AI for Rail Defect Detection

Degree programme: BSc in Computer Science
Thesis advisor: Prof. Marcus Hudritsch
Expert: Adrian Egli
Industrial partner: SBB AG, Bern

Rail inspections are crucial for ensuring the safety of the Swiss Federal Railways (SBB). This thesis is a comparative study of deep learning architectures using auto- and expert-labeled data for enhancements in automated rail inspections.

Introduction

The rail tracks at SBB require monthly manual inspections to ensure the early detection and repair of damage, preventing further infrastructure issues and accidents. However, these inspections are risky and take place during normal operations, posing a danger to inspectors. Therefore the SBB initiated the project “AI Streckeninspektion” (AISI) to enable rail inspections without workers on the track. The diagnosis train captures rail images. Although the current AISI system using Faster R-CNN performs well, SBB is exploring newer architectures like DETR and YOLO for potential use. Thus, this thesis serves as a foundation for evaluating the potential use of these architectures. Additionally, to enhance damage detection on rails, an approach was developed for generating auto-labeled data for machine learning models, addressing the high costs of manual data labeling.

Approach

To compare the performances of YOLO, DETR, and Faster R-CNN, all models were trained and evaluated using the same datasets. The approach for creating auto-labeled datasets used recurring defects detected by the model as training data, assuming that a monthly recurring defect is real and reliable. Real-life augmentations provided from these recurring defects under different weather conditions were used to make the models more robust towards the mentioned variations, leveraging real-life augmentations instead of only artificial ones.

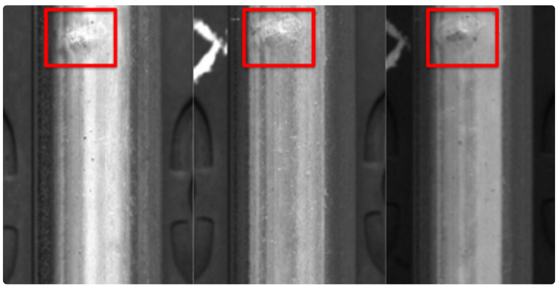

Same Defect on Different Measurement Trips Showing Real-Life Augmentations (Captured with MERMEC Cameras)

Results

This thesis provides AISI a basis for deciding whether to further develop DETR or YOLO models. Each architecture has its pros and cons: DETR requires more training data compared to YOLO and performs better on larger datasets, while YOLO trains faster than Faster R-CNN and DETR and can be used to quickly compare new training sets. Moreover, DETR trains less accurately on smaller datasets compared to YOLO but outperforms it on larger ones for some defects. Overall, the best models from YOLO and DETR, as of May 29, 2024, showed only a 5% performance gap compared to the best Faster R-CNN model on the test set. Furthermore, automated training data generation has proven its effectiveness for rapidly creating large, high-quality datasets comparable to expert-labeled data, facilitating the training of data-intensive architectures like DETR.

Further Research

Based on the findings of this thesis, AISI will decide whether to continue further development on one of the evaluated architectures. Additionally, due to the successful data enhancement approach, AISI will expand future expert-labeled datasets with real-life augmentations for each specific defect. This expanded dataset, enriched with real-life augmentations, will be used to train the models, with the goal of achieving even more accurate and robust damage detection capabilities.

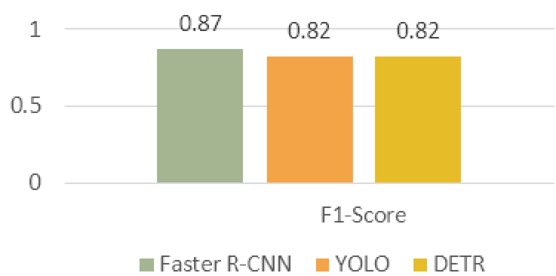

Comparison of F1-Scores Between Different Architectures

Endrit Kastrati
Computer Perception and Virtual Reality