volume to further validate the method. The final result is a scalar field displaying snow depth data across an entire roof surface at a high spatial resolution.



Figure 4. Ellis Hall on Queen's campus, Kingston, with snow depth measurements identified with black circles (SD_x) and thermal control points identified with red triangles (TCP_x).

4.1.1 Data collection. Data was collected using the same flight paths for the no-snow and snow models to ensure consistent model accuracy and coverage. For the photogrammetry flights, double-grid boustrophedon flight paths were used. The first two used flight parameters of 85% overlap, 0.75 - 1.25 cm/pixel ground sampling distance (GSD), and a camera angle of 65 degrees. These paths were shifted 45 degrees from one another to provide eight oblique viewing directions of the building surface. The third path is a lower flight to capture surface details at a higher resolution, with an overlap of 80%, GSD of 0.5 - 1.0 cm/pixel, and nadir camera angle.

The LiDAR flights consisted of five flights, with the first flight being an orthographic boustrophedon pattern at nadir camera angle, with a flight altitude of approximately 70 meters over the target surface, and a LiDAR side overlap of 87%. The following four flights are linear oblique paths at the same altitude and a camera angle of 70 degrees, with flights facing the building from each side. UAV localization was done using Real-Time Kinematics (RTK) GPS through a mobile base station, and inertial measurement units (IMUs) on-board the UAV. The flight plans included calibration segments of the flight at the start, end, and between direction changes to improve IMU accuracy.

4.1.2 Error metrics. Snow depth measurements were made at 17 points on the roof, as shown in Figure 4. Measurements were made using a snow probe inserted into the snow down to the roof surface to record the depth. Laser rangefinders were used to identify the exact position of measurement locations. Using the position data of the control points, equivalent points on the model were identified and the model depths were compared to the field measurements.

4.1.3 Data processing. Photogrammetry models are processed using Reality Capture, a photogrammetry software developed by Epic Games (Reality Capture Development Team, 2025). Following the alignment, meshing, and texturing stages, a PLY model is exported and brought into CloudCompare, a 3D point cloud processing software, for analysis (CloudCompare Development Team, 2025). After both the no-snow and snow

models have been brought into CloudCompare, they are manually segmented of their surroundings and aligned using equivalent point pairs and an iterative closest point (ICP) algorithm set to continue until an RMS of 10^{-5} . The mesh vertices are used for the analysis, and these point clouds are spatially sub-sampled to a minimum point distance of 5 mm for uniformity across the mesh surface and between compared models. The volume analysis is done using a '2.5D Volume' tool. This refers to the principal direction that must be specified on which the volume is computed. The tool uses a rasterization process to divide the clouds into an n by m grid of cells, specified by a user-defined step size Δx , in this case set to 0.01 m. Where there are points in the principle plane identified in both clouds (ground and ceiling), the difference in height is computed as the planar distance between these two points. Volume across the surface can then be defined by the following:

$$V = \sum_{i=1}^n \sum_{j=1}^m \Delta x^2 d_{ij}, \quad (1)$$

where V = volume difference
 Δx = step size
 d_{ij} = difference in height at grid cell (i, j)
 $n \times m$ = grid dimensions
 i, j = indices in the grid

and in the case of a grid step missing either a ground or ceiling value, $d_{ij} = 0$. Following the computation in Equation (1), a global volume difference can be determined, and the grid can be exported as a new point cloud, showing the difference in height, or in this case snow depth, for every cell in the XY-plane (CloudCompare Development Team, 2015).

4.2 Building Envelope Thermal Reconstruction

This paper proposes a joint RGB-IR 3D reconstruction method capable of producing models of high geometric and radiometric accuracy. The higher spatial resolution and image contrast of the RGB images is used to create a high-quality mesh surface, and the thermal images are used for the model texture.

Previous methods of thermal-RGB reconstruction require complex homography, using feature matching between image pairs for pose estimation of the thermal cameras within the RGB point cloud environment. The method used in this paper instead uses strategically selected flight parameters and post-processing to improve the clarity and dynamic contrast of the thermal images so that features can effectively be extracted and camera poses can be accurately estimated.

4.2.1 Data collection. The thermal data was collected using the DJI Zenmuse H20T, which produces radiometric 640 x 512 RJPEG images using an uncooled VOx microbolometer sensor. The thermal camera has a 40.6 degree FOV and a 58 mm focal length, and functions in a thermal range of -40 °C to 150 °C. H20T specifications claim a $\pm 2^\circ\text{C}$ or $\pm 2\%$ accuracy at 5 m in laboratory conditions. Preliminary testing indicated a loss of fine detail when flying at distances greater than 20 m to the target surface, so a flight height of 5 m - 15 m over target surfaces was used. A reduced flight speed of 0.85 m/s was used to reduce the likelihood of motion blur, a phenomenon more common in thermal imagery. With the proposed thermal reconstruction workflow, it was critical that the thermal images alone could

establish enough tie points to achieve a successful alignment. Because of this fact, combined with the narrow thermal FOV and low flight altitude, higher image overlaps of 80 % and 70 % were used for the orthogonal and oblique thermal flights, respectively. Thermal images were taken in conjunction with RGB images using the H20T's wide camera lens, featuring a 1/2.3" CMOS 12 MP sensor and producing 4056 x 3040 RGB images.

4.2.2 Error metrics. Radiometric control points were taken at 8 locations on the building envelope and rooftop to validate the UAV-based thermal images and final model. A commercial laser thermometer with a temperature range of -40 °C to 600 °C was used, with a reported reading accuracy of ± 2 °C below 0 °C and ± 1 °C above 0 °C. The control points were compared with points on the mesh surface of the model. In addition to model validation through field measurements, temperature Mean Average Error (MAE) at various camera poses across the model was used to evaluate whether the thermal data captured was accurately represented in the 3D reconstruction.

4.2.3 Pre-processing. In DJI's RJPEG output files, thermal information is encrypted and difficult to access. A thermal tool application developed by UAV4GEO was used to convert the RJPEG files to TIF files (UAV4GEO Development Team, 2025). Image metadata is preserved in this process, and temperature information is stored for each pixel of the images. A script was then used to produce a set of PNG files from the TIF dataset, where a global relationship is defined between a pixel temperature and its color value. This method provides 255 distinct values for temperature to be assigned, linearly spread out between the highest and lowest pixel temperatures in the dataset. These maximum and minimum values were clipped to 18 °C and -25 °C for this experiment, respectively, which maintained 99.5% of pixel value temperatures while removing outliers that significantly reduced the contrast of the thermal images. These outliers were likely due to very low-emissivity materials reflecting ambient temperature, and heat vents producing small regions of high temperature. Regions outside the specified range were assigned the relevant maximum or minimum temperature.

4.2.4 Photogrammetry. Reality Capture was used for the photogrammetry pipeline. A typical end-to-end photogrammetry process would consist of the same source dataset being used for alignment, meshing, and texturing. Because of the higher image contrast and spatial resolution of the RGB images, this dataset was used for the meshing stage to provide a more geometrically accurate reconstruction of the building. In a separate project, the thermal images were imported and aligned. Due to the lower altitude flight and the pre-processing steps, as well as the RTK GPS metadata to establish position priors, a sparse cloud of tie points was produced. The sparse clouds of both reconstructions were exported as components, along with the estimated camera poses for both datasets, and the components were aligned using ICP and bundle adjustment global alignment techniques within Reality Capture. For the following stages, the RGB images were restricted from texturing, and the thermal images were restricted from meshing in order to produce the thermal reconstruction.

4.2.5 Post-processing. The resulting mesh was exported as a PLY file and brought into CloudCompare for post-processing. The surrounding environment was manually removed to provide an accurate representation of the buildings temperature range. The RGB colour channel values of all mesh vertices (for gray-scale, they are the same across the three channels) are converted into a scalar field.

For the point cloud of mesh vertices P , each point p has a colour value $c(p) \in [0, 255]$. A linear transformation can be defined:

$$T(p) = \frac{T_{\max} - T_{\min}}{255} c(p) + T_{\min}, \forall p \in P, \quad (2)$$

that maps from $c(p)$ to $T(p)$ such that $T(p) \in [T_{\min}, T_{\max}]$,

where $T(p)$ = temperature in °C at p
 $T_{\min}$ = minimum temperature
 $T_{\max}$ = maximum temperature
 $c(p)$ = color channel intensity value at p

using the known relationship between the color value and temperature determined using the pre-processing minimum and maximum temperatures. The model can then be displayed with a color map that better displays the temperature range, and points on the mesh can be sampled for temperature accuracy metrics.

5. RESULTS

5.1 Snow Load Results

Results from the Ellis Hall case study can be seen in Figure 5. Seen in the lighter colors, both the LiDAR models and photogrammetry models provided snow depth maps that effectively captured regions of built-up snow and drifts that had formed on the rooftop. Additionally, fine snow details were captured, included footprints left in the snow by an animal, seen on the highest roof section to the left of the telescope.

The field measurements collected on the rooftop were compared with depth information from the models, and the snow depth error at each control point is shown for LiDAR (δ_l) and photogrammetry (δ_p) in Table 2. Average snow depth error for these methods was 1.0 cm and 0.9 cm for LiDAR and photogrammetry, respectively.

Table 2. Measured Snow Depths Compared with Photogrammetry and LiDAR Model Depths

Measured		Lidar		Photogrammetry	
Point	Depth (cm)	Depth (cm)	δ_l (cm)	Depth (cm)	δ_p (cm)
SD_1	6.5	5.7	0.8	6.5	0
SD_2	7.5	6.9	0.6	8.2	0.7
SD_3	18.0	18.9	0.9	19.5	1.5
SD_4	12.5	12.0	0.5	12.5	0
SD_5	10.0	10.4	0.4	10.0	0
SD_6	11.5	12	0.5	12.1	0.6
SD_7	12.0	12.6	0.6	11.0	1.0
SD_8	9.5	10.2	0.7	9.6	0.1
SD_9	10.0	10.4	0.4	9.1	0.9
SD_10	13.5	13.4	0.1	12.0	1.5
SD_11	7.0	8.6	1.6	7.1	0.1
SD_12	11.0	12.1	1.1	10.7	0.3
SD_13	8.5	9.1	0.6	8.6	0.1
SD_14	15.5	13.1	2.4	14.3	1.2
SD_15	11.0	12.8	1.8	8.4	2.6
SD_16	14.0	17.9	3.9	15.5	1.5
SD_17	10.5	9.6	0.9	7.9	2.6
Average	-	-	1.0	-	0.9

Additional flights were conducted at three other buildings with varying degrees of snow cover to observe how photogrammetry performed in different conditions and settings. Rooftop access




Figure 5. Photogrammetry snow depth map (left) and LiDAR snow depth map (right).

was not permitted at these locations, and thus additional depth measurements were not taken. Resulting models appeared to show equivalent snow detail to the case study, from qualitative analysis of the reconstructions and depth fields.

5.2 Thermal Reconstruction Results

5.2.1 Geometric model accuracy. The geometric model accuracy determines the quality of the surface on which the temperature information is displayed. This can be measured qualitatively through visual examination and quantitatively through control points and cloud-to-cloud comparisons with a ground truth model. Figure 6 shows the RGB-IR 3D reconstruction. Close-up images on the right provide qualitative evidence of thermal re-projection accuracy and demonstrate details free from aliasing and distortion.

5.2.2 Thermal measurement accuracy. Absolute temperature error can occur due to incorrect calibration of environmental parameters such as ambient temperature and the emissivity of the target material. The absolute temperature error can be determined using field measurements with calibrated instruments. Absolute temperature error was evaluated at eight different locations on the building envelope and rooftop using a calibrated laser thermometer. The results can be seen in Table 3, and the testing yielded an average temperature error of 0.9 °C.

Table 3. Temperature Field Measurements Compared with Temperatures Sampled from the Model

Measured		Model	
Point	Temp (°C)	Temp (°C)	δ (°C)
TCP_1	7.3	7.6	0.3
TCP_2	10.2	10.0	0.2
TCP_3	-12.7	-8.9	3.8
TCP_4	9.5	8.5	1.0
TCP_5	11.3	11.3	0.0
TCP_6	14.9	15.7	0.8
TCP_7	-6.3	-5.6	0.7
TCP_8	9.6	8.9	0.7
Average	-	-	0.9

5.2.3 Aliasing and distortion. Aliasing and other distortions can occur during the texturing process as the thermal imagery is projected onto the existing mesh. This can be caused by a misalignment in the estimated camera poses or by inaccuracies in mesh geometry, lens distortion effects, or insufficient texture

resolution relative to the mesh. This can result in a loss of fine details displayed on the model. Aliasing and distortion can be observed both visually, and quantitatively by comparing pixel-wise temperature values from the IR source images with corresponding RGB values from rendered model views, enabling the evaluation of thermal and geometric model accuracy relative to the input images (Hassan et al., 2024).

Temperature MAE, considering every pixel of the compared images, was measured for five different camera poses. The randomly chosen images, the error map, and the MAE for each camera pose can be seen in Figure 7. The average MAE across all camera poses was 2.7 °C.

6. DISCUSSION

The results from the snow depth experiment yielded promising results, with both photogrammetry models and LiDAR models producing depth errors equal to or less than 1 cm. Additionally, the average depth difference between the two methods was less than 1.3 cm. LiDAR producing a higher error than photogrammetry could have been due to the higher levels of noise produced in the LiDAR point cloud, when compared to the photogrammetry cloud that was derived from a mesh surface. Qualitative observation of varying roof types and snow coverage indicate that this method can be applied to a range of snow conditions, but more measurement field tests are required to confirm this.

The proposed thermal reconstruction method produced a model with high geometric accuracy, high radiometric accuracy, and minimal aliasing and distortion. Additionally, the proposed method does not require complex homography or multi-modal image matching. Field measurements yielded an average absolute temperature error of 0.9 °C, an average temperature MAE of 2.7 °C between IR source images and the model. This error could be caused by incorrect projection of the thermal texture over the mesh resulting from misalignment of the camera poses or geometric deformations in the mesh. Additionally, it was observed that areas where shadows were present exhibited higher error. This is because the temperature of these areas shifts as the flight progresses, and these values are averaged to produce the final texture in the model.

Another notable observation in the thermal reconstruction shown in Figure 6 is the presence of significantly lower temperature




Figure 6. Joint RGB-thermal 3D reconstruction. After images from both datasets are aligned separately, the resulting sparse clouds of tie points are merged. The RGB images are used for meshing, due to their higher spatial resolution and image contrast. The texture is then projected onto the mesh using the thermal images, and a linear transformation can be applied to yield the corrected temperature values.




Figure 7. Grayscale images (top) display the thermal images collected from the UAV. Using the estimated camera poses from the photogrammetry alignment stage, a pixel-wise comparison between renders of the model and the IR source images was conducted to yield a temperature error for each pixel (bottom). The mean average error (MAE) for each camera pose is displayed in the bottom left.

readings on reflective surfaces such as the HVAC vent and telescope. These lower temperature values likely result from the low emissivity characteristics of metallic or polished surfaces. Emissivity is the efficiency of a surface in emitting thermal radiation relative to an ideal thermal radiator. Due to their reduced capacity to emit infrared radiation, these surfaces predominantly reflect ambient thermal radiation, causing them to appear artificially colder in thermal imagery. Consequently, the measured temperatures on these reflective surfaces may not accurately represent their true surface temperature.

7. CONCLUSION

This paper demonstrates the successful application of UAV-based 3D reconstruction techniques for assessing rooftop snow depths and thermographic properties of building envelopes in cold region climates. Photogrammetry-based reconstruction achieves accurate snow depth measurements, aligning closely with manual field observations and LiDAR scans. The presented thermal reconstruction method effectively combines high-resolution RGB data with thermal imagery, providing precise temperature mapping with low geometric and radiometric errors.

References

- Adamopoulos, E., Volinia, M., Girotto, M., Rinaudo, F., 2020. Three-Dimensional Thermal Mapping from IRT Images for Rapid Architectural Heritage NDT. *Buildings*, 10(10).
- Alhawari, A., Mukhopadhyaya, P., 2018. Thermal bridges in building envelopes – An overview of impacts and solutions. *International Review of Applied Sciences and Engineering*, 9, 31–40.
- Andras, M., Lovas, I., Domozi, Z., 2019. Photogrammetry on low resolution thermal pictures.
- Biegus, A., Rykaluk, K., 2009. Collapse of Katowice Fair Building. *Engineering Failure Analysis*, 16, 1643–1654.
- Brown, I., Ghaly, M., Greiser, C., Lam, N., Lehmann, P., 2024. Seasonal optimisation of drone-based photogrammetry in a heterogeneous boreal landscape. *Applied Vegetation Science*, 27.
- Burdziakowski, P., Bobkowska, K., 2021. UAV Photogrammetry under Poor Lighting Conditions—Accuracy Considerations. *Sensors*, 21(10).
- Bühler, Y., Marty, M., Egli, L., Veitinger, J., Jonas, T., Thee, P., Ginzler, C., 2015. Snow depth mapping in high-alpine catchments using digital photogrammetry. *The Cryosphere*, 9, 229–243.
- Chiba, T., Thiis, T., 2016. Accuracy of snow depth measurements on roofs measured with photogrammetry. *International Conference on Snow Engineering 2016*.
- Chodura, N., Greeff, M., Woods, J., 2025. Evaluation of flight parameters in uav-based 3d reconstruction for rooftop infrastructure assessment.
- CloudCompare Development Team, 2015. Cloudcompare wiki - 2.5d volume. Last edited: 19 October 2015, 13:50.
- CloudCompare Development Team, 2025. Cloudcompare – 3d point cloud and mesh processing software (version 2.12) [computer software].
- Croce, P., Formichi, P., Landi, F., Mercogliano, P., Bucchignani, E., Dosio, A., Dimova, S., 2018. The snow load in Europe and the climate change. *Climate Risk Management*, 20.
- Daffara, C., Muradore, R., Piccinelli, N., Gaburro, N., de Rubeis, T., Ambrosini, D., 2020. A Cost-Effective System for Aerial 3D Thermography of Buildings. *Journal of Imaging*, 6, 76.
- Geis, J., Strobel, K., Liel, A., 2012. Snow-Induced Building Failures. *Journal of Performance of Constructed Facilities*, 26(4), 377–388.
- Graybeal, B., Phares, B., Rolander, D., Moore, M., Washer, G., 2002. Visual Inspection of Highway Bridges. *Journal of Nondestructive Evaluation*, 21, 67–83.
- Hassan, M., Forest, F., Fink, O., Mielle, M., 2024. Thermonerf: Joint rgb and thermal novel view synthesis for building facades using multimodal neural radiance fields.
- Holický, M., Sýkora, M., 2009. Failures of roofs under snow load: Causes and reliability analysis. 444–453.
- Iwaszczuk, D., Stilla, U., 2017. Camera pose refinement by matching uncertain 3D building models with thermal infrared image sequences for high quality texture extraction. *ISPRS Journal of Photogrammetry and Remote Sensing*, 132, 33–47.
- Javadnejad, F., Gillins, D. T., Parrish, C. E., and, R. K. S., 2020. A photogrammetric approach to fusing natural colour and thermal infrared UAS imagery in 3D point cloud generation. *International Journal of Remote Sensing*, 41(1), 211–237.
- Jiang, S., Jiang, C., Jiang, W., 2020. Efficient Structure from Motion for Large-Scale UAV Images: A Review and a Comparison of SfM Tools. *ISPRS Journal of Photogrammetry and Remote Sensing*, 167, 230–251.
- Lagüela, S., Díaz Vilariño, L., Roca, D., Armesto, J., 2014. Aerial oblique thermographic imagery for the generation of building 3d models to complement geographic information systems. *2014 Quantitative InfraRed Thermography*.
- Lin, D., Yang, N., Miao, Q., Cui, X., Xu, D., 2025. True 3D thermal inspection of buildings using multimodal UAV images. *Journal of Building Engineering*, 100, 111806.
- Lobkina, V., 2021. Analysis of Roof Collapse Cases Caused by Snow Loads in Russia (2001–2021). *Sustainability*, 13, 13580.
- Luna Torres, A., Vergara Olivera, M., Almeida Del Savio, A., Gracey Bambarén, G., 2024. Effect of Climatological Factors on the Horizontal Accuracy of Photogrammetric Products Obtained with UAV. *Sensors*, 24(22).
- López, A., Jurado, J. M., Ogayar, C. J., Feito, F. R., 2021. An Optimized Approach for Generating Dense Thermal Point Clouds from UAV-Imagery. *ISPRS Journal of Photogrammetry and Remote Sensing*, 182, 78–95.
- Maier, K., Nascetti, A., van Pelt, W., Rosqvist, G., 2022. Direct photogrammetry with multispectral imagery for UAV-based snow depth estimation. *ISPRS Journal of Photogrammetry and Remote Sensing*, 186, 1–18.
- Motayyeb, S., Samadzadegan, F., Dadrass Javan, F., Hosseinpour, H., 2023. Fusion of UAV-based infrared and visible images for thermal leakage map generation of building facades. *Heliyon*, 9(3), e14551.
- Nilsson, C., Friberg, O., 2012. Collapses and damage to swedish farm buildings due to snow loads. *International Conference of Agricultural Engineering*.
- Price, E., Liu, Y. T., Black, M. J., Ahmad, A., 2022. Simulation and control of deformable autonomous airships in turbulent wind. M. H. Ang Jr, H. Asama, W. Lin, S. Foong (eds), *Intelligent Autonomous Systems 16*, Springer International Publishing, Cham, 608–626.
- Ramón-Constantí, A., Adan, A., Castilla Pascual, F., 2022. Thermal Point Clouds of Buildings: A review. *Energy and Buildings*, 274, 112425.
- Reality Capture Development Team, 2025. Realitycapture (version 2025.1) [computer software].
- Seitz, S., Curless, B., Diebel, J., Scharstein, D., Szeliski, R., 2006. A comparison and evaluation of multi-view stereo reconstruction algorithms. *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, 1, 519–528.
- Theodosiou, T., Papadopoulos, A., 2008. The impact of thermal bridges on the energy demand of buildings with double brick wall constructions. *Energy and Buildings*, 40(11), 2083–2089.
- Tomita, K., Chew, M., 2022. Infrared Thermography for Delamination Detection on Infrastructures and Buildings: A Review. *Buildings*, 12(2), 125.
- UAV4GEO Development Team, 2025. Thermal-tools.
- Yang, M. D., Su, T. C., Lin, H. Y., 2018. Fusion of infrared thermal image and visible image for 3d thermal model reconstruction using smartphone sensors.