2024 European Conference on Computing in Construction
Chania, Crete, Greece
July 15-17, 2024

Capturing Reality Changes from Point Clouds for Updating Road Geometric Digital Twins

Diana Davletshina[1], Varun Kumar Reja[1], and Ioannis Brilakis[1]
[1]University of Cambridge, Cambridge, UK

Abstract

Proactive maintenance of roads enables increased asset lifespan with improved safety and minimised downtime. However, the absence of up-to-date structured information leads to expensive reactive maintenance. Geometric digital twins (GDT) provide digital replicas of object geometry, but no automatic tools exist to maintain them. This paper develops a method for detecting and applying geometric changes to road GDTs. Our solution performs a distance-based comparison of point clouds, estimates the changes and then applies them to GDT. It achieves a 77.87 F1-score in detecting changes and significantly saves time by automating the process, making such digital twins viable and practically applicable.

Introduction

The average travel delay in the United Kingdom has increased by 40% over the past three years due to roadwork disruptions, causing decreased journey time reliability and slower highway speeds (National Highways, 2023). These delays are not just inconvenient but also economically impactful, potentially leading to a £0.50 loss in Gross Value Added per hour per person in Britain (Confederation of British Industry, 2017). Efficient transportation infrastructure, crucial for economic productivity, requires up-to-date information on road assets, with road geometry playing a vital role in traffic safety and operational efficiency (Design Manual for Roads and Bridges, 2020). Road geometric digital twins (GDT) enable such monitoring of road conditions and infrastructure health, offering near real-time data integration and detailed replicas of road systems (Steyn and Broekman, 2022). The life-cycle of GDTs includes construction from point clouds and optionally images and ongoing maintenance to reflect changes in the road infrastructure (Drobnyi et al., 2023; Osadcha et al., 2023). This paper focuses on automating the stage of GDT maintenance, which is a substantial part of the GDT lifecycle.

The main challenge in integrating GDTs is their high cost of maintenance that outweighs potential benefits, introducing a significant barrier to their widespread adoption. This research aims to reduce these costs, eliminating the manual work involved. Specifically, maintaining kilometres of roads requires manual effort in identifying changes and modelling various changed road assets (Davletshina and Brilakis, 2023). Although current research has automated parts of this process, it falls short in scope, often failing to cover a wide range of objects and identifying specific changes that happened to the assets. Additionally, there is a need for a comprehensive and automatic end-to-end pipeline for change detection, followed by applying these changes to GDT.

This paper aims to facilitate efficient, proactive maintenance by providing a tool to update a geometric digital twin automatically, thus delivering a fresh version of GDT quickly and cheaply. We achieve this by first comparing point clouds for changes, identifying the types of changes and applying them in GDT in a streamlined, fully automatic pipeline. Exemplar input and outputs can be seen in Figure 1. The solution covers all the object classes that are already present in GDT. Thus, the solution is flexible and not hard-coded to the specific entities. Furthermore, we consider the wide range of change types in the scope of our work: a) the same object is there, but it underwent pose (rotation and translation) changes; b) the same object is there, but it underwent shape deformations; and c) the object is removed or added. As a result, we detect most of these changes and achieve 77% F1 score.

The organisation of this paper is as follows. The next section discusses the background literature on maintaining GDTs. Then, section Proposed Solution introduces the proposed automatic digital twin maintenance method. Research Methodology, followed by the Results and Discussion section, can be found further. The conclusions of the study are presented in the last section.

(a) GDT (old).

(b) GDT (new).

Figure 1: Examples of the old-version GDT and updated GDT achieved by our proposed solution. Pose transformations and removals are detected in the road furniture and applied in the new version.

Background

The evolution of Digital Twins (DTs) presents a complex challenge, particularly in their efficient updating, which remains unresolved. Schrotter and Hürzeler (2020) highlighted the difficulty in updating urban DTs, highlighting the vast volume of 3D data as a significant obstacle to process. Nevertheless, ongoing research within the built environment sector is actively investigating various methods to effectively maintain and manage digital twins and 3D models of infrastructure assets.

Updating GDTs of built assets starts with collecting appropriate data on the physical counterpart. Osadcha et al. (2023) found that photogrammetry and laser scanning are the primary means for data collection and capturing 3D information. While photogrammetric methods are cheaper, they produce less accurate depth data due to lighting and shadow artefacts (Stal et al., 2013). Mobile laser scanning could serve as a compromise between various methods from images to static scanning, allowing faster operation while driving but producing sparser point clouds than static counterparts (Soilán et al., 2019). Collected data then requires further processing.

A critical step in maintaining GDTs is effectively detecting changes in the data collected at different timestamps. Stilla and Xu (2023) explored this, identifying several types of changes in urban objects using 3D point clouds. They distinguished the following possible change types: appeared, disappeared, partially and fully moved, deformed and unchanged, highlighting that the most common types were the first two, while partial movement was used when detailed changes were required upon use case. Their work estimates point-based, segment- or object-based spatial differences after having point clouds registered. While point-based methods rely on accurate distance measurements, segment- and object-based methods depend on the accuracy of the classification provided. Recent works use deep learning solutions to provide semantic labels and discover further opportunities to match changes with machine learning.

In this realm, Zhu et al. (2023) introduced a method for multi-object relocalization and reconstruction in the segmented into instances 3D point clouds. They could match, register and reconstruct the objects of the input scenes using a deep learning approach, namely an encoder-decoder network. Although solving three tasks simultaneously, they envision the challenge of large-scale spatiotemporal changes yet to be addressed.

Developing dynamic Building Information Models (BIMs) from 3D point clouds is another key research direction in this field. Rausch and Haas (2021) developed a dynamic BIM by updating BIM elements' shape and pose from 3D point clouds, utilising optimisation of geometric object parameters. Despite achieving below-centimetre accuracy, their method's reliance on generic shapes limits its effectiveness with free-form objects. Alternatively, Ghahremani et al. (2018) used 3D point clouds for structural condition monitoring, comparing different-timestamp point clouds to detect deterioration. They identified changed areas by discrepancies between as-built and as-is point clouds, using meshing to update the model. Shellshear et al. (2015) adopted a simpler approach for 3D model maintenance, employing repeated scans from fixed viewpoints and markers to automatically update regions in the target model with point clouds. While effective for built assets, these methods highlight the need for focused approaches in maintaining GDTs for roads.

Knowledge Gaps, Research Objectives and Questions

Maintaining GDTs for roads involves two primary phases: detecting changes and applying these changes. However, a fully automated system for this task is yet to be realized. Also, current methods do not encompass most road assets, such as the frequent objects identified in the work of Davletshina and Brilakis (2023), highlighting a significant gap. While researchers have primarily focused on automating aspects such as identifying temporal changes, further development is still needed. Although advancements in built asset management provide insights, there remains a critical need for domain-specific adaptation and refinement of these solutions to cater specifically to road infrastructure.

This context sets the stage for our objective: to propose and develop a fully automated method for maintaining road GDTs. We outline the following key research questions in addressing our objective: 1) What approaches are most effective for detecting changes in road environments? 2) What strategies should be employed to efficiently apply detected changes to GDTs? 3) Is it possible to devise a versatile method adaptable to various road environments and infrastructures? Addressing these questions will be pivotal in advancing the field and ensuring the accuracy and relevance of road GDTs in a dynamic and evolving landscape.

Proposed Solution

Our approach to detecting and analysing temporal changes in geometric data employs a dual-layer methodology, as illustrated in Figure 2. This is necessitated by the requirement to initially discern the nature of any changes (Change Detection Layer), followed by appropriately addressing the identified changes (Change Application Layer). Initially, the method involves a comparative analysis of point cloud data, representing different states of a physical environment captured at distinct times. As the two point clouds will not contain the same point set, we take advantage of comparative distances to estimate high discrepancies. The framework commences with the older Geometric Digital Twin (GDT) version, incorporating the original point cloud and the newly acquired point cloud for subsequent change detection and processing.

The procedure begins with a thorough analysis of the input point clouds to identify changes by computing distances between points. The algorithm targets areas exhibiting significant alterations, i.e., large distances, indicative of substantial environmental changes. Thus, we filter these points using a distance threshold, which is a hyperparameter to be tuned. Upon pinpointing these areas, we anticipate that most changes, unless due to extensive road reconstruction, will be isolated to individual objects, as road infrastructures typically undergo minimal alterations. As a result, we have these isolated change points, i.e., filtered with the threshold and as a result isolated in space point clusters featuring high distance values such as in Figures 3d and 3e. Then, we apply the DBSCAN clustering algorithm (Ester et al., 1996) on these points – known for its effectiveness with noisy data and not requiring pre-set cluster numbers – which facilitates the segregation and extraction of clusters (e.g., denoted by different colours in Figures 3d and 3e) and their associated objects as identified in the current GDT, enabling a more focused examination.

Subsequently, the nature of the detected changes guides the ensuing steps. For each cluster, the algorithm first finds the closest correspondences between the two versions and ascertains whether the change is an addition, a removal, or a transformation of an existing object. Additions and removals are identified similarly, as they can be seen as the same operation but in opposite sequences. We consider that addition or removal happened if the closest corresponding cluster was already linked to another cluster with a smaller distance. Otherwise, we consider that there was a pose change.

Objects identified as removed are excised from the dataset while new additions are integrated. In cases where existing objects have undergone modifications, the Iterative Closest Point (ICP) algorithm (Zhang, 1994) is deployed on the found cluster correspondences (i.e., correspondence pairs) to compute the transformation, encompassing both rotational and translational adjustments between them. This process begins with estimating the initial transformation for a global match, then minimising the distance between two shapes to locate the closest points. This phase is pivotal in deciphering how the object has shifted or evolved over time. Depending on the nature of the change, the algorithm either applies a transformation to align the old and new states of objects with pose changes if any pose changes were found or otherwise updates the GDT to mirror the new geometry of deformed objects by replacing the object. As a result, an automatic command is sent to perform these described changes to the GDT.

Consequently, our solution presents a systematic and flexible framework for temporal analysis of point cloud data concerning roads, resulting in an up-to-date GDT reflecting reality. This fully automated and scalable solution can be applied over extensive road lengths, significantly enhancing our comprehension and monitoring of environmental alterations across various applications, from asset maintenance to safety analysis.

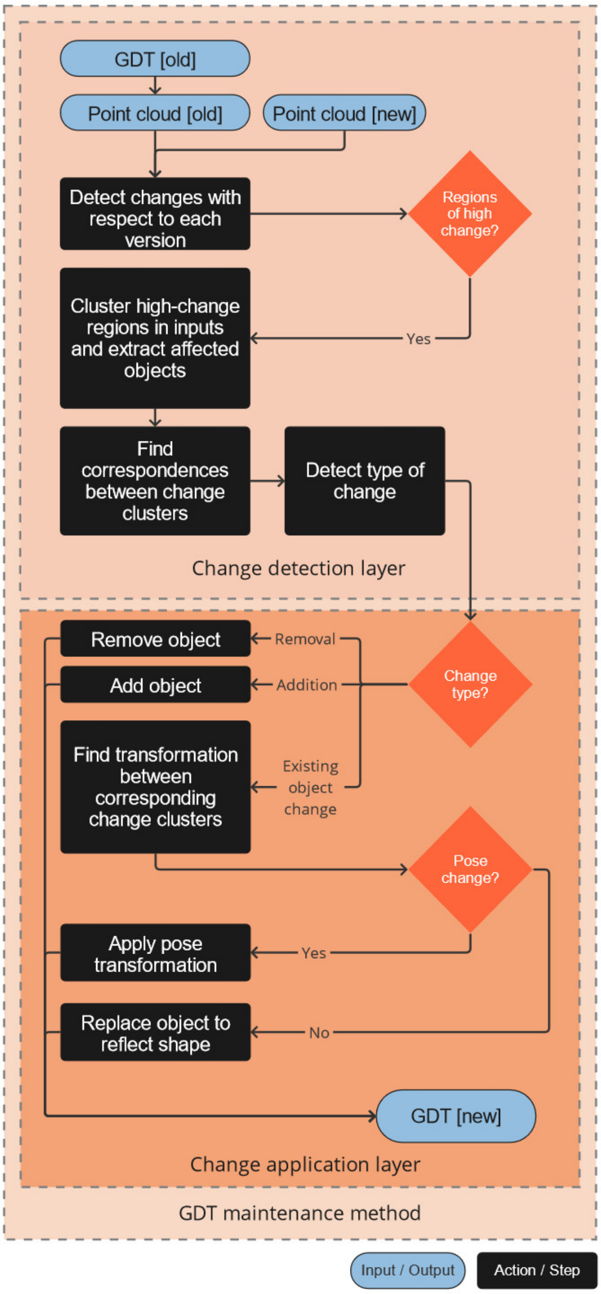

Figure 2: Solution overview.

Research Methodology

We use the Digital Roads (Digital Roads of the Future, 2023) point cloud dataset, which contains 11 labelled classes. We limit the data to the A11 road subset, which comprises 20 road segments, each 200 meters long. As the study conducted in this paper requires time-series data with changes, we synthetically generate the changed version of the point clouds (the *new* state) from the existing dataset (the *old* state) due to the unavailability of a real-life data. We introduce specific perturbations to the data for sufficient differences between data to simulate real-life scanning. For that, we add random noise to the coordinates of both versions of the data (the *old* and *new* states). The noise is normally distributed around mean 0 with a standard deviation of 0.01. Furthermore, we randomly downsample both versions such that 20% of the points are removed. We simulate the changes to the road infrastructure, namely road barriers, traffic signs and lamps, such as fallen assets, leaning assets, missing assets and moved assets in the *new* version of the point cloud. We consider removals of the existing objects and additions of the new objects in the case of the missing objects. These scenarios are seen to be equivalent as one object is missing in one of the versions of scanning.

Our assumptions are as follows. Firstly, we assume that the GDT is semantically labelled and has its original point cloud stored as meta information. If that is not the case, then a synthetic point cloud is expected to be created from the GDT's 3D model. Secondly, we assume that point clouds are registered with each other; data registration is out of the scope of this work. Lastly, we assume that new scanning is performed before new furniture is installed to avoid confusion between removed and newly added assets' pairs.

We use the threshold-based method to identify high-change regions, with the threshold hyperparameter equalling the 0.99 quantiles. Further, we use the DBSCAN clustering algorithm (Ester et al., 1996) to cluster these high-change regions with the following hyperparameters: epsilon equal to 3 and minimum points per cluster equal to 10.

We provide the metrics of precision (Formula 1) showing how many changes were correctly classified, recall (Formula 2) giving the number of correctly predicted items among all correct items and their harmonic mean as F1 score (Formula 3). In the formulas, *TP* refers to true positives, denoting the number of data points classified correctly as positives, *FP* – false positives, denoting the number of data points classified incorrectly as positives; *FN* – false negatives, denoting the number of positive points predicted as negatives. We calculate precision and recall values for each sample and report averaged final metrics.

$$Precision = \frac{TP}{TP+FP} \tag{1}$$

$$Recall = \frac{TP}{TP+FN} \tag{2}$$

$$F1 = \frac{2*Precision*Recall}{Precision+Recall} \tag{3}$$

Results & Discussion

Figure 3 demonstrates the solution's inputs, intermediate and final outputs. Our method can find the changed regions, such as in Figure 3c, where the two input point clouds (Figures 3a and 3b) are aligned. These divergent regions are clustered, as in Figures 3d and 3e, and matched, as well as their transformations are estimated (Figure 3f). As a result, our solution finds pose changes with 90.97% precision and 68% recall (Table 1). The method misses out on a few samples where the changes are minimal, such as small translation and rotation angles. Further, vegetation represents a challenge, introducing additional high discrepancies. Thus, vegetation should be considered for filtering out beforehand using vegetation detection solutions. Additionally, the intermediate found change clusters (e.g., Figures 3d and 3e) come with a certain amount of noise either within clusters or outside of them, which might be considered as a potential scene alteration. As this confuses the method, noise removal at this stage must be included in future work.

(a) Input point cloud (old). (b) Input point cloud (new).

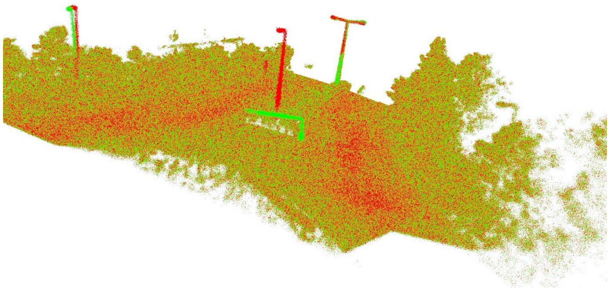

(c) Aligned input point clouds coloured in red and green.

(d) Change clusters in (a). (e) Change clusters in (b).

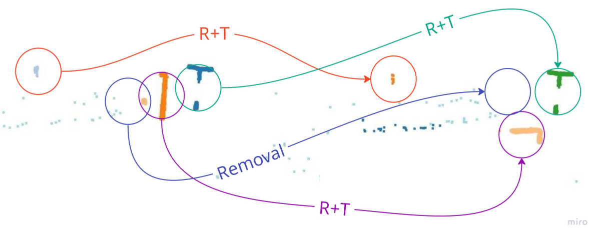

(f) Actions visualised. R+T stands for the estimated rotation and translation matrices.

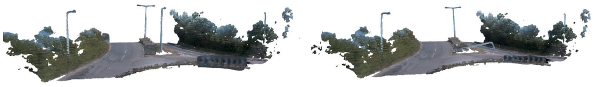

(g) GDT (old). (h) GDT (new).

Figure 3: Examples the inputs, intermediate and final outputs of the method.

Table 1: Change detection evaluation, in %

Precision	Recall	F1 score
90.97	68.07	77.87

Furthermore, the algorithm faces challenges in accurately matching changed clusters when they are significantly distant from each other, particularly in situations where multiple similar changes occur concurrently within the same scene. Handling such intricate change scenarios effectively requires more advanced algorithms for object matching. Incorporating prior knowledge could be crucial in these contexts, aiding in differentiating between scenarios where objects have been moved as opposed to those where objects were removed and replaced with similar ones.

Conclusions

In summary, this work presents a method for maintaining up-to-date geometric digital twins of roads in a fully automatic manner. This method is distinguished by its flexibility and adaptability, which is capable of handling a wide range of objects based on the semantic labels provided in the input GDTs. Our approach goes beyond mere change detection in point clouds; it integrates these changes directly into the GDTs, thereby forming a seamless, end-to-end pipeline that extends from initial scanning to the final GDT for road infrastructures. This development paves the way for the broader implementation of geometric digital twins in the transportation sector, offering significant benefits from the design phase to asset management. This is achieved with reduced costs, courtesy of the automation process. Enhancing the productivity of the sector will yield positive outcomes for road users, including improved connectivity, fewer road closures, and reduced emissions, ultimately leading to a more efficient and environmentally friendly transportation experience. Looking forward, we propose that future research should concentrate on collecting comprehensive datasets that capture real-life temporal changes, tackle the challenges presented by complex and large-scale transformation scenarios, and harness the potential of deep learning for enhanced accuracy in change detection.

Acknowledgments

This work was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) Industrial CASE in partnership with National Highways and Costain [grant No EP/V056441/1] and the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie [grant agreement No 101034337].

References

Confederation of British Industry (2017). Shaping regional infrastructure. Priorities for growth.

Davletshina, D. and Brilakis, I. (2023). A review on constructing and maintaining geometric digital twins of highways.

Design Manual for Roads and Bridges (2020). CD 109 - Highway link design.

Digital Roads of the Future (2023). Digital Roads Dataset. `https://drf.eng.cam.ac.uk/research/digital-roads-dataset`.

Drobnyi, V., Hu, Z., Fathy, Y., and Brilakis, I. (2023). Construction and maintenance of building geometric digital twins: State of the art review. Sensors, 23(9):4382.

Ester, M., Kriegel, H.-P., Sander, J., Xu, X., et al. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. In kdd, volume 96, pages 226–231.

Ghahremani, K., Khaloo, A., Mohamadi, S., and Lattanzi, D. (2018). Damage detection and finite-element model updating of structural components through point cloud analysis. Journal of Aerospace Engineering, 31(5):04018068.

National Highways (2023). Annual Benchmarking Report. `https://nationalhighways.co.uk/media/dnen52ja/abr-v5-171023.pdf`.

Osadcha, I., Jurelionis, A., and Fokaides, P. (2023). Geometric parameter updating in digital twin of built assets: A systematic literature review. Journal of Building Engineering, page 106704.

Rausch, C. and Haas, C. (2021). Automated shape and pose updating of building information model elements from 3d point clouds. Automation in Construction, 124:103561.

Schrotter, G. and Hürzeler, C. (2020). The digital twin of the city of zurich for urban planning. PFG–Journal of Photogrammetry, Remote Sensing and Geoinformation Science, 88(1):99–112.

Shellshear, E., Berlin, R., and Carlson, J. S. (2015). Maximizing smart factory systems by incrementally updating point clouds. IEEE computer graphics and applications, 35(2):62–69.

Soilán, M., Sánchez-Rodríguez, A., del Río-Barral, P., Perez-Collazo, C., Arias, P., and Riveiro, B. (2019). Review of laser scanning technologies and their applications for road and railway infrastructure monitoring. Infrastructures, 4(4):58.

Stal, C., Tack, F., De Maeyer, P., De Wulf, A., and Goossens, R. (2013). Airborne photogrammetry and lidar for dsm extraction and 3d change detection over an urban area–a comparative study. International Journal of Remote Sensing, 34(4):1087–1110.

Steyn, W. J. V. d. M. and Broekman, A. (2022). Development of a digital twin of a local road network: a case study.

Stilla, U. and Xu, Y. (2023). Change detection of urban objects using 3d point clouds: A review. ISPRS Journal of Photogrammetry and Remote Sensing, 197:228–255.

Zhang, Z. (1994). Iterative point matching for registration of free-form curves and surfaces. International journal of computer vision, 13(2):119–152.

Zhu, L., Huang, S., Schindler, K., and Armeni, I. (2023). Living scenes: Multi-object relocalization and reconstruction in changing 3d environments. arXiv preprint arXiv:2312.09138.