Autonomous Drone Navigation for Structural Inspection Using 3D Gaussian Splatting and Optimization-based Path Planning

Hui Zuo[1], Thasanka Kandage[1], Hongyang Xu[1,2], Nima Shirzad-Ghaleroudkhani[1], Qipei Mei[1], José Fiestas[3] and Luis A. Bedriñana[4]

[1]Department of Civil and Environmental Engineering, University of Alberta, Edmonton, AB, Canada
[2]School of Energy and Mining Engineering, China University of Mining & Technology-Beijing, Beijing, China
[3]Department of Computer Science, Universidad de Ingenieria y Tecnologia - UTEC, Lima, Perú
[4]Department of Civil and Environmental Engineering, Universidad de Ingenieria y Tecnologia - UTEC, Lima, Perú
{polo.huizuo, thasanka, xhongyan, shirzadg, qipei.mei}@ualberta.ca, {jfiestas, lbedrinana}@utec.edu.pe

Abstract –

Structural inspections increasingly rely on unmanned aerial vehicles (UAVs) to improve access and safety, however, UAV-based bridge inspections often suffer from inconsistent data coverage, pilot-dependent flight paths, and limited repeatability across inspection campaigns. This study presents a fully autonomous drone inspection framework that integrates 3D Gaussian Splatting (3DGS) with optimization-based path planning using a 3D A* algorithm. The framework reconstructs a photorealistic 3D model from video data to identify Regions of Interest and automatically generates collision-free flight trajectories that balance surface coverage, flight distance, and inspection time. Field test on the Low Level Bridge in Edmonton, Canada demonstrates that a high-fidelity 3DGS preview can be produced in around five minutes, enabling rapid on-site rendering. The optimized helical path achieves 91.4% coverage with a 45.8% reduction in mission time, offering a scalable, model-driven solution for safe and efficient structural health monitoring.

Keywords –

Autonomous Drone Navigation; Structural Inspection; 3D Gaussian Splatting; 3D A* Optimization; Path Planning; Georeferenced 3D Reconstruction; Structural Health Monitoring.

1 Introduction

Structural inspection plays a critical role in ensuring the safety, reliability, and longevity of civil infrastructure [1]. Many civil engineering structures, such as bridges, are now operating way beyond their originally intended design service lives [2]. As these structures age, regular inspection and maintenance become vital to detect damage, evaluate deterioration, and prevent catastrophic failures [3]. Traditional inspection methods, however, rely heavily on manual access such as scaffolding, under-bridge units, or rope-supported inspections, which expose personnel to significant safety risks and often require prolonged traffic closures [4]. Moreover, the quality of data acquired through manual methods is inconsistent, depending on the inspector's line of sight, environmental conditions, and subjective judgment [5]. These limitations have prompted growing interest in autonomous and data-driven inspection technologies.

Unmanned aerial vehicles (UAVs), or drones, have emerged as a practical solution for visual inspection because of their flexibility, accessibility, and cost-effectiveness [6]. Modern UAVs equipped with high-resolution cameras can quickly collect detailed imagery from angles that are otherwise difficult or unsafe to reach [7]. As a result, drone-based inspections are increasingly used for bridges, towers, and building facades [8]. Yet despite their advantages, most current UAV workflows remain semi-manual [9]. Operators must manually define flight paths, ensure sufficient image overlap, and visually identify defects post-flight, often resulting in redundant trajectories and inefficient battery usage. These factors make repeatable inspections difficult and introduce variations between missions.

To address these challenges, recent research has focused on integrating photogrammetry and 3D reconstruction techniques into inspection workflows. Structure-from-Motion (SfM) [10] and Multi-View Stereo (MVS) [11] approaches can generate detailed point clouds and textured meshes from overlapping imagery. These reconstructions allow engineers to measure geometric properties [12], identify defects [13], and visualize damage locations [14] in three dimensions. However, they often require substantial computation time and lack the capability for real-time visualization in the field [15]. More advanced neural rendering methods, such as Neural Radiance Fields (NeRF) [16], have demonstrated photorealistic reconstruction quality but remain computationally intensive [17], limiting their use for practical inspection applications where real-time feedback is crucial.

3D Gaussian Splatting (3DGS) [18] has recently emerged as a promising alternative to neural radiance fields for efficient, high-fidelity scene reconstruction. Instead of modeling volumetric radiance through deep networks, 3DGS represents scenes as a collection of Gaussian primitives that encode position, orientation, color, and opacity. This representation is optimized through differentiable rendering, enabling real-time visualization with significantly lower computational cost.

While 3DGS provides an efficient means to reconstruct and visualize structures, it alone does not automate the inspection process. In practice, the effectiveness of UAV-based inspection also depends on flight-path design. These predefined paths often result in redundant image overlap and inefficient flight trajectories.

To overcome these limitations, this study introduces a fully autonomous drone inspection framework that integrates 3D Gaussian Splatting with optimization-based path planning. The proposed framework contributes three major advancements to UAV-based inspection. First, it demonstrates that 3DGS can deliver high-fidelity, georeferenced 3D models directly from video capture within minutes, enabling near real-time field deployment. Second, it establishes a robust link between reconstructed models and autonomous path planning, allowing the drone to execute optimized, repeatable flight missions derived from previously captured data. Third, it validates the practicality of the framework through both simulation and field testing, confirming its potential to enhance safety, efficiency, and data consistency in structural health monitoring.

2 Methodology

The proposed framework establishes a fully autonomous pipeline for structural inspection that integrates data acquisition, 3D Gaussian Splatting, and optimization-based path planning. As illustrated in Figure 1, the process begins with an initial drone flight around the target structure to capture a high-resolution video sequence. The collected footage is processed using 3D Gaussian Splatting to reconstruct a dense and photorealistic 3D representation of the structure, enabling identification of defected zones, defined as Regions of Interest (ROI). This digital model serves as the foundation for automated mission planning: the reconstruction is converted into a LAS file via COLMAP to facilitate geometric analysis and path generation.

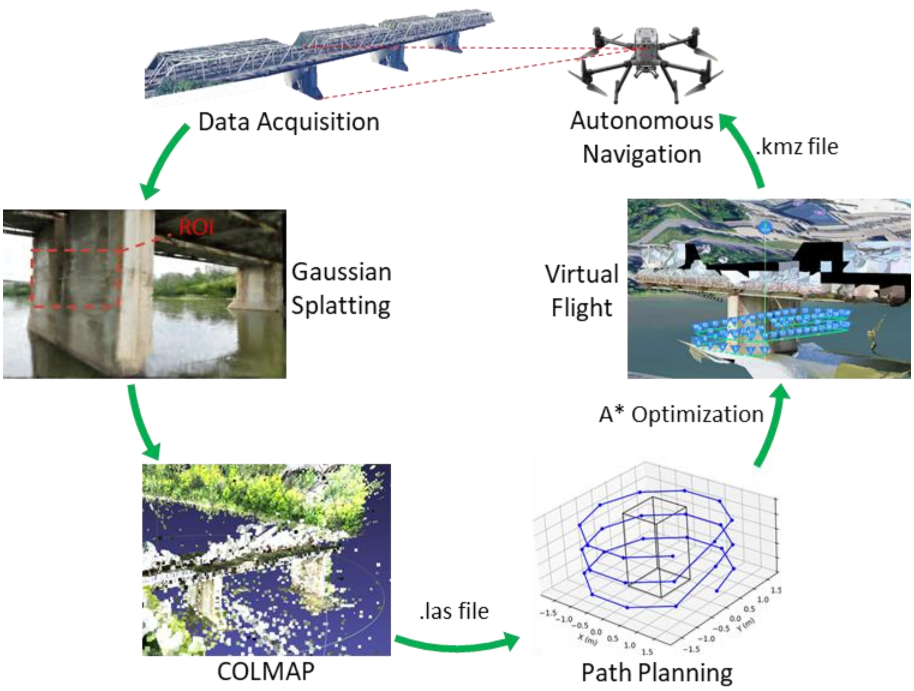

Figure 1. Overview of the proposed autonomous drone inspection framework

Candidate trajectories are generated based on the structure's geometry and inspection requirements, and subsequently optimized using a 3D A* algorithm to balance multiple objectives, including surface coverage, total flight distance, and total flight time. The optimized path is exported as a KMZ file and uploaded to the drone system, closing the loop between data acquisition, 3D reconstruction, path planning and mission execution. Once the initial model and corresponding ROI have been established, subsequent inspections can seamlessly reuse the optimized routes, enabling repeatable and consistent operations with minimal manual oversight.

2.1 3D Gaussian Splatting for Reconstruction

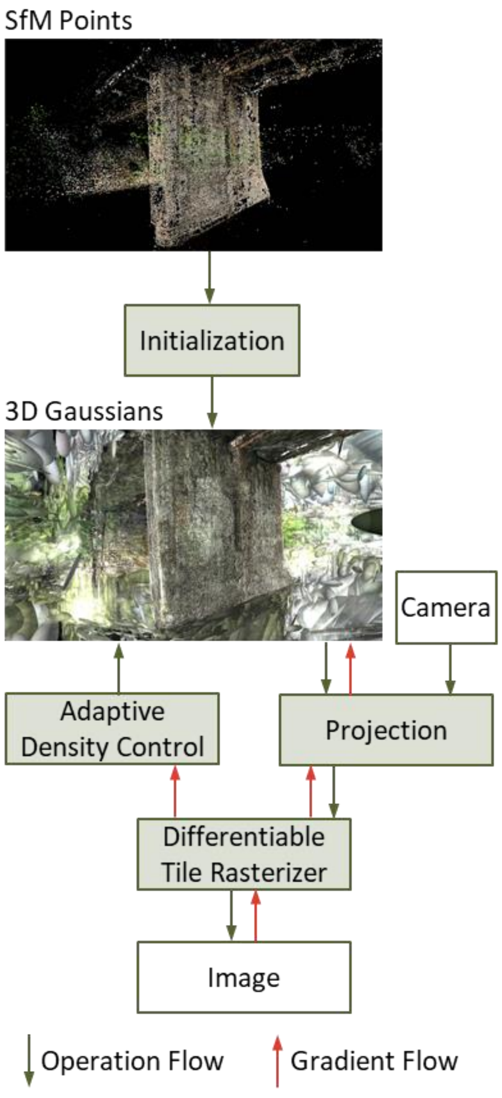

Figure 2. 3D Gaussian splatting workflow

As illustrated in Figure 2, the 3D Gaussian Splatting process begins by initializing a set of Gaussian primitives from the sparse point cloud generated through Structure-from-Motion. Each Gaussian represents a small volumetric element characterized by position, orientation, color, and opacity. These primitives collectively encode both the geometry and appearance of the scene in a continuous and differentiable form.

During optimization, the Gaussian parameters are iteratively refined through a differentiable rendering process. Each Gaussian is projected onto the image plane, and the rendered result is compared with the input images to minimize photometric error. The adaptive density control module dynamically adjusts the number and distribution of Gaussians, concentrating higher densities in geometrically complex or textured regions while maintaining efficiency in uniform areas.

The differentiable tile rasterizer allows gradients to flow back from the image domain to the 3D Gaussians, enabling end-to-end optimization of color, opacity, and spatial parameters. This tile-based renderer provides a substantial acceleration compared with traditional neural radiance-field methods while preserving comparable visual fidelity. Once optimization converges, the resulting 3D Gaussian model offers a photorealistic and spatially coherent reconstruction of the target structure, suitable for subsequent inspection analyses and path-planning.

All 3D Gaussian Splatting training and rendering are performed on a workstation running Ubuntu 22.04 and equipped with an NVIDIA A6000 GPU with 48 GB memory. The framework is also compatible with consumer-grade GPUs, with proportionally longer training times.

2.2 Path Planning and Optimization

Following the 3D Gaussian Splatting reconstruction, ROIs are manually identified by the inspector from the rendered 3D model. This manual delineation is the only human-dependent step in the pipeline, once defined, the ROI coordinates are stored and reused for subsequent inspections, ensuring repeatability. The reconstructed scene is then processed through COLMAP to generate a LAS-format point cloud that preserves geometric and spatial information. From this model, the boundary of each ROI is extracted and exported as a CSV file containing latitude, longitude, and altitude data. These ROI coordinates define the spatial targets for the subsequent flight path optimization process.

The path planning algorithm operates in three main stages: ROI sampling, trajectory computation, and route optimization. Initially, the structure's surface is represented as a convex hull derived from the LAS point cloud, and candidate waypoints are sampled directly from the ROI file. Each waypoint is offset from the surface by a configurable standoff distance to ensure safe flight clearance.

For waypoint-to-waypoint navigation, a 3D A* algorithm is applied to compute collision-free trajectories within a voxelized representation of the surrounding space. The voxel grid incorporates structural boundaries and clearance zones, ensuring that all generated paths maintain a safe distance from the structure's surface. Each potential path segment is evaluated for visibility using a line-of-sight check, and if direct visibility is obstructed, the A* planner computes an alternative route around the structure.

Once the shortest feasible paths between all waypoint pairs are determined, a nearest-neighbor Traveling Salesman Problem (TSP) heuristic determines the visiting sequence to minimize total flight distance. The optimized tour is visualized in 3D and exported as a set of georeferenced waypoints. Each waypoint includes a corresponding yaw angle calculated relative to the structure's centroid, ensuring that the drone's camera maintains consistent orientation toward the inspection target. Together, these steps establish a fully automated pipeline from ROI detection to flight execution.

To verify inspection completeness, a coverage validation is performed using a simulated camera model with defined focal length and field of view. A 3D frustum is projected at each waypoint, and the structure surface is sampled to estimate visible areas, enabling quantitative assessment of inspection completeness. The final output is a flight route file in KMZ format that can be directly uploaded to the drone's remote controller.

3 Field Test and Results

Figure 3. Field data collection setup on the pier

To evaluate the performance of the proposed autonomous inspection framework, a field test is conducted on the Low Level Bridge in Edmonton, Canada, as illustrated in Figure 3. The bridge, opened in 1900, spans the North Saskatchewan River and carries vehicular traffic between downtown and the southern districts of the city. It features reinforced concrete piers supporting a Pratt truss superstructure. Owing to its age and complex geometry, the structure presents an ideal test case for validating the proposed 3D reconstruction and inspection workflow.

A DJI Matrice 350 RTK equipped with a Zenmuse H30T camera is deployed to capture high-resolution visual data from the first concrete pier on the southern approach of the bridge. The drone is operated to execute a single orbital flight around the pier, recording continuous 4K video at 3840 × 2160 resolution to ensure full surface coverage. Individual frames are extracted from the recorded video at a rate of one frame per second, producing a sequence of high-quality still images. The corresponding GPS information from the flight log is embedded into each extracted image, enabling the generation of a georeferenced 3D model. These geotagged images are then used as input for 3D Gaussian Splatting reconstruction, forming the basis for a detailed and spatially accurate digital representation of the bridge pier.

3.1 3DGS Results

The 3D Gaussian Splatting is performed using 160 image frames with embedded GPS coordinates. The dataset is divided into an 8:2 ratio for training and testing. Initial camera poses are estimated through COLMAP's SfM pipeline, whose output format is compatible with the 3DGS training framework. The model is trained for 30,000 iterations on an A6000 GPU, completing in 23 minutes and 33 seconds, and converges to 988,970 Gaussian primitives that encode the pier's geometry and surface appearance.

Quantitative evaluation indicates good reconstruction fidelity, with a Structural Similarity Index of 0.787, Peak Signal-to-Noise Ratio of 19.240 dB, and Learned Perceptual Image Patch Similarity of 0.328 on the testing set. These metrics are consistent with complex outdoor scenes where illumination changes, water reflections, and metallic members challenge photometric consistency. The resulting model preserves the GPS coordinate, allowing seamless integration with subsequent path planning procedure.


(a) Ground truth (b) 7k training (05 min 07 s) (c) 30k training (23 min 33 s)

Figure 4. Visual progression of 3DGS

A practical benefit for on-site workflows is early visual readiness. As shown in Figure 4, by about 7,000 iterations, which is roughly 5 minutes of training on our hardware, the renderer already produces photorealistic previews sufficient to localize and mark ROI. Final refinement to 30,000 iterations further enhances sharpness and fine-detail consistency but is not essential for ROI selection or mission planning.

3.2 Path Planning and Optimization Results

In order to evaluate the performance of the proposed path planning algorithm, the optimized flight trajectory is compared with the default geometric route generated by DJI Flighthub 2 through virtual simulations. Both executed at a nominal speed of 0.5 m/s. As summarized in Table 1, the default route requires 1433 s to complete a 999.0 m trajectory, achieving full coverage but with extensive overlap between adjacent waypoints. In contrast, the proposed algorithm produces a helical trajectory derived from the 3D model, maintaining a 1.5 m standoff distance from the pier surface. As illustrated in Figure 5, the optimized route achieves 91.4% surface coverage with a total flight distance of 387.9 m and a duration of 776 s. This corresponds to a 45.8% reduction in flight time relative to the default route, while preserving comprehensive inspection coverage and structural safety margins.

Table 1. Comparison of Flight Path Performance

Mission Type	Default Route	Optimized Route
Flight Distance (m)	999.0	387.9
Flight Time (s)	1433	776
Coverage (%)	100.0	91.4

The camera frustum is modelled based on the Zenmuse H30T specifications (focal length 6.72 mm, diagonal field of view 82.1°), ensuring consistent visibility during coverage estimation. The slight reduction in overall coverage arises from intentional clearance at the upper and lower boundaries of the structure, where proximity to the bridge superstructure and water surface could increase collision risk.

Subsequent field verification on the Low Level Bridge confirms that the algorithm-generated trajectory performs reliably in practice. During the field test, the drone maintains stable flight along the planned path under calm weather conditions, with the DJI M350 RTK's onboard obstacle avoidance providing an additional safety layer. Note that the current field test is conducted under favorable environmental conditions, and further evaluation under adverse conditions such as high winds, variable lighting, and GPS-degraded environments is warranted. Nevertheless, once the initial 3D model and ROI are established, the same optimized mission can be repeatedly executed in future inspections without manual replanning, offering a consistent and efficient solution for ongoing structural monitoring.

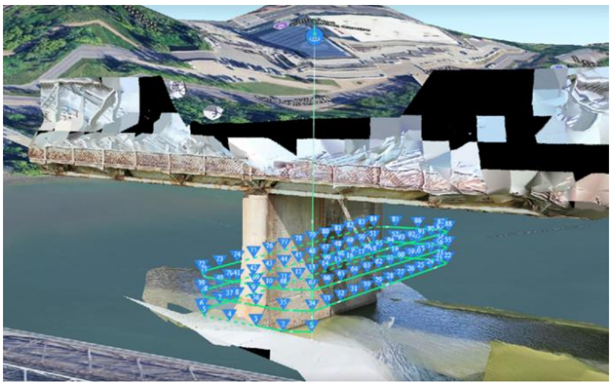

Figure 5. Optimized flight path

4 Conclusions

This study presents a fully autonomous drone inspection framework that integrates 3D Gaussian Splatting with optimization-based path planning for efficient structural assessment. The framework establishes a closed loop from visual data acquisition to automated mission generation, enabling reliable and repeatable inspections with minimal manual intervention. The combination of photorealistic 3D reconstruction and algorithmic path optimization demonstrates strong potential for improving inspection efficiency, reducing operational risk, and standardizing data collection in structural health monitoring. The key findings of this study can be summarized as follows:

1. The proposed pipeline successfully reconstructs georeferenced 3D models of bridge structures from video capture using 3D Gaussian Splatting. A photorealistic preview with high visual fidelity can be generated in around five minutes of training, allowing inspectors to localize defects and define ROI directly in the field.
2. The optimization-based path planning module, built upon a 3D A* algorithm and TSP sequencing, effectively generates collision-free flight routes with controllable standoff distances.
3. Virtual flight simulations in DJI Flighthub 2 demonstrate a 45.8 % reduction in total mission time compared with the default geometric route, while maintaining 91.4 % coverage under safe clearance constraints.
4. Field verification on the Low Level Bridge confirms the reliability and repeatability of the optimized mission, validating its applicability for ongoing inspections using the same reconstructed 3D model and planned path.

Despite the promising results, several limitations remain. The current validation is based on a single bridge pier, and further testing across diverse structure types (e.g., towers, tanks, and long-span bridges) and environmental conditions is needed to confirm the generalizability of the framework. The current framework also relies on a manually piloted initial flight for data acquisition, which could be replaced by an autonomous exploration mode in future development. In addition, the altitude measurements in this study are derived from the drone's onboard barometer, which can introduce offset errors due to air-pressure fluctuations. Incorporating a GNSS Mobile Station or RTK-based positioning system would enable precise altitude calibration and centimeter-level georeferencing accuracy. Future work will also focus on benchmarking the planning module against additional baselines (e.g., a zigzag scan pattern and a random waypoint sampling strategy) and evaluating the sensitivity of key planning parameters such as standoff distance and waypoint density, as well as integrating real-time defect detection with adaptive flight replanning, potentially transforming UAV-based inspection into a fully autonomous and adaptive system.

Acknowledgements

The authors gratefully acknowledge the support from the 2024 University of Alberta, Faculty of Engineering - Universidad de Ingeniería y Tecnología (UTEC) Seed Grant Program, as well as the NSERC Discovery Grant (Grant No. RGPIN-2022-04160).

References

[1] X. Chen *et al.*, "Innovative life-cycle inspection strategy of civil infrastructure: Smartphone-based public participation," *Structural Control and Health Monitoring*, vol. 2023, no. 1, p. 8715784, 2023, doi: 10.1155/2023/8715784.

[2] D. Reagan, A. Sabato, and C. Niezrecki, "Feasibility of using digital image correlation for unmanned aerial vehicle structural health monitoring of bridges," *Structural Health Monitoring*, vol. 17, no. 5, pp. 1056–1072, Sep. 2018, doi: 10.1177/1475921717735326.

[3] S. S. Khedmatgozar Dolati, N. Caluk, A. Mehrabi, and S. S. Khedmatgozar Dolati, "Non-destructive testing applications for steel bridges," *Applied Sciences*, vol. 11, no. 20, p. 9757, Jan. 2021, doi: 10.3390/app11209757.

[4] T. Whitley *et al.*, "Design of a small unmanned aircraft system for bridge inspections," *Sensors*, vol. 20, no. 18, p. 5358, Jan. 2020, doi: 10.3390/s20185358.

[5] R. Xiong, P. Liu, and P. Tang, "Human reliability analysis and prediction for visual inspection in bridge maintenance," pp. 254–262, May 2022, doi: 10.1061/9780784483893.032.

[6] P. Nooralishahi *et al.*, "Drone-based non-destructive inspection of industrial sites: A review and case studies," *Drones*, vol. 5, no. 4, p. 106, Dec. 2021, doi: 10.3390/drones5040106.

[7] S. Harwin and A. Lucieer, "Assessing the accuracy of georeferenced point clouds produced via multi-view stereopsis from unmanned aerial vehicle (UAV) imagery," *Remote Sensing*, vol. 4, no. 6, pp. 1573–1599, Jun. 2012, doi: 10.3390/rs4061573.

[8] J. F. Falorca and J. C. G. Lanzinha, "Facade inspections with drones–theoretical analysis and exploratory tests," *International Journal of Building Pathology and Adaptation*, vol. 39, no. 2, pp. 235–258, May 2020, doi: 10.1108/IJBPA-07-2019-0063.

[9] S. Tsachouridis, F. Pavloudakis, C. Sachpazis, and V. Tsioukas, "Monitoring slope stability: A comprehensive review of UAV applications in open-pit mining," *Land*, vol. 14, no. 6, p. 1193, Jun. 2025, doi: 10.3390/land14061193.

[10] J. L. Schonberger and J.-M. Frahm, "Structure-from-motion revisited," presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 4104–4113. Accessed: Oct. 12, 2025. [Online]. Available: https://www.cv-foundation.org/openaccess/content_cvpr_2016/html/Schonberger_Structure-From-Motion_Revisited_CVPR_2016_paper.html

[11] D. Cernea, "OpenMVS: Multi-View Stereo Reconstruction Library." [Online]. Available: https://github.com/cdcseacave/openMVS
[12] A. Koutsoudis, B. Vidmar, G. Ioannakis, F. Arnaoutoglou, G. Pavlidis, and C. Chamzas, "Multi-image 3D reconstruction data evaluation," *Journal of Cultural Heritage*, vol. 15, no. 1, pp. 73–79, Jan. 2014, doi: 10.1016/j.culher.2012.12.003.
[13] Z. Xu, R. Kang, and R. Lu, "3D reconstruction and measurement of surface defects in prefabricated elements using point clouds," *Journal of Computing in Civil Engineering*, vol. 34, no. 5, p. 04020033, Sep. 2020, doi: 10.1061/(ASCE)CP.1943-5487.0000920.
[14] T. Yamane, P. Chun, and R. Honda, "Detecting and localising damage based on image recognition and structure from motion, and reflecting it in a 3D bridge model," *Structure and Infrastructure Engineering*, vol. 20, no. 4, pp. 594–606, Apr. 2024, doi: 10.1080/15732479.2022.2131845.
[15] H. Zuo, T. Sun, H. Xie, X. Ma, N. Shirzad-Ghaleroudkhani, and Q. Mei, "Transformer-based multi-resolution fast 3D reconstruction for structural damage detection," in *ISARC. Proceedings of the International Symposium on Automation and Robotics in Construction*, Montreal, Canada: IAARC Publications, 2025, pp. 988–995.
[16] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "NeRF: Representing scenes as neural radiance fields for view synthesis," *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, Dec. 2021, doi: 10.1145/3503250.
[17] H. M. Wijaya, H. H. Handayani, and T. Hariyanto, "Gaussian splatting for rapid and precise 3D modeling in mapping application: Literature review," in *2024 IEEE Asia-Pacific Conference on Geoscience, Electronics and Remote Sensing Technology (AGERS)*, Dec. 2024, pp. 111–117. doi: 10.1109/AGERS65212.2024.10932937.
[18] B. Kerbl, G. Kopanas, T. Leimkuehler, and G. Drettakis, "3D gaussian splatting for real-time radiance field rendering," *ACM Trans. Graph.*, vol. 42, no. 4, pp. 1–14, Aug. 2023, doi: 10.1145/3592433.