Review of: "CNN-Based Road Damage Detection"

Teerapong Panboonyuen

Potential competing interests: No potential competing interests to declare.

This paper presents a CNN-based method for detecting road damage using**the YOLOv5 object recognition algorithm**. It involves gathering and preprocessing images, training a CNN, and evaluating its performance, achieving approximately 90% accuracy. The system identifies various types of road damage, including cracks, potholes, and bumps, and sends real-time notifications to road engineers to facilitate timely maintenance. This approach aims to improve road safety and reduce inspection time and costs.

Strengths

- The paper is well-written, structured, and easy to read.
- This paper demonstrates a highly accurate method (around 90%) for detecting road damage, indicating the robustness and reliability of the proposed system.
- The integration of real-time notifications to road engineers enhances the practicality and utility of the system in real-world road maintenance and safety applications.

Weaknesses

Although this work has a very attractive starting point, it also has some obvious limitations

- The writing and presentation should be improved. First, the authors should elaborate more on the underlying intuition or motivation behind the YOLOv5.
- I expect random trials across calibration/test splits for the validity experiments. Especially for evaluating the false detection rate, I expect to put calibration and test windows in one bucket and then randomly draw multiple calibrations and test sets. Note that this also applies to the Proposed performance. This is more or less standard in conformal prediction work (and should be lightweight as no re-training is involved).
- Please add and discuss some of the mentioned references. What is the advantage/drawback of the presented approach?
- The compared methods are very old. The authors will need to consider recent methods.
- Testing the proposed attacks on more recent and effective backbones than YOLOv5 might enrich the experimental results.
- This paper lacks detailed explanation and quantitative analysis of the data augmentation techniques used, which is crucial for understanding their impact on the model's performance.
- There is insufficient comparison with other state-of-the-art road damage detection methods, which could provide a clearer context of the model's performance relative to existing solutions.
- The discussion on the scalability of the proposed system to different road conditions and environments is limited, raising

questions about its broader applicability.

- This paper does not address potential challenges in implementing the system on a larger scale, such as hardware requirements, deployment costs, and maintenance.
- The paper mentions achieving a high accuracy but lacks a comprehensive discussion of other evaluation metrics, such as precision, recall, and F1-score, which are essential for a thorough assessment of the model's performance.

method.

- Fine-tuning the parameters should be discussed.
- Lack of computational cost analysis (both training and inference time). According to the implementation details and the overall inference process (testing time), the authors are suggested to conduct some analysis on this issue.