Semantic Segmentation of Terrestrial Laser Scans of Railway Catenary Arches: A Use Case Perspective

Bram Ton – PhD Candidate at Saxion University of Applied Sciences

Having access to accurate and recent *digital twins* of infrastructure assets benefits the maintenance, condition monitoring and construction planning of these assets. There are many cases where such a digital twin does not yet exist, such as for legacy structures. In order to create such a digital twin, a mobile laser scanner can be used to capture the current, accurate geometric representation of the structure. With the aid of semantic segmentation, the scene can be decomposed into different object classes.

This decomposition can then be used to retrieve CAD models from a CAD library to create an accurate digital representation. This study explores three deep learning-based models for semantic segmentation of point clouds in a practical real-world setting: PointNet++, Super Point Graph and Point Transformer.

This study focuses on the use case of catenary arches of the Dutch railway system in collaboration with Strukton Rail, a major contractor for rail projects. A diverse, high-resolution, and annotated dataset for evaluating point cloud segmentation models in railway settings is presented. The dataset contains 14 individually labelled classes and is the first of its kind to be made publicly available. A modified PointNet++ model achieved the best overall mean Intersection over Union (mIoU) of 71% for the semantic segmentation task on this new, diverse, and challenging dataset.

Introduction

Renovation, maintenance, condition monitoring and construction of infrastructural projects demand assessments of the current situation. These processes are necessary for the evaluation of the existing situation, possibly leading to advice for re-designing aspects like structural integrity, optimisation of traffic flow, and safety. In addition, the introduction of Building Information Modelling (BIM) and 3D design in general have created an increased demand for accurate, up-to-date, 3D information of existing infrastructure.

3D information is often outdated as the actual constructed infrastructure over time starts to deviate from the original design plans. Furthermore, blueprints do not always exist with a sufficient level of detail, are not available in a digital format, or only exist in 2D. These factors underline the need for methods to accurately, and efficiently capture up-to-date, 3D information of infrastructural assets.

At present, assessments and the subsequent translation to 3D are done mostly manually which is a time-consuming and error-prone task. This has given rise to technology aimed at automating the digitisation of infrastructure, such as photogrammetry and mobile laser scanning. Laser scanning is a method which provides immediate 3D geometric information without any elaborate processing.

Semantic segmentation of point cloud data provides a crucial stepping stone towards automated condition monitoring. For instance, if catenary masts are part of the segmentation process, their tilt can automatically be determined. Maintenance can be planned if certain thresholds for tilt are exceeded.

Multi-epoch data in combination with segmentation also paves the way for even more advanced solutions such as change detection and deformation monitoring. These solutions can contribute to predictive maintenance. For instance, in the case of mast tilt, it would be possible to determine a tilt velocity. This can be used to make projections in the future, and can aid the creation of optimal maintenance plans. A similar approach can be applied to the catenary wires, for instance to monitor sag and stagger.

This paper evaluates several state-of-the-art approaches to semantic segmentation for the digitisation of infrastructure. It does so by evaluating a use case of railway catenary arches in the Netherlands. Catenary arches are the supporting structures above the railway track which carry the power lines for the trains, see Figure 1.

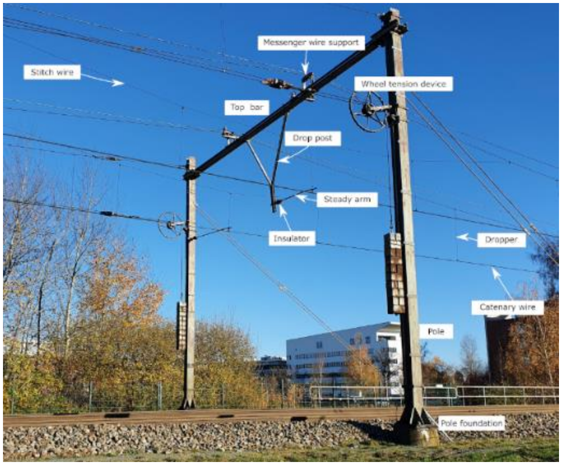

Figure 1: An example of a catenary arch (not in dataset) which shows the labels of the majority of the classes.

The catenary system in the Netherlands consists of a variety of new and legacy arches, with custom and standardised components being mixed. Digitising the physical arches into their 3D, digital counterparts is an ongoing task.

"With the aid of automated digital twinning, it is estimated that planning time for large infrastructural projects can be reduced by one year. – Strukton Rail"

As part of this undertaking, mobile laser scans have been made of a small piece of railway track in the Netherlands.

Catenary arch dataset

To the best of our knowledge, there are no publicly available point cloud datasets of railway catenary arches. Therefore, our work is based on a dataset provided by Strukton Rail, containing an 800 m stretch of railway track near Delft, the Netherlands, containing 15 catenary arches, which has been digitised into a point cloud. The point cloud data was collected with a high resolution terrestrial laser scanner.

The scanned stretch of railway track is made available by the data provider in four chunks of data. A semi-automated method (Figure 2) was used to detect the location of the catenary arches within these chunks.

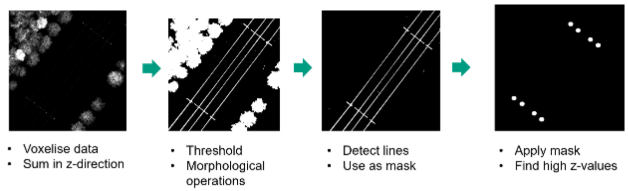

Figure 2: Processing steps for locating catenary arches within a large scene.

After extracting the individual arches, each of the samples was manually labelled in 14 different classes.

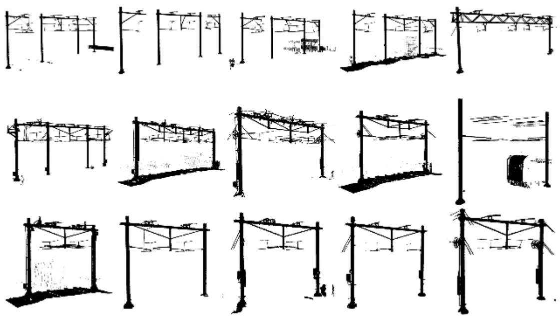

Figure 3: Overview of the dataset. Note the large variation of catenary arch types present in the dataset.

The number of points in a catenary arch ranges between 1.6 and 11 M points. In total, the dataset contains roughly 55.4 M points. The dataset has a large imbalance in the distribution of the classes, which is inherent to the type of object. The three largest classes (unlabelled, pole and top bar) jointly constitute 72.3 % of the points in the dataset. On the other hand, the three smallest classes (dropper, stitch wire and wheel tension device) constitute only 1 % of the dataset.

FULL TEXT

Methods and approach

This paper considers three well established and fundamentally different state-of-the-art approaches that semantically segment a dataset of catenary arches in the Netherlands and compares their efficacy. Specifically, we address variations on PointNet, an implementation of SuperPoint Graph, and Point Transformer.

As the number of samples of the dataset is limited, a leave-one-out cross-validation procedure is used. In total, 14 different components are segmented within the scene. The *mean Intersection over Union* (mIoU) per class is used as a performance metric. This metric weighs all the classes equally and is independent of the class size.

Results and observations

Overall, the modified PointNet++ model performed best, achieving an average class mIoU of 71 %. But when counting the number of best performing metrics, the non-weighted Point Transformer is superior. Its mean class performance is dragged down by just a few classes. The SuperPoint Graph model was not deemed appropriate for this use case as it was very prone to calculation imprecision and had high computational demands.

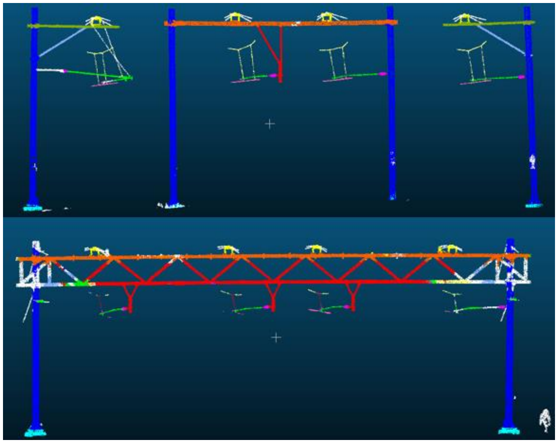

Figure 4: Top: best scoring arch with an mIoU of 0.85, bottom: worst scoring arch with an mIoU of 0.41. Each colour represents a certain predicted class.

To counter the substantial class imbalance of the dataset, the models were also trained using class weights. Surprisingly, this had a negligible effect on the result for the PointNet++ models, yet for the Point Transformer model it was even destructive.

Conclusion and outlook

This work has evaluated three deep learning-based point cloud segmentation methods (PointNet++, SuperPoint Graph and Point Transformer) in a real-world scenario. Overall, the modified PointNet++ model performed best, achieving an average class mIoU of 71 %.

The current dataset was collected using a mobile laser scanner mounted on a tripod; such a solution would not be viable when moving towards a production-ready solution. Future work will focus on using data from mobile or airborne laser scanning solutions to cover large lengths of track.

Though the focus of this work has been on rail infrastructure, the proposed methods can easily be adapted to other domains such as road transport infrastructure. Please get in touch to discuss the possibilities.

Facts

Researchers	Bram Ton * Berend-Jan van der Zwaag Wouter Teeuw Paul Havinga
Academic partners	Saxion University of Applied Sciences
Industrial partners	Strukton Rail
	* b.t.ton@saxion.nl