



Research Article

A Case Study of Thermal Texture Transfer to Drone-Derived 3D Models with Neuralangelo and Instant-NGP: Using Photogrammetry, Neural Networks, and Data Fusion to Derive Thermal Twins

Ammar MEMARI, Linus HENKEL and Jerome AGATER

Jade Hochschule Wilhelmshaven/Oldenburg/Elsfleth, Germany

Correspondence should be addressed to: Jerome AGATER; jerome.agater@jade-hs.de

Received date:11 June 2025; Accepted date:12 February 2026; Published date: 31 March 2026

Copyright © 2026. Ammar MEMARI, Linus HENKEL and Jerome AGATER. Distributed under Creative Commons Attribution 4.0 International CC-BY 4.0

Abstract

Infrastructure diagnostics, energy efficiency analysis, and remote inspection can benefit from access to thermal 3D models of buildings and other installations. However, photogrammetrically reconstructing 3D models from thermal image data alone usually produces insufficient results due to a lack of texture, uniform regions, overexposure, and inadequate resolution. Our approach decomposes the problem into two parts: reconstruction of the 3D model and association of the thermal information with the 3D model. For reconstruction, we investigated the use of RGB imagery to photogrammetrically reconstruct the 3D model using COLMAP for traditional Structure-from-Motion (SfM), as well as the use of emerging neural reconstruction methods such as Neuralangelo. For subsequent thermal information transfer, we propose two fusion strategies: (1) Manual registration of image-pairs using an affine transform derived from user-selected correspondences to map thermal values to mesh vertex attributes; (2) Calibration and correction of lens distortion (radial/tangential, principal-point shift) and sensor misalignment to consistently project thermal imagery onto the reconstructed model. We use a real-world case study of a university service building at Jade Hochschule (Wilhelmshaven, Germany), with imagery captured by a DJI Mavic 3T drone, to validate the methodology, assessing the trade-offs in reconstruction fidelity, processing time, and thermal data accuracy. While our results show that vertex-based fusion is feasible, labor-intensive manual registration and information loss due to cropping and vertex sparsity limit accuracy and usability. Requiring only one-time parameter estimation, calibration and texture-based thermal projection yields more realistic, higher-fidelity thermal overlays, while enabling largely automated processing after the initial estimation. As a fast alternative for ad hoc thermal scene inspection, we evaluated Instant-NGP and 3D Gaussian splatting. While both can visualize thermal appearance, they suffer from reduced quality compared to our RGB-based versions and exhibit artifacts. Overall, our results indicate that an RGB-first reconstruction pipeline with camera-calibrated thermal texturing is the most viable path to usable thermal digital twins from drone data, while neural reconstructions may capture complex surfaces but remain computationally costly. Our modular, reproducible framework for thermal 3D digital twins is available as open source.

Keywords: Photogrammetry, Neuralangelo, Gaussian Splatting, thermal 3D reconstruction

Introduction

The creation of digital twins of real objects like buildings, machine parts, power plants, and other industrial systems provides many advantages for maintenance and construction (Grieves and Vickers, 2017). A digital 3D model can include more than just geometrical properties, i.e., material properties like stress, strength, elasticity, and so on. Additionally, information like cost or weight can be annotated on digital materials. These properties can be visualized through color textures. The visualization of all this information helps to diagnose or preemptively detect problems arising during operation or maintenance of the corresponding real-world object, e.g., damages in wind power plants (von Benzon and Chen, 2025), assessing quality and state of the construct, and exploring meta-information and structure. Furthermore, a digital twin allows the object to be inspected in its entirety securely and remotely, e.g., in a virtual reality setup, thereby improving cooperation with international experts when designing new systems or inspecting existing structures (Jeong et al., 2024).

Some digital twins are already created during the creation and design phase of, e.g., new buildings or manufacturing setups, but it is often necessary to create 3D models of already existing objects. This is usually done using photogrammetry. Photogrammetry is the method for reconstruction of environments or 3D objects as 2D maps or 3D models from photography and other optical sensor data using algorithms from computer vision. In contrast to this classical, algorithmic approach, machine learning implemented through neural nets has started to become a viable option, promising some advantages over the mathematical reconstruction. In 2023, NVIDIA research released Neuralangelo (Li et al., 2023)¹, one such a machine-learning based solution. Neuralangelo based reconstruction can offer more precision and also improved realism regarding complex structures, such as ruffles, edges or irregular surfaces. Furthermore, diverse lighting situations can be

modeled with improved results due to extensive training of lighting variants in the training data.

Since their widespread availability, commercial off-the-shelf drones have been used for capturing the imagery necessary for the photogrammetry. Meanwhile, the availability of multi spectral and infrared sensors for the creation of thermal imaging has opened up many possibilities, for example, the detection of cold spots in buildings (Autelitano et al., 2023). However, existing computer vision algorithms are not practically usable to create 3D models from thermal imaging, as crucial information is often missing (Paziewska and Rzonca, 2022): When a thermal image sensor is sampling an area that does not emit any or very little infrared radiation, these sampled real-world areas are represented uniformly as the same value, i.e., zero. At the same time, extremely hot areas are lacking details due to overexposure (Hou, Chen, Volk and Soibelman, 2022). Because of these shortcomings in the captured imagery, algorithms cannot find defining points in the pictures, which are necessary for the inference of the geometry. Furthermore, infrared cameras typically lack resolution when compared to RGB cameras, which would result in less detailed 3D models, even when the thermal source image would have enough variance to be usable.

In this paper, we explore the idea of combining the high resolution of RGB cameras with the information from thermal cameras into a digital twin via data fusion, constructing a 3D thermal digital twin using Neuralangelo (Li et al., 2023) and COLMAP (Schönberger and Frahm, 2016; Schönberger et al., 2016)² for the 3D reconstruction.

Related Work

The fusion of photogrammetric and thermal imaging data for 3D digital twin creation has become a vital tool in infrastructure diagnostics, energy audits, and remote inspections. Several studies have proposed

workflows to integrate RGB and thermal data into 3D models using UAV platforms. Parracho et al. (2023) developed a drone-based photogrammetric and thermographic workflow tailored for integration into (Building Information Modelling) and BEM (Building Energy Modelling) environments, highlighting the feasibility of combining RGB and infrared imagery for building energy performance analysis.

Dadras Javan and Savadkouhi (2019) proposed an approach where thermal and visible point clouds are separately generated and then fused via coarse and fine registration. This methodology significantly improves thermal mesh resolution by interpolating thermal data into high-resolution RGB point clouds. Similarly, Paziewska and Rzonca (2022) demonstrated the feasibility of integrating dense thermal and RGB-based point clouds for creating accurate models of heritage buildings using multiple drone systems.

For larger-scale or critical infrastructure monitoring, Gaspari et al. (2023) discussed a comprehensive integration of data from terrestrial laser scanning, UAV-based photogrammetry, and LiDAR for creating precise digital twins of dams, which can support structural analysis in natural hazard-prone environments.

Hou, Yao, Ramos, Lu, Volk and Soibelman (2022) proposed an optimized framework for fusing point cloud data from RGB and thermal images, highlighting the importance of pixel location, angle of capture, and drone flight parameters on mapping accuracy. Hou, Chen, Volk and Soibelman (2022) also investigated the effects of altitude, camera angle, and flight

path design on thermal-RGB fusion, concluding that optimal data quality depends on dynamic environmental and mission-specific parameters.

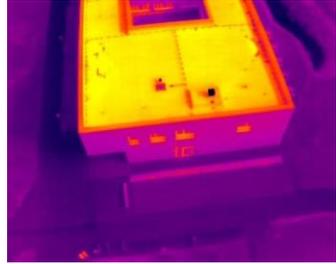
López et al. (2021) tackled the challenge of generating dense thermal point clouds by registering thermal images to RGB point clouds using the ECC algorithm, introducing occlusion-handling and dispersion-aware interpolation methods.

Finally, Autelitano et al. (2023) proposed an end-to-end pipeline for drone-based 3D thermal mapping of building envelopes, incorporating defect detection and comparison with idealized models for diagnostics and maintenance planning.

These prior works provide the foundation upon which our study builds: employing contemporary neural rendering techniques such as Neuralangelo and Gaussian Splatting to fuse RGB and thermal imagery in photogrammetry pipelines for digital twin enhancement.

Source Material and Capturing of Video Data

We selected the student service center (SSC) building of the Jade Hochschule located at the campus in Wilhelmshaven, Germany, as a modeled object for the development and the experiments, allowing easy access and quick iteration. The foundation stone of the building was laid on Dec. 7th 2020 and the official opening ceremony was on Sep. 22nd 2023, making the building relatively new.



(a) Thermal image capture



(b) RGB image capture

Figure 1: Examples of thermal and RGB images captured at the same time, illustrating the differences in sensor resolution as well as aspect ratio

For the capturing of image data, a DJI Mavic 3T (Thermal) drone was used, providing a gimbals-corrected sensor platform with a thermal camera (see Fig. 1a) as well as an RGB camera (see Fig. 1b). While RGB cameras measure light in the visible spectrum in the red-, green-, and blue-tones using photodiodes, common thermal cameras measure infrared radiation (typically in the mid- to far-infrared), by detecting the temperature changes resulting from absorption of the infrared

radiation in vanadium oxide components, and thereby indicating the thermal energy emitted from surfaces. In Photogrammetry, the properties of the camera sensors are important for the quality and precision of the reconstructed 3D models. The interesting sensor specifications are listed in Table 1. For the capturing of the image data, the drone was flown in a pattern around and above the building, gathering data from different angles and heights in the form of both pictures and video streams.

Table 1: Image sensor specification in the DJI Mavic 3T drone

Feature	Wide-Angle Camera	Thermal Camera
Sensor	1/2" CMOS, 48 MP	Uncooled VOx Microbolometer
35 mm Focal Length	24 mm	40 mm
Aperture	f/2.8	-
Resolution	4000 x 3000 Pixels	640 x 512 Pixels

Development

Two different approaches for the construction of the digital twin via data fusion were undertaken. In Approach 1, the thermal data are stored in the vertex properties of the reconstructed 3D model. In Approach 2, the thermal imagery is used as a texture for the 3D model. In the following section, the two approaches are presented as well as the results.

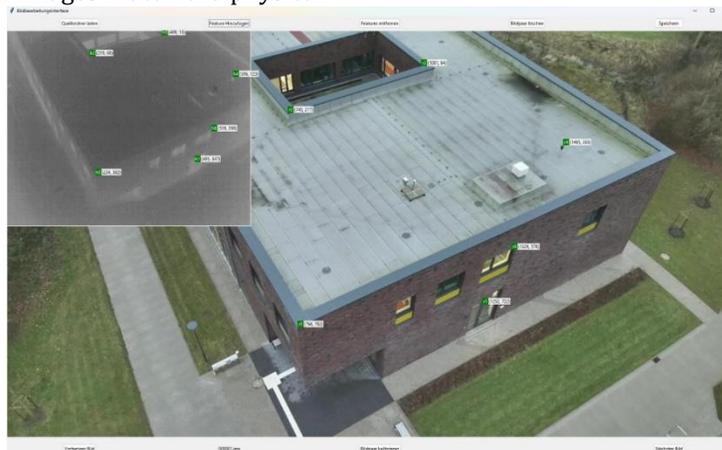
Approach 1: Thermal Information in Vertex Properties

For the implementation of the first approach, the images taken by the two different cameras need to be registered exactly. As the RGB and thermal cameras have different properties and are located at different positions/angles on the sensor head of the drone, coupling of the image material from each camera is necessary, the so-called registration. With this method, the different images resulting from different cameras are brought into the same context with the goal to align the properties of the

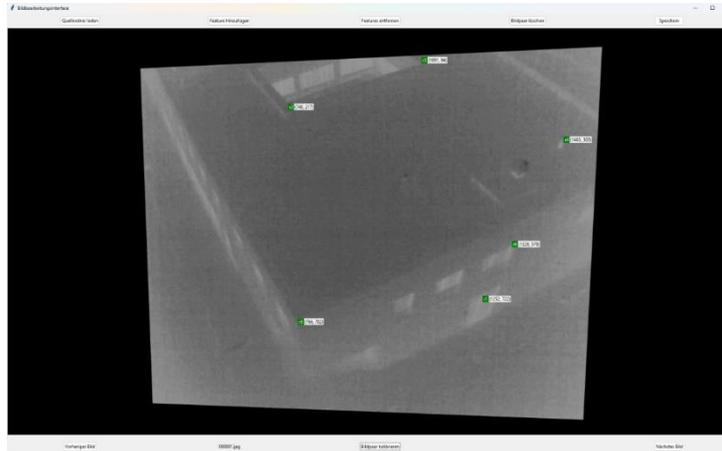
image pairs (Thermal,RGBt) taken at timestamp t . We used the video streams as the data source because they provided more source material. For the registration, it was necessary to extract individual frames from the streams. Unfortunately, it's not enough to change the same properties such that they have the same values, as the pictures have slightly different positions. To alleviate the problem, a GUI tool was developed using Python and Tkinter, allowing to register the images manually using only a few points. For this, prominent points are marked in both images, representing the same physical point in the scene. Heuristically, it is wise to select points for this, which are easy to be recognized and found in each member of an image pair. For optimal results, the chosen points should furthermore span the entire space of the image. Using these points as a reference, an affine transformation matrix is automatically derived (Dadras Javan and Savadkouhi, 2019), which is then used to scale and warp one image to match the other. In Fig. 2a, a screenshot of the GUI program illustrates how the points are selected in the image pairs; Fig. 2b illustrates how by clicking "Bildpaar kalibrieren" (calibrate image pair) the generated transformation matrix is applied to the thermal image. Fig. 2c shows the calibration result, where both images match the physical

points in the captured scene. Manual registration of all the required points is labor-intensive. To ease the process of marking the required points, an algorithm was implemented, which tracks the initially set points and moves the markers to the correct position, when the physical point is visible in the image. Once all images are registered in this way, all image pairs are warped once again, so that all the thermal images regain their rectangular shapes. The RGB image member of the pair is then cropped, so that the dimensions of both members are exactly equal. An example of a resulting image is presented in Fig. 3, which is now suitable for a reconstruction pipeline using Neuralangelo or traditional photogrammetry software alike.

Since computational resources were missing in the first phase, the program COLMAP based on the classical photogrammetry algorithms was used to reconstruct the 3D model instead of Neuralangelo. However, since Neuralangelo and COLMAP are oriented in the same coordinate space, the same experiment can be used for models reconstructed using Neuralangelo. Fig. 4 shows the result of a classical photogrammetrical reconstruction using the RGB-members of the registered image pairs.



(a) Marking the reference points for each image representing the same physical point



(b) Affine transformation matrix applied to the thermal image



(c) Overlay of the registered image pair

Figure 2: Developed GUI program for the registration of captured images



Figure 3: Overlay of the two members of an image pair after registration, calibration, transformation, and cropping.

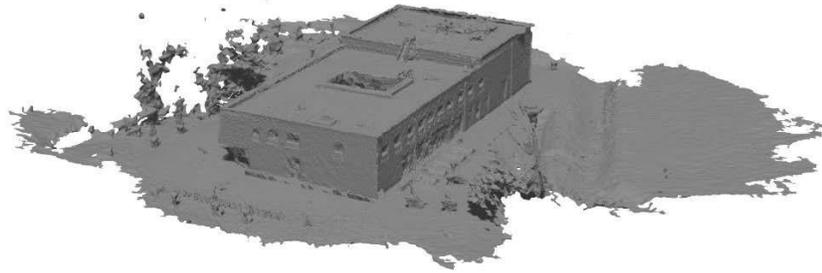


Figure 4: Reconstructed 3D model using the classical photogrammetry algorithms in COLMAP

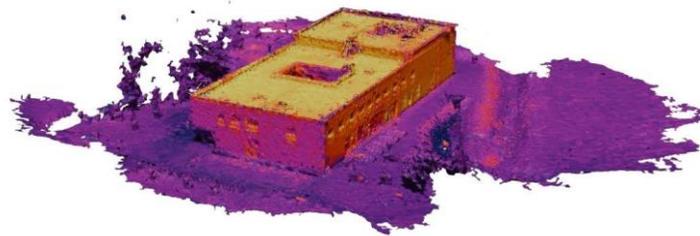


Figure 5: Reconstructed 3D model with changed vertex color information, now representing the thermal data and not the RGB data.

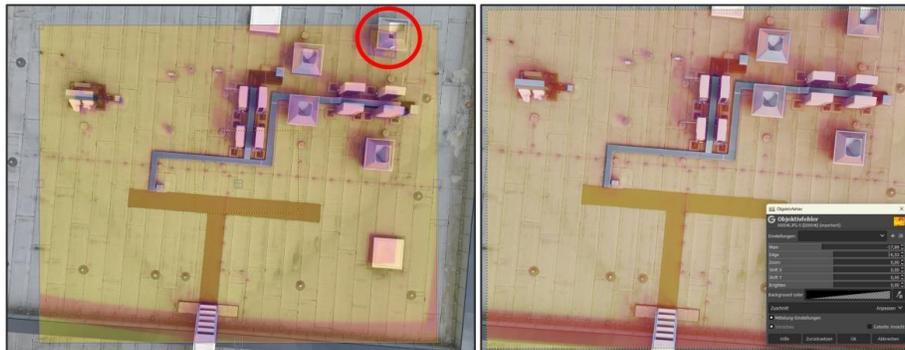


Figure 6: Overlay of the distorted thermal image over the RGB image using the GIMP filter Lens Distortion.

Once a 3D-Model has been reconstructed, the model is read by the program trimesh as well as auxiliary files, documenting in which image each point in the sparse point-cloud can be seen. For every point of the 3D model, the closest point in the parse point-cloud was identified, thereby associating each point in the 3D model with an image pair instance, in which this point is highly likely visible. Then, the 2-D coordinates of the 3D point in the 3D model of the individual image were calculated and the corresponding value from the thermal image for that coordinate extracted. This thermal information was then annotated at the individual vertex and the updated 3D model

saved. Fig. 5 shows the resulting 3D model with colors representing the thermal information extracted and applied to the vertices instead of the RGB information used for the construction of the 3D model.

Approach 2: Thermal Textures

For this approach, the registration of the image data was done using calibration of the camera properties. The parameters of the cameras can be determined using different methods. To validate the initial idea of correcting the camera induced distortion, we opted to use an image editing software, i.e., the GNU Image

Manipulation Program (GIMP), by applying the filter Lens Distortion (GIMP Documentation Team, 2025) manipulating the thermal image, such that there is geometrical congruence between the thermal image and the physical scene represented in the RGB image. Once the correct parameters of the filter have been determined, they can be re-applied to a checkerboard pattern using the same filter, distorting it. Using the open-source software OpenCV, distortion parameters can be readily inferred from the distortion in the checkerboard pattern.

Once this method had been validated via this manual prototyping, a GUI program was developed to ease the process of determining the correct parameters and remove the dependence of a third-party application. The program allows overlaying the two members from each image pair in the same coordinate system and then corrects it by applying scaling, distortion, and offset shifting (Hou, Yao,

Ramos, Lu, Volk and Soibelman, 2022). Once the correct parameters for these transformation operations have been determined using the GUI tool, they can be applied to all images in need of correction.

Usually, distortion is characterized by the parameters $k1$, $k2$, and $k3$, accounting for the radial distortion, and the parameters $p1$ and $p2$, accounting for the tangential distortion. These parameters can be changed and applied in the program directly using the interface elements on the right side of the application shown in Fig. 7. Streamlining another step of the image enhancement, a method for the correction of the displacement of the center of the image by cropping the thermal image was integrated, as the physical center of the RGB image does not match the physical center of the thermal image in the cameras of the drone used, DJI Mavic 3T.

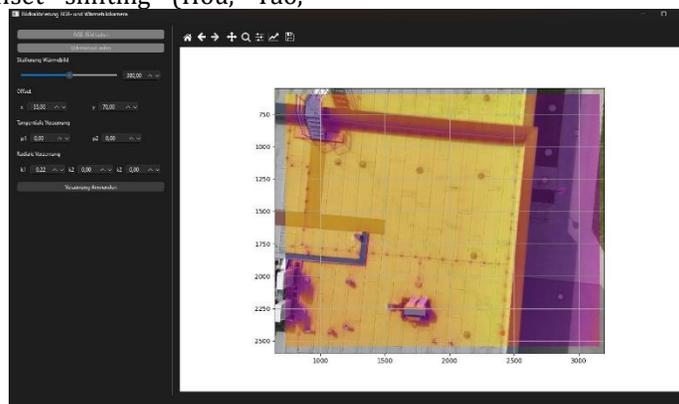


Figure 7: GUI tool showing a thermal image with the determined calibration parameters applied.

The center is displaced approx. 70 pixels to the left and 70 pixels down. To bring the two images into good congruence, the thermal image needs to be rotated by $k1 = -0,22$ and scaled by 380%. By applying this distortion, the pixels of the image are moved away from the image center. The application of these distortions diverts the actual focal length from the initially calculated focal length, requiring a manual intervention.

Results

Approach 1 is a viable solution to apply the thermal information to a reconstructed 3D model. However, the approach is only efficient and precise, when the image series do not need to be registered manually, and are instead simply singular image-pairs that can be registered through generic changes automatically. The manual registration introduces a fundamental inaccuracy in the

result. Furthermore, the cropping of the images can remove too much information from the RGB image, hindering the reconstruction. Approach 2 has proven as effective and delivers good performance. Even though the process of visual calibration is necessary, the

entire process can be automated, as the calibration of the camera needs to be done just once. Further reconstructions can be done automatically on the base of the determined values for the parameters.

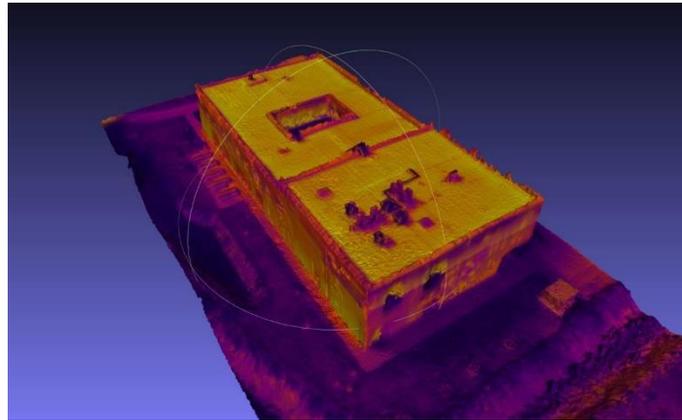
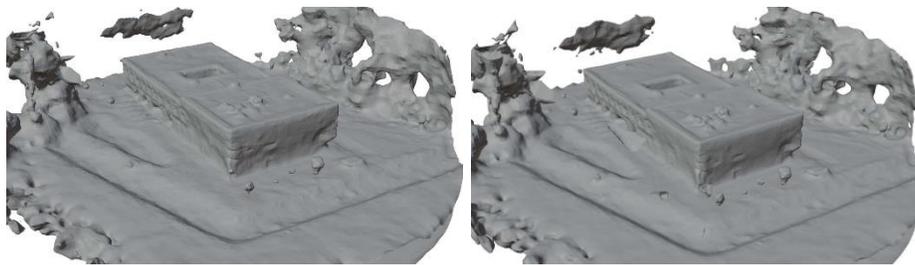


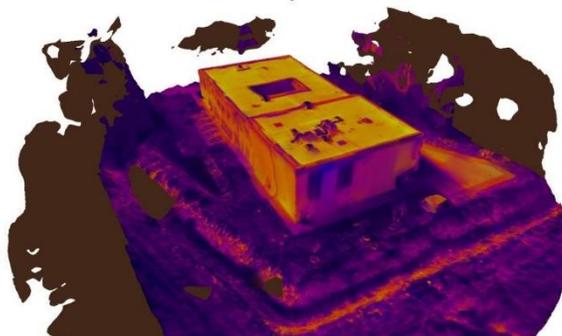
Figure 8: The reconstructed 3D model from COLMAP with projected thermal imaging textures.

Reconstructing the 3D model using Neuralangelo



(a) after 250,000 iterations

(b) after 500,000 iterations



(c) after 500,000 iterations with the projected thermal textures

Figure 9: Reconstruction of the SSC building using Neuralangelo

An instance of a high-performance system was rented to reconstruct the 3D model using a machine learning-based approach instead of classical photogrammetry. The system hosted a docker container with the specifications listed in Table 2. Neuralangelo was installed in the container and the data-set uploaded. The model was trained for 22 hours with approx. 3 iterations per second. The training was stopped after 250,000 iterations. This represents half of the training duration in the standard configuration of Neuralangelo, which is 500,000 iterations. The result after 250,000 iterations is shown in Fig. 9a. Then the training was continued for another 22 hours until the iteration count of the standard configuration was reached. The result after 500,000 iterations is shown in Fig. 9b.

Using Approach 2, the generated thermal textures were projected onto the

reconstructed 3D model using Meshlab. The resulting model is depicted in Fig. 9c.

Reconstructing the 3D model using Instant Neural Graphics Primitives

The Instant Neural Graphics Primitives (Instant NGP) are another method to enhance 3D scenes using neural networks, simulating the light rays in the scene in order to construct realistic views from different viewpoints (Müller et al., 2022). Instant NGP is another technology from NVIDIA, providing a speed-up compared to the original Neural Radiance Field (NeRF) approach (Müller et al., 2022; Mildenhall et al., 2021). While this method does not reconstruct a 3D-model directly, it can be used to infer an inspectable digital representation nevertheless.

Table 2: Properties of the Docker Container used for Neuralangelo

Storage	50 GB
RAM	50 GB
CPU Type	Virtual
Cores	12
GPU	Nvidia RTX 4090
VRAM	24 GB

Data preparation is analog to the pipeline used for Neuralangelo, which can be re-used for Instant NGP, taking about 10 minutes on the hardware listed in Table 2. The training itself happens in real time. After the training, the resulting scene can be viewed immediately. While using thermal imaging is not supported directly, the images used as input for Instant NGP can be replaced with the registered

thermal members of the image pairs after the initial training. Adaptation of the pipeline is only needed to generate a transforms.json file containing the properties of the thermal camera and the corresponding image pair members. Fig. 10 shows the interface of Instant NGP while viewing the re-constructed SSC building, with thermal imaging replacing the RGB images used for training.

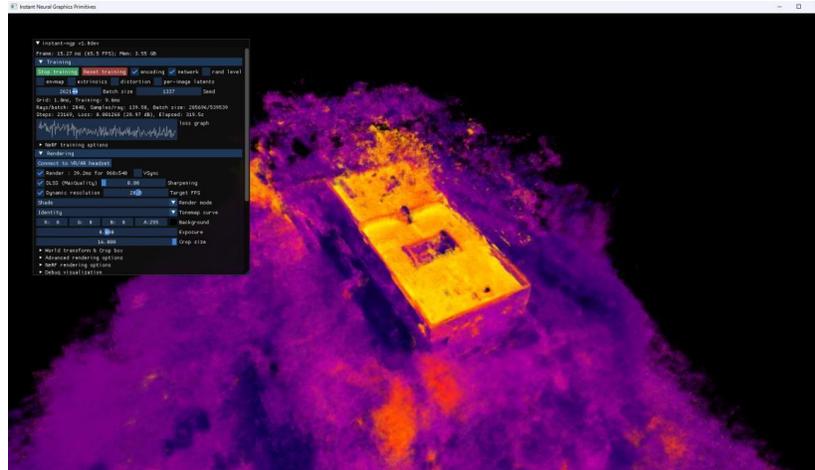


Figure 10: View of a scene reconstructed in Instant NGP using the thermal imaging after training with the RGB images.

The building itself is directly recognizable, and the color projection of the thermal images matches the physical scene. However, the resolution is too small and several artifacts are visible. Different reasons might contribute to these problems, e.g., problems in calibration of the virtual cameras or insufficient computational resources. In general, Instant NGP shows potential for the use case of digital inspection, especially because the data are viewable without a long waiting time for training and processing. Further research is necessary to study the usability aspects compared to a traditional 3D model, as well as for improving the quality of the resulting 3D scene.

Reconstructing the 3D model using 3D Gaussian Splatting

3D Gaussian Splatting (3D-GS) uses point-cloud representation for the rendering of 3D scenes (Kerbl et al., 2023). In this approach, the entire scene is modelled by discrete Gaussian distributions that surround the individual points in the point cloud. Every point is defined as a Gaussian Kernel with a specific spatial extent, color, and weighting. These Gaussian Kernels are projected along the camera view angle and combined using Soft Blending, resulting in a viewable image representation.

In contrast to Instant NGP, Gaussian splatting uses pre-generated point clouds, which are generated through Structure-from-Motion (Schönberger and Frahm, 2016; Schönberger et al., 2016) algorithms with external programs, e.g., COLMAP. This approach is efficient and does not need complex neural networks, reducing the processing time and improving the rendering. However, gaps in the point clouds can result in a high reduction in quality (Chen and Warren, 2020). The drone could only create images of the building from a steep angle, due to the surrounding buildings and objects like trees, resulting in such gaps in the point cloud. The resulting reduction in quality is visible, especially on the facade of the building.

For the experiments with 3D-GS, two datasets were created to be used as input parameters for the training of the models. The data sets contain the model export of the point cloud and the camera parameters, both from COLMAP, and the images used for the creation of the point cloud in COLMAP. For the training of the model, 24 GB VRAM is required. With increasing training time, more points are created, increasing the quality. For the training, an NVIDIA RTX 3090 with 24 GB VRAM was used, resulting in a training speed of approx. 11 iterations per second and a total training time of approx. 45 minutes. The

training was also tried on a less-powerful GPU, i.e., an NVIDIA RTX 3070 with 8 GB VRAM, resulting in a training speed of only 0.02 iterations per second.

The reconstruction of the RGB scene is photorealistic and rich in details. Especially

uniform surfaces like vegetation and smaller building details are represented very well. Similar to Instant NGP, however, thermal imaging results in some geometries not being detected correctly and resolution and texture offer less precision. Fig. 11 shows the resulting scene in the viewer of 3D-GS.

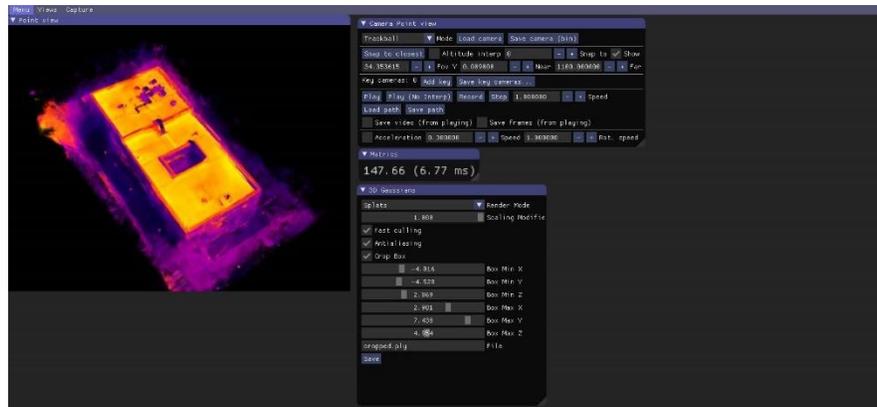


Figure 11: View of a scene reconstructed using 3D-GS with the thermal imaging replacing the RGB images

Conclusion

Interpretation of Results

The effort needed for a high-quality reconstruction using Neuralangelo is currently disproportionately high, especially when compared with the traditional photogrammetry applications already utilizing the GPU most of the heavy lifting processing tasks, considering the complex process, the necessary hardware resources and the long duration for the actual processing. Even after 44 hours of training utilizing the entire GPU, no result could be generated, that reached or exceeded the precision of other applications, potentially due to the available computational resources. The projection of the thermal data through changes in the vertex parameters (Approach 1) is only usable as an approximation, with more disadvantages than advantages. The projection of the thermal data using existing texturing algorithms into an external texture (Approach 2) has shown to be the better method, offering higher flexibility in

changing the applications and algorithms. Furthermore, the results of Approach 2 are more realistic and the data open for modification.

Outlook and Lessons Learned

As with all data pipelines, the quality of the input data is crucial for the quality of the results. For drone flights, the best choice is using waypoint-based flight involving tridimensional planning to capture the object from all angles. The process of the assimilation of the image pairs could be improved, removing inaccuracies and reducing the manual effort necessary. By using quaternion vectors, the virtual camera could be adjusted, removing the need to crop images and thereby losing the information from these areas.

Summary

Due to the increased use of drones providing RGB- and thermal imaging, the projection of thermal imaging on a model constructed using

photogrammetry is increasingly relevant. The big challenge is combining the different sensors and data in a unified geometrical context, allowing the use of machine learning and traditional photogrammetry software to construct the 3D model from RGB data and applying the thermal imaging.

The reconstruction of surfaces of 3D models with Neuralangelo can produce good results and can be a good choice for use cases with a need for high precision. Especially for buildings with surfaces hard to reconstruct, Neuralangelo has a distinct advantage. However, since the demand on the hardware is quite high, traditional reconstruction algorithms and methods, which are already using the GPU for most of the work, offer a lot more efficient solution and are therefore better suited for most use cases.

The first approach of using the RGB-imaging for the 3D reconstruction of the depicted scenes and then projecting the thermal imaging on the model has shown as a practical solution. The Approach 2, modifying the NVM file and texturing the 3D model, is the better choice, mainly due to the following reasons:

1. the details of the thermal images are not lost when reducing the vertexes,
2. the result looks more natural to human viewers,
3. the handling of the data is easier.

When the use case does not require an actual 3D model, but the 3D representation of the covered scene is enough, Instant NGP or Gaussian Splatting are an interesting alternative. The experiments detailed in this paper have shown to be viable options to visualize a 3D scene with thermal information. Better results could be obtained by further optimizing the input data.

Acknowledgment

The work of the author Jérôme Agater was funded by the JadeProf project.

Notes

¹ Neuralangelo on GitHub, <https://github.com/NVlabs/neuralangelo>, last accessed 2025-05-06

² COLMAP on GitHub, <https://github.com/colmap/colmap>, last accessed 2025-05-06

³ Friedrich-Paffrath-Straße 101, 26384 Wilhelmshaven, Germany, latitude: 53.5472, longitude: 8.0890

References

- Autelitano, K., Bolognini, M., De Angelis, E., Fagiano, L. and Scaioni, M. (2023), 'On the use of drones for vision-based digitalization, diagnostics and energy efficiency assessment of buildings', *Science and Technology for the Built Environment* 29(10), 985–997.
- Chen, Y. and Warren, W. (2020), '3D Fusion of Infrared Images with Dense RGB Reconstruction from Multiple Views – with Application to Fire-fighting Robots'. *arXiv preprint*. [Online], [Retrieved Jun 2025], <https://arxiv.org/abs/2007.14606>
- Dadras Javan, F. and Savadkouhi, M. (2019), 'Thermal 3D Models Enhancement Based on Integration with Visible Imagery', *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences XLII-4/W18*, 263–269.
- Gaspari, F., Barbieri, F., Ioli, F., Pinto, L. and Valgoi, P. (2023), 'Integration of 3D surveying approaches for critical infrastructure digital twins in natural hazard-prone scenarios'. [Online], [Retrieved Jun 2025], <https://meetingorganizer.copernicus.org/EGU23/EGU23-7773.html>
- GIMP Documentation Team (2025), 'Gnu image manipulation program 3.0 user manual, chapter 5.5 lens distortion', [Online], [Retrieved Jun 2025], <https://docs.gimp.org/3.0/en/gimp-filter-lens-distortion.html>.
- Grieves, M. and Vickers, J. (2017), *Digital Twin: Mitigating Unpredictable*,

- Undesirable Emergent Behavior in Complex Systems, in F.-J. Kahlen, S. Flumerfelt and A. Alves, eds, 'Transdisciplinary Perspectives on Complex Systems', Springer International Publishing, Cham, 85–113.
- Hou, Y., Chen, M., Volk, R. and Soibelman, L. (2022), 'Investigation on performance of RGB point cloud and thermal information data fusion for 3D building thermal map modeling using aerial images under different experimental conditions', *Journal of Building Engineering* 45, 103380.
 - Hou, Y., Yao, S., Ramos, L., Lu, W., Volk, R. and Soibelman, L. (2022), Optimization of Point Clouds' RGB and Thermal Information Data Fusion for Buildings' 3D Thermal Mapping, in 'Computing in Civil Engineering 2021', American Society of Civil Engineers, Orlando, Florida, 663–670.
 - Jeong, I., Jang, Y., Park, J., Nietfeld, K., Pakbaz, A. and Moon, H. (2024), Remote Virtual Inspection of Infrastructure in the Metaverse, in 'Construction Research Congress 2024', American Society of Civil Engineers, Des Moines, Iowa, 319–328.
 - Kerbl, B., Kopanas, G., Leimkühler, T. and Drettakis, G. (2023), '3d gaussian splatting for real-time radiance field rendering', *ACM Transactions on Graphics* 42(4).
 - Li, Z., Muller, T., Evans, A., Taylor, R. H., Unberath, M., Liu, M.-Y. and Lin, C.-H. (2023), Neuralangelo: High-Fidelity Neural Surface Reconstruction, in '2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)', IEEE Computer Society, Los Alamitos, CA, USA, 8456–8465.
 - López, A., Jurado, J. M., Ogayar, C. J. and 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.
 - Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R. and Ng, R. (2021), 'Nerf: representing scenes as neural radiance fields for view synthesis', *Commun. ACM* 65(1), 99–106.
 - Müller, T., Evans, A., Schied, C. and Keller, A. (2022), 'Instant neural graphics primitives with a multiresolution hash encoding', *ACM Trans. Graph.* 41(4).
 - Parracho, D. F. R., Poças Martins, J. and Barreira, E. (2023), A Workflow for Photogrammetric and Thermographic Surveys of Buildings with Drones, in M. D. L. N. González García, F. Rodrigues and J. Santos Baptista, eds, 'New Advances in Building Information Modeling and Engineering Management', Springer Nature Switzerland, Cham, 77–95.
 - Paziewska, J. and Rzonca, A. (2022), 'Integration of Thermal and RGB Data Obtained by Means of a Drone for Interdisciplinary Inventory', *Energies* 15(14), 4971. [Online], [Retrieved Jun 2025], <https://www.mdpi.com/1996-1073/15/14/4971>
 - Schönberger, J. L. and Frahm, J.-M. (2016), Structure-from-Motion Revisited, in 'Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)', 4104–4113
 - Schönberger, J. L., Zheng, E., Pollefeys, M. and Frahm, J.-M. (2016), Pixelwise View Selection for Unstructured Multi-View Stereo, in 'European Conference on Computer Vision (ECCV)', 501–518.
 - von Benzon, H.-H. and Chen, X. (2025), 'Mapping damages from inspection images to 3d digital twins of large-scale structures', *Engineering Reports* 7(1), e12837