Carinthia II

Nature Tech

Part 3

Published since 1811
215th Year of Carinthia

Volume 2

Issue 2

2025 Verlag des Naturwissenschaftlichen Vereins für Kärnten

Predictive maintenance in infrastructure: Utilizing 3D point clouds for efficient damage detection

Christina Petschnigg, Alexander Pamler, Kazim Onur Arisan, Jan Morten Loës, Torsten Ullrich

ABSTRACT

Growing urbanization is driving the demand for infrastructure such as parking lots, roads, and bicycle lanes. While green spaces and trees are often integrated into these development projects to mitigate negative climate impacts, they can cause root-related damage that poses safety risks and requires costly monitoring. Public road networks are typically inspected with advanced but expensive surveillance vehicles that are too costly for private applications, leaving private infrastructure such as parking lots, private roads, and storage areas without comparable solutions. Thus, this paper presents a methodology for detecting and classifying damage areas in 3D point clouds of parking lots, distinguishing root-related damage from construction joints using a combination of deep learning and classical statistics. The approach is evaluated on data from Vienna International Airport and validated against manually labeled ground truth data. Results show that accurate localization and classification of damage is feasible using only a single laser scanner, providing a cost-effective alternative to conventional monitoring. Moreover, the method facilitates predictive maintenance by automatically detecting damage and enabling integration into Building Information Modeling software.

Vorausschauende Instandhaltung von Infrastruktur: Effiziente Schadensdetektion mithilfe von 3D-Punktwolken

ZUSAMMENFASSUNG

Die wachsende Urbanisierung erhöht die Nachfrage nach Infrastruktur wie Parkplätzen, Straßen und Radwegen. Grünflächen und Bäume werden häufig in diese Entwicklungsprojekte integriert, um negative Effekte des Klimawandels abzumildern. Allerdings können diese wurzelbedingte Schäden verursachen, die Sicherheitsrisiken bergen und eine kostenintensive Überwachung erfordern. Das öffentliche Straßennetz wird in der Regel mit fortschrittlichen, jedoch teuren Überwachungsfahrzeugen inspiziert, die für private Infrastruktur wie Parkplätze, Privatstraßen und Lagerflächen zu kostspielig sind. Daher wird in diesem Paper eine Methodik zur Erkennung und Klassifizierung von Schäden in 3D Punktwolken von Parkplätzen vorgestellt, bei der wurzelbedingte Schäden mithilfe einer Kombination aus Deep Learning und klassischer Statistik von Baufugen unterschieden werden. Der Ansatz wird anhand von Daten des Flughafens Wien evaluiert und mit manuell annotierten Ground-Truth-Daten validiert. Die Ergebnisse zeigen, dass eine präzise Lokalisierung und Klassifizierung von Schäden mit einem einzelnen Laserscanner möglich ist, was eine kostengünstige Alternative zu herkömmlichen Monitoringverfahren darstellt. Darüber hinaus unterstützt die vorgeschlagene Methode Predictive-Maintenance-Maßnahmen, indem Schäden automatisch erkannt und in Building Information Modeling-Software integriert werden können.

KEYWORDS

- 3D point cloud
- tree roots
- predictive maintenance
- damage detection

INTRODUCTION

Since 2008, more than half of the global population lives in urban areas, a proportion that is expected to increase to 68% by 2050 [1]. While urbanization drives economic and social progress, it also increases soil sealing, impacting the environment and society [2]. The expansion of impervious surfaces – such as roads, buildings, and parking lots – at the expense of green and open spaces is a global trend with profound environmental consequences. It exacerbates the effects of heavy rainfall by contributing to flooding and also intensifies the urban heat island effect. A common strategy to mitigate these effects is incorporating green spaces, especially trees, into sealed areas, as they promote cooling [3], water evaporation, air purification [4], and hydrological protection [5]. While essential for addressing climate change-related challenges, plant and root growth can damage infrastructure, causing cracks or bulges in surfaces. Consequently, systematic infrastructure monitoring and inspection are essential for early damage detection and predictive maintenance, enabling cost-effective repairs, reducing accident risks, extending infrastructure lifespan, and enhancing public safety. However, these monitoring and inspection efforts demand substantial time and financial resources from infrastructure operators.

A significant portion of infrastructure, including parking lots, roads, and bike lanes, is publicly owned and is monitored using expensive surveillance vehicles. While effective for inspections, their high cost makes them inaccessible to private sector organizations, businesses, and operators of commercial or private infrastructure who are responsible for managing and maintaining private roads, parking lots, and other outdoor spaces. Consequently, inspecting infrastructure in these settings presents a more significant challenge. Recent studies on low-cost road and infrastructure inspections mainly utilize two-dimensional (2D) images, as they are cost-effective and easily captured with readily available devices such as smartphones [6]. However, 2D imagery lacks depth information, which limits its ability to capture detailed geometric characteristics of surface conditions. In contrast, three-dimensional (3D) point clouds provide rich geometric and color information, enabling a more comprehensive analysis, which not only improves the detection and classification of surface anomalies but also allows for a more accurate assessment of their extent. Consequently, 3D data support more informed decision-making by enabling the prioritization of maintenance tasks based on the geometric severity of the detected damage. For instance, pavement damage inspection using 3D point clouds is described in [7] and [8].

Against this background, our paper presents a cost-effective methodology for detecting and classifying damage to parking infrastructure using 3D point cloud data, acquired using a terrestrial laser scanner instead of a full surveillance vehicle. The proposed approach combines deep learning with classical statistical techniques to distinguish between damage caused by root growth and construction joints. Specifically, a PointNet-based [9] neural network is applied to segment the parking lot into relevant and non-relevant points. Damage is then identified at the relevant points through the geometric approximation of the ground surface and the classification and directional analysis of damage. Since the damage detection takes place within a georeferenced 3D point cloud, the results can be integrated directly into Building Information Modeling (BIM) systems, enabling efficient planning and management of maintenance activities. In summary, our contributions are:

- **Framework**: We propose a common framework for detecting, documenting, and classifying infrastructure damage using georeferenced 3D point cloud data. The framework segments relevant surface points and distinguishes damage caused by root growth from construction joints. This offers a more cost-efficient solution compared to the use of surveillance vehicles and delivers greater accuracy than conventional 2D image-based monitoring methods.
- **Experiment:** We assess the accuracy and completeness of damage detection by comparing the model's output with manually annotated ground truth data, using a real-world parking lot at Vienna International Airport as a test site.

MATERIALS AND METHODS

Data Collection and Pre-Processing

The study was conducted in Parking Lot D at Vienna International Airport, which spans an area of approximately 13,000 m² and is depicted in Figure 1. Data acquisition was performed using the terrestrial laser scanner VZ-400i (Riegl International GmbH, Vienna, Austria). For accurate georeferencing, existing control points in the surrounding area were additionally captured. In total, 98 scan positions were recorded and subsequently processed using RiSCAN PRO software (RIEGL Deutschland Vertriebsgesellschaft mbH, Gilching bei München, Germany) . The raw point cloud data were filtered to exclude points with reflectivity values below -25 dB or above 5 dB, as well as those with a deviation value exceeding 15. Low reflectivity often results from weak signal returns on dark or absorbent surfaces, whereas high reflectivity may indicate sensor artifacts. Elevated deviation values typically reflect measurement noise. Filtering these outliers improves geometric accuracy for subsequent analysis and reliability. The individual scan positions were then registered relative to each other to align and merge them into a unified coordinate system. The resulting point cloud was colorized using RGB values extracted from images taken by the scanner's internal camera. Movable objects, such as vehicles and pedestrians, were removed by the software, while residual noise was manually eliminated using the Terrain Filter tool.

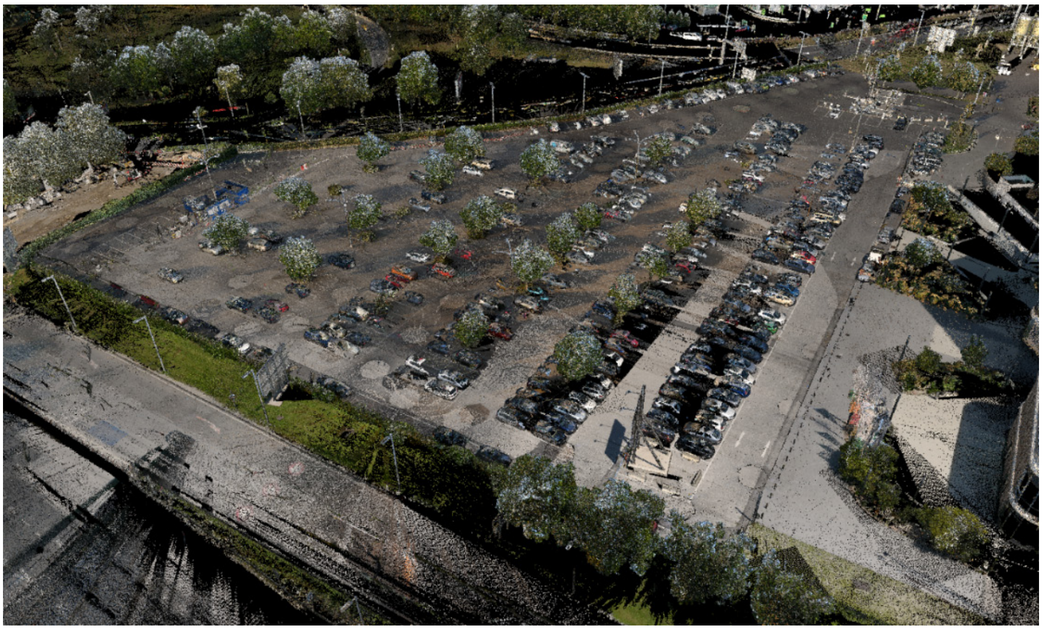

Fig. 1

Figure 1: Parking Lot D at Vienna International Airport. Source: own figure.

Abbildung 1: *Parkplatz D am Flughafen Wien. Quelle: eigene Abbildung.*

The consolidated point cloud with a resolution of 1 cm was imported into the software program CloudCompare [10] for further refinement. Statistical outlier removal was conducted to further eliminate noise points that compromise the integrity of the dataset. Subsequently, the point cloud was manually annotated to assign each point to one of the classes: ground, tree, low vegetation, and car. Figure 2 presents the point cloud of the test site at Vienna International Airport in two formats: one representing the cleaned data and the other depicting the class labels.

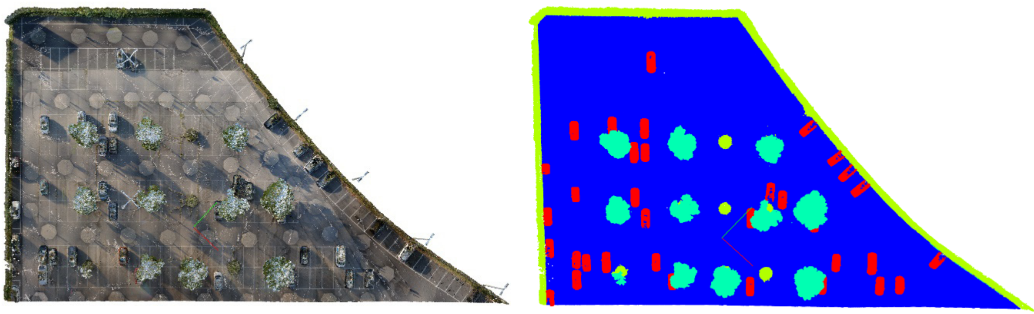

Fig. 2

Figure 2: Cleaned (left) and labeled (right) point cloud of the test site at the Vienna International Airport. Source: own figure.

Abbildung 2: *Gesäuberte (links) und klassifizierte (rechts) Punktwolke des Testgeländes am Flughafen Wien. Quelle: eigene Abbildung.*

Model Implementation and Training

Our methodology comprises three primary steps: 1) segmentation of the point cloud to identify relevant features (such as ground, vegetation, and trees) for further analysis, while excluding irrelevant points (such as vehicles); 2) tree documentation and damage detection, which leverages geometric features such as positive and negative protrusions to identify surface anomalies; and 3) classification and documentation of the identified areas of damage, distinguishing between those caused by root growth and construction joints. Future research will extend this analysis to include additional damage types, such as those induced by heavy usage or freeze-thaw cycles, including potholes and subsidence.

Segmentation

In the first step, we employ deep learning techniques for automatic point cloud segmentation. Specifically, we adopt a PointNet-based approach [9], which processes raw point cloud data directly [11], eliminating the need for intermediate steps like voxelization [12] or 2D image conversion [13]. To improve computational efficiency, the point cloud is first down-sampled using uniform grid sampling and then split into training (≈70%), validation (≈20%), and testing (≈10%) sets. The PointNet-based model is implemented using Python's PyTorch framework [14] and is trained for 100 epochs on a GeForce RTX 3090 GPU (NVIDIA Corporation, Santa Clara, CA, USA) for parameter optimization. Post-inference inaccuracies, which primarily result from limited training data and labeling noise, are reduced using a multi-step refinement of car and vegetation detection. Initially, the inference results are projected onto a 2D bird's-eye view of the scene, where the respective objects are delineated through density-based clustering. The boundaries of these regions are subsequently refined, and the corrected segmentations are reprojected into the 3D point cloud to enhance spatial accuracy. Subsequently, all detected vehicles are removed from the dataset, as they are not relevant for assessing ground surface damage.

Damage Detection

In the second step, tree positions are extracted from the segmented "tree" point cloud. Individual trees are identified and isolated using density-based clustering. The position of each tree is defined as the ground-level coordinate located at the center of the tree stem. This georeferenced information can be directly integrated into a BIM system, which is increasingly important in a smart facility planning context, where accurate, up-to-date models of both built and natural assets are essential. Such integration helps in assessing risks related to root intrusion, visibility obstructions, and vulnerability to storms.

The segmented point clouds corresponding to the "ground" and "vegetation" classes are further analyzed. The ground point clouds are divided into disjoint square regions for ground plane approximation. Within each subdivision, a fourth-degree polynomial regression model is applied to capture nonlinearities in the data, such as surface irregularities. To ensure robustness against outliers, a RANSAC (Random Sample Consensus) regressor is employed. This process generates a fourth-degree polynomial function for each square segment of the ground, which models its complex and uneven surface. The resulting surfaces are then connected using bicubic interpolation, providing a smooth and continuous approximation of the entire ground surface. This approximation facilitates the identification of damage, with root-induced defects manifesting as positive protrusions and construction joints corresponding to either positive or negative protrusions. These deviations are identified by comparing the point cloud data to the approximated ground surface.

Damage Classification

In the third step, we differentiate whether the identified areas of damage are caused by tree root growth or construction. This involves a detailed analysis of the damage structure. Initially, individual damage areas are delineated from one another. The point cloud is then converted into a graph representation using a k-d tree, which is a data structure optimized for efficient nearest-neighbor searches and partitioning of points in k-dimensional spaces. Each connected component in the point cloud is initially assigned to a single area of damage. However, due to the inherent incompleteness of point clouds, collinear areas of damage are merged, as missing points often cause discontinuities that actually represent one single contiguous area of damage. After this merging process, areas of damage are classified based on their geometric properties. Root-induced damage tends to be curved and aligns with the direction of the recorded tree positions, distinguishing them from linear construction joints, which are typically straight. Other damage types, such as potholes, were not present in the test data.

RESULTS

As outlined in the previous section, the proposed infrastructure monitoring framework consists of three main stages: segmentation of the point cloud, documentation of trees and detection of damage, and subsequent damage classification. To ensure a comprehensive evaluation of the framework's effectiveness, each stage is assessed independently. As previously mentioned, the evaluation is conducted on Parking Lot D at Vienna International Airport.

Segmentation Results

Training and optimizing the PointNet-based neural network over 100 epochs on the training set—which includes approximately two rows of the parking area—yielded a point-wise accuracy of 97.16% on the training data and 95.32% on the validation data. Evaluation of the unseen test area, which corresponds to a short parking row, resulted in a point-wise accuracy of 92.61%. Table 1 summarizes the corresponding per-class accuracies and the mean Intersection over Union (mIoU) for each class. Qualitative visualization of the segmentation results is provided in Figure 3. The model demonstrated a strong ability to distinguish between the predefined classes, with particularly high separability observed for the "ground" and "tree" classes. However, due to the limited extent of the training data and some labeling noise in the "vegetation" and "ground" classes, the segmentation exhibited inaccuracies at class boundaries, especially between the classes "car", "ground", and "vegetation". These boundary ambiguities were subsequently mitigated by incorporating geometric features of the point cloud, as described in the previous section, resulting in the refined segmentation illustrated in Figure 4 and yielding an overall classification accuracy of 97.74% on the test set.

Tab. 1

	After Segmentation		After refinement	
Class	**Class Accuracy**	**IoU**	**Class Accuracy**	**IoU**
Ground	0.944	0.925	0.999	0.978
Tree	0.943	0.921	0.989	0.963
Vegetation	0.726	0.644	0.739	0.704
Car	0.824	0.602	0.917	0.916

Table 1: Class accuracy and class IoU after segmentation and refinement.

Tabelle 1: *Class Accuracies und Class IoU nach der Segmentierung und Verfeinerung.*

Damage Detection

As tree detection is performed through straightforward density-based clustering on a clearly distinguishable point cloud class, it was considered a trivial task in this context

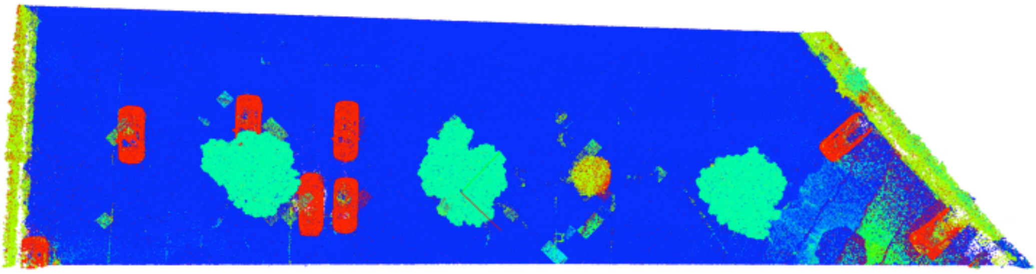

Fig. 3

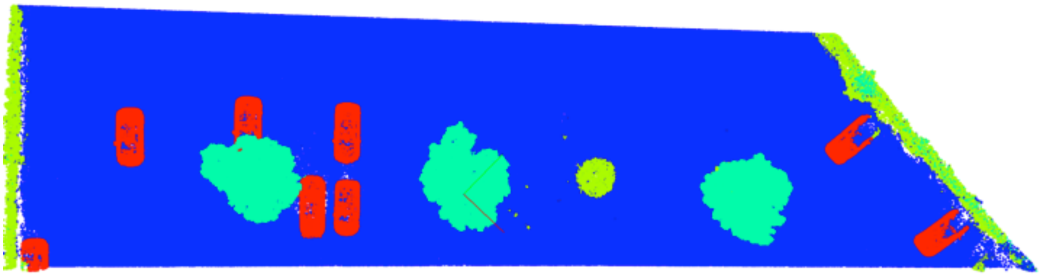

Fig.4

Figure 3: Qualitative segmentation results on the test data set. Source: own figure.

Abbildung 3: *Qualitative Segmentierungsergebnisse auf dem Testdatensatz. Quelle: eigene Abbildung.*

Figure 4: Qualitative segmentation results on the test data set after refinement. Source: own figure.

Abbildung 4: *Qualitative Segmentierungsergebnisse auf dem Testdatensatz nach der Verfeinerung. Quelle: eigene Abbildung.*

and therefore not subject to quantitative evaluation.

The previously described method, which detects damage by analyzing deviations between the point cloud and the approximated ground surface, effectively localized damaged regions, achieving a mutual overlap of 95.9% with manually annotated ground truth data. The overlap was calculated by measuring nearest-neighbor distances between points in both point clouds, with overlap defined as points within a distance threshold set at 1% of the bounding box diagonal. The mutual overlap ratio was computed by averaging the directional overlaps from each cloud to the other. Figure 5 presents a comparison between manually labeled damage and examples of damage identified by the system. Overlapping points are shown in green, while red and blue points represent those unique to the ground truth and system output, respectively. Although the overall overlap was high, our method occasionally failed to detect fine construction joints characterized by subtle negative surface deviations. In contrast, it demonstrated high sensitivity to positive protrusions, detecting root-induced damage with high accuracy.

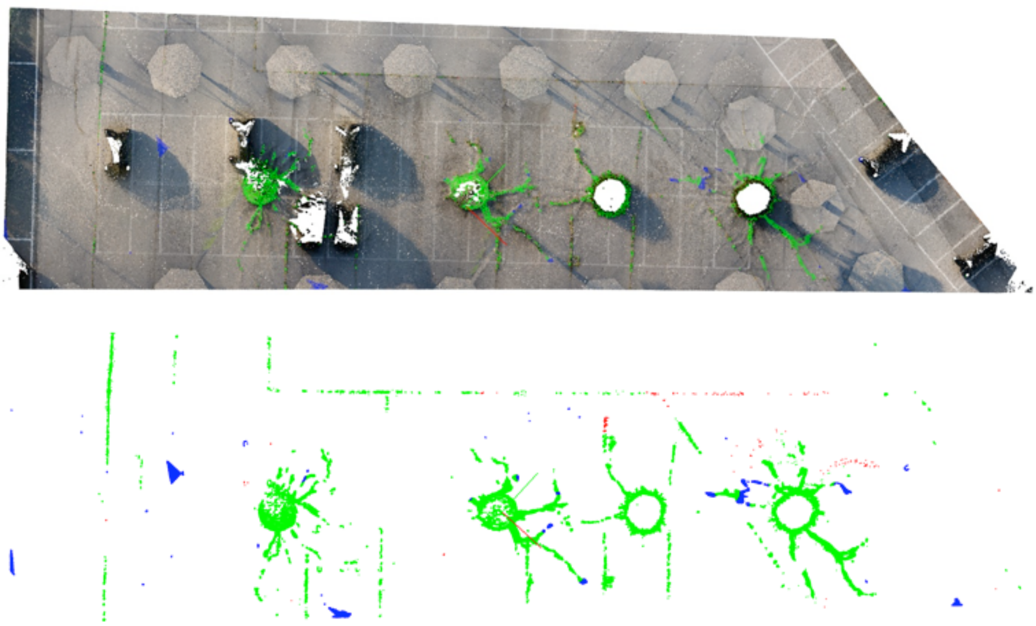

Fig. 5

Figure 5: Comparison of ground truth and detected damage. Overlapping points are shown in green, while red and blue indicate points unique to the ground truth and system output, respectively. Full point cloud (top), damage only (bottom). Source: own figure.

Abbildung 5: *Vergleich von manuell erkannten und automatisch detektierten Schadstellen. Überlappende Punkte sind in Grün dargestellt, Rot und Blau zeigen jeweils Punkte, die ausschließlich in der Ground Truth bzw. in der Auswertung enthalten sind. Gesamte Punktwolke (oben), nur Schäden (unten). Quelle: eigene Abbildung.*

Damage Classification

Subsequently, individual areas of damage were classified based on their geometric characteristics to differentiate between curved, root-induced damage and the typically linear patterns associated with construction-related damage. Figure 6 illustrates the qualitative results of the classified damage, where root-induced damage is visualized in red and construction joints in green. The overall classification achieved an accuracy of 95.15%. The accuracy was determined by establishing spatial correspondence between the predicted and ground truth point clouds. This was achieved by identifying the nearest neighbor in the ground truth for each point in the predicted dataset using a k-d tree. Since the point clouds differed in sampling density and exhibited spatial deviations between predictions and annotations, they also differed in the number of points, making one-to-one correspondence infeasible. Consequently, the nearest-neighbor approach offered a robust solution for establishing spatial correspondence and evaluating geometric overlap between the two datasets.

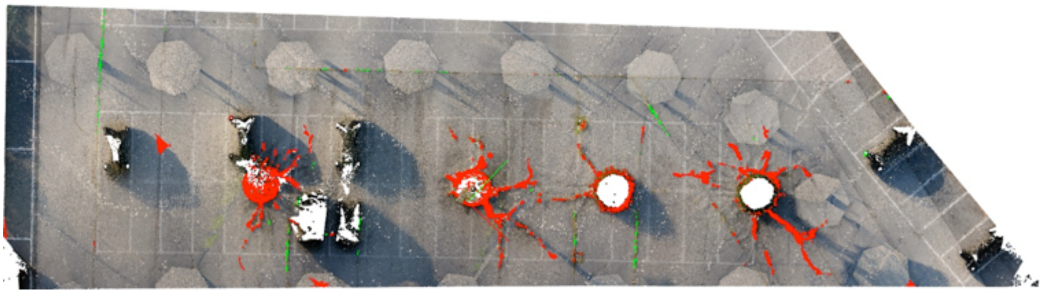

Fig. 6

Figure 6: Qualitative results of damage classification. Root-induced damage is depicted in red, while construction joints are shown in green. Source: own figure.

Abbildung 6: *Qualitative Ergebnisse der Schadensklassifikation. Wurzelschäden sind in Rot dargestellt, während Baufugen in Grün dargestellt werden. Quelle: eigene Abbildung.*

DISCUSSION AND CONCLUSION

This study presents a low-cost and scalable methodology for detecting and classifying surface-level infrastructure damage using 3D point cloud data. Our framework represents a significant advancement in making infrastructure monitoring more accessible. By eliminating the need for expensive surveillance vehicles and relying instead on terrestrial laser scanners and point cloud processing, it enables smaller organizations and private entities to conduct reliable, automated inspections. This has the potential to improve predictive and preventive maintenance efforts as well as safety, especially in private parking lots, campuses, and commercial zones. Moreover, preliminary estimates indicate that surveying with terrestrial laser scanning is significantly more cost-efficient per square meter than vehicle-based systems, while simultaneously providing richer geometric information than low-cost 2D imaging methods.

The segmentation results demonstrated high accuracy (>92%) even in unseen test areas, affirming the effectiveness of the neural network approach in extracting relevant infrastructure features from raw point cloud data. However, challenges remain—particularly at the boundaries of segmented classes—where label ambiguity and point cloud noise can impair performance. The subsequent damage detection and classification processes effectively differentiated between root-induced and construction-related damage, achieving high detection overlap (95.9%) and classification accuracy (95.15%).

The results presented in the paper demonstrate that the proposed method achieves high accuracy in damage detection in the test area, particularly in detecting root-induced damage. Another notable aspect is the integration of the methodology with BIM systems, which supports more effective predictive maintenance planning. By classifying and documenting both root-induced surface damage and construction joints, the approach facilitates the development of a holistic digital representation of a facility, including maintenance-relevant information for outdoor infrastructure. Overall, the findings suggest that the proposed methodology constitutes a promising and cost-effective solution for infrastructure monitoring.

Nonetheless, there remains potential for improvement, especially in refining the segmentation of complex surfaces, improving the detection of subtle construction-related damage, and extending the framework to accommodate a broader range of damage types. Despite the high overall accuracy, the segmentation process exhibited limitations, particularly in handling ambiguities at class boundaries—most notably between the classes “car”, “ground”, and “vegetation”. These issues are expected to be mitigated in future work through an expanded and more diverse dataset, which will provide the deep learning model with increased variability and improved generalization. Additional improvements will focus on extending the framework’s capabilities to include further damage types, such as those resulting from heavy mechanical wear or environmental effects like freeze-thaw cycles—an effort that will also benefit from a larger dataset. Furthermore, the damage detection component showed reduced sensitivity to subtle construction joints, including fine cracks and slight surface depressions, which the current approach struggles to detect. To address this, future developments will incorporate additional data types, such as image data, into the analysis pipeline. The inclusion of visual information is expected to enhance the detection and validation of fine cracks. Additionally, the current method of data collection—terrestrial laser scanning—is time-consuming. The substitution of terrestrial scanners with drone-based surveys offers a promising alternative to reduce acquisition time and increase coverage; however, it may introduce additional sources of error like motion artifacts or positional inaccuracies that must be considered. Moreover, the application of semi-supervised learning techniques can help alleviate the annotation workload and enable better generalization across diverse surface types and environmental conditions.

While the Vienna International Airport parking lot provided a controlled environment for testing, the applicability of the proposed method to more safety-critical areas of airport infrastructure—particularly runways and taxiways, where surface integrity is essential for safe aircraft operations—remains a subject for future investigation. Beyond the aviation domain, the methodology can be adapted for a wide range of outdoor infrastructure applications, including bridges, building exteriors, and industrial sites. This versatility, in combination with the system’s cost-efficiency, positions the framework as a valuable instrument for proactive infrastructure monitoring and maintenance.

ACKNOWLEDGEMENTS

The research was funded by the Austrian Research Promotion Agency (FFG) within the project “AMAzE 2.0” (#897833).

REFERENCES

[1] United Nations Department of Economic and Social Affairs, “2018 Revision of World Urbanization Prospects,” United Nations, Department of Economic and Social Affairs, 2019

[2] R. Scalenghe and F. A. Marsan, “The anthropogenic sealing of soils in urban areas,” *Landscape and Urban Planning*, vol. 90, no. 1-2, p. 1–10, 2009, doi: 10.1016/j.landurbplan.2008.10.011

[3] Z. Ring, D. Damyanovic and F. Reinwald, “Green and open space factor Vienna: A steering and evaluation tool for urban green infrastructure,“ *Urban Forestry & Urban Greening*, vol. 62, p. 127131, 2021, doi: 10.1016/j.ufug.2021.127131

[4] S. Jaafari, A. Shabani, M. Moeinaddini, A. Danehkar and Y. Sakieh, "Applying landscape metrics and structural equation modeling to predict the effect of urban green space on air pollution and respiratory mortality in Tehran," *Environmental Monitoring and Assessment*, vol. 192, no. 412, 15 p., 2020, doi: 10.1007/s10661-020-08377-0

[5] L. Sturiale and A. Scuderi, "The role of green infrastructures in urban planning for climate change adaptation," *Climate*, vol. 7, no. 10, p. 119, 2019, doi: 10.3390/cli7100119

[6] H. Maeda, Y. Sekimoto, T. Seto, T. Kashiyama and H. Omata, "Road damage detection and classification using deep neural networks with smartphone images," *Computer-Aided Civil and Infrastructure Engineering*, vol. 33, no. 12, p. 1127–1141, 2018, doi: 10.1111/mice.12387

[7] M. E. Mohammadi, D. P. Watson and R. L. Wood, "Deep learning-based damage detection from aerial SfM point clouds," *Drones*, vol. 3, no. 3, p. 68, 2019, doi: 10.3390/drones3030068

[8] Z. Feng, A. El Issaoui, M. Lehtomäki, M. Ingman, H. Kaartinen, A. Kukko, et al., "Pavement distress detection using terrestrial laser scanning point clouds – Accuracy evaluation and algorithm comparison," *ISPRS Open Journal of Photogrammetry and Remote Sensing*, vol. 3, p. 100010, 2022, doi: 10.1016/j.ophoto.2021.100010

[9] C. Petschnigg, M. Spitzner, L. Weitzendorf and J. Pilz, "From a point cloud to a simulation model—Bayesian segmentation and entropy based uncertainty estimation for 3D modelling," *Entropy*, vol. 23, no. 3, p. 301, 2021, doi: 10.3390/e23030301

[10] CloudCompare, version 2.14, CloudCompare, available: https://www.cloudcompare.org/release/, 2025

[11] R. Q. Charles, H. Su, M. Kaichun, and L. J. Guibas, "PointNet: deep learning on point sets for 3D classification and segmentation," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 77-85, 2017, doi: 10.1109/CVPR.2017.16

[12] C. Choy, J. Gwak, and S. Savarese, "4D spatio-temporal convnets: Minkowski convolutional neural networks," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3070-3079, 2019, doi: 10.1109/CVPR.2019.00319

[13] J. Deng, W. Zhou, Y. Zhang and H. Li, "From multi-view to hollow-3D: hallucinated hollow-3D R-CNN for 3D object detection," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 31, no. 12, p. 4722–4734, 2021, doi: 10.1109/tcsvt.2021.3100848

[14] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, et al., "PyTorch: An imperative style, high-performance deep learning library," in *33rd Conference on Neural Information Processing Systems*, 12 p., 2019, doi: 10.48550/arXiv.1912.01703

ABOUT THE AUTHORS

Christina Petschnigg
Fraunhofer Austria Research GmbH, Klagenfurt
E-Mail: christina.petschnigg@fraunhofer.at

Alexander Pamler
Fraunhofer Austria Research GmbH, Graz

Kazim Onur Arisan
VIE Build GmbH, Vienna

Jan Morten Loës
VIE Build GmbH, Vienna

Torsten Ullrich
Fraunhofer Austria Research GmbH, Graz

BACK TO CONTENT