StreetSense: ML-Powered Sidewalk Monitoring

Brian Li, Vanessa Pierre, Daventzky Edouard
Advisors: Suvosree Chatterjee, Jenifar Selim, Sara Sevilla, Serena Hoermann, Ph.D., Jason Hallstrom, Ph.D.

BACKGROUND

Public infrastructure conditions strongly affect mobility, accessibility, and safety, particularly for pedestrians, cyclists, and other vulnerable groups. Although existing AI and crowdsourcing tools can map sidewalk accessibility, they often lack direct integration with local city reporting systems.

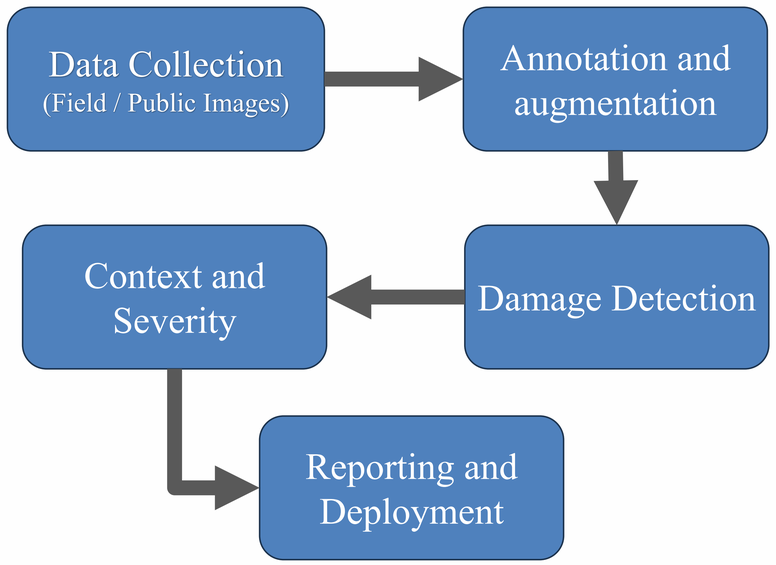
- Our proposed system addresses this gap by helping pedestrians identify, organize, and report physical obstructions more clearly, complementing existing city services rather than replacing them.

Research Question: How can machine learning and computer vision improve pedestrian agency over physical obstructions?

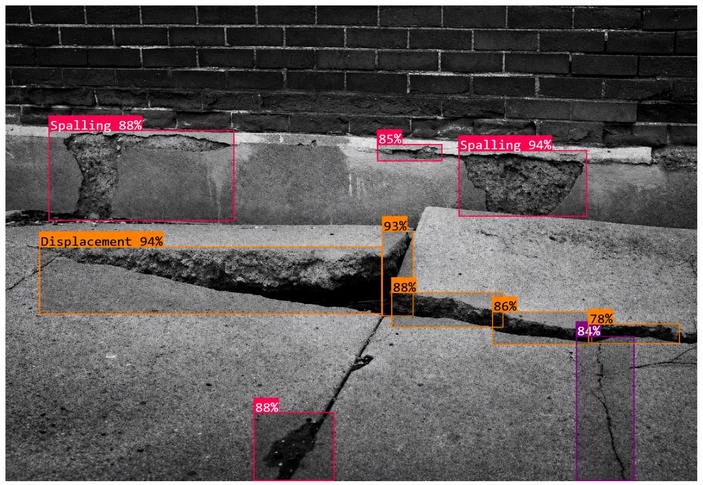
PROJECT GOALS

- Build an AI system that detects sidewalk damage from uploaded images.
- Rank problems by severity so dangerous or high-impact issues can be prioritized.
- Create a map-based dashboard showing where infrastructure problems are in West Palm Beach.
- Test the system on a small pilot area before expanding to more neighborhoods.

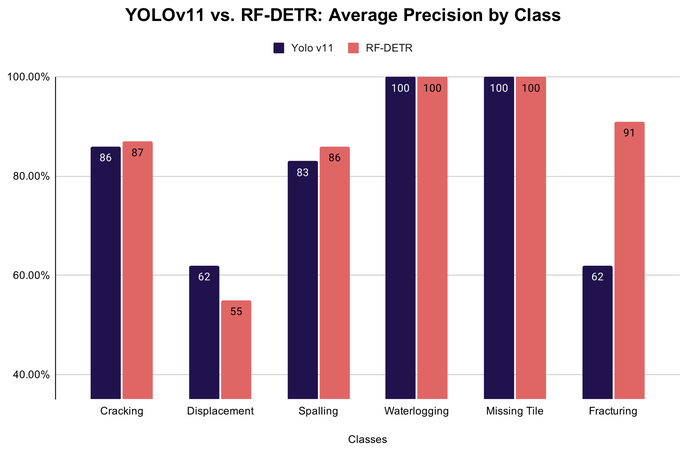
METHODS



RESULTS



The sample output above demonstrates that the model can simultaneously locate and classify multiple types of sidewalk damage within a single street-level image.



- Both YOLOv11 and RF-DETR achieved similar overall performance across the six sidewalk-defect classes.
- YOLOv11 performed better at detecting displacement, while RF-DETR performed better in some other categories.
- Performance varied more by defect type than by the model used. The detection example shows that multiple sidewalk defects can be located and classified in a single image.

Key finding: Both models performed comparably, but performance varied substantially by defect class.

QR CODES

Demo Site

https://tinyurl.com/374ezyaj

Data

Precision Graph

Interview Results

DISCUSSION

In creating a six-class sidewalk detection model, we found that running two different models on the same photo yielded the most accurate results. We chose to run YOLO v11 and RF-DETR because, as shown in the figures above, there are specific classes that one model detects better than the other. Instead of us choosing one and risking sidewalk damage going undetected, we made the decision to utilize both.

Future work includes:

- Extending this technology to roads and bike lanes
- Expanding the application to a comprehensive app for pedestrians

ACKNOWLEDGMENTS

The work was supported through (1) the NSF Engineering Research Center for Smart Streetscapes (CS3), funded through NSF Cooperative Agreement EEC-2133516.

REFERENCES

[1] D. Morra et al., “Mapping sidewalk accessibility with smartphone imagery and computer vision,” Philosophical Transactions of the Royal Society A, vol. 382, no. 2285, 2024.
[2] M. Saha et al., “Project Sidewalk: A web-based crowdsourcing tool for collecting sidewalk accessibility data at scale,” Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems, 2019
[3] D. Castells, “Obstacle detection and avoidance on sidewalks,” Proceedings of the International Conference on Computer Vision Theory and Applications, 2010.
[4] A. Boussihmed, K. El Makkaoui, I. Ouahbi, Y. Maleh, and A. Chetouani, “A TinyML model for sidewalk obstacle detection: aiding the blind and visually impaired people,” Multimedia Tools and Applications, 2024