Street Parking Presence Inference from Street-Level Imagery via Multi-Cue Detection and Geo-Aggregation

Chirag Jain (9087168606), Ritik Singh (9087047321)

February 23, 2026

1) What problem are we trying to solve?

Street parking availability is governed by distributed curbside cues such as parking signs, parking meters/pay stations, and curb markings. These cues are often sparse, viewpoint-dependent, and frequently occluded. Consequently, single-image classification of "street parking present?" is brittle: a negative prediction can simply mean the relevant cue is outside the camera field of view.

We propose to infer **street parking presence at the street-segment level** using a two-stage pipeline:

1. **Image-level cue detection:** classify whether an image contains (a) parking sign(s), (b) parking meter/pay station(s) ***if feasible***, and (c) curb-related cues (curb presence/appearance; curb color ***if feasible***).
2. **Segment-level aggregation**(during inference)**:** cluster nearby street-level images into a segment and aggregate cue evidence across images to infer whether the segment supports on-street parking.

2) Why is this problem important? Why are we interested in it?

Street parking plays a significant role in urban mobility and traffic congestion. A substantial portion of urban traffic congestion is often attributed to drivers cruising in search of parking, leading to increased fuel consumption, emissions, and travel delays. Despite its importance, reliable and up-to-date information about street parking infrastructure is frequently incomplete, inconsistent, or manually maintained by municipalities.

Curbside space must support competing uses including parking, loading zones, ride-hailing pickup/dropoff, and emergency access. As cities attempt to modernize curb management policies, there is a growing need for scalable, automated tools that can monitor and interpret curbside conditions from visual data. Vision-based curb analytics has recently been used for curb lane monitoring and illegal parking impact estimation, demonstrating the feasibility and societal relevance of computer vision for curbside analysis [1].

However, most current parking information systems rely on manual surveys, static GIS databases, or crowdsourced updates that are difficult to scale and maintain. Automatically inferring street parking presence from street-level imagery provides a scalable alternative that can assist urban planners, navigation systems, and smart-city infrastructure.

From a computer vision perspective, this problem is particularly compelling because parking-related cues are small, sparse, and viewpoint-dependent. Relevant evidence may be distributed across multiple images along a street segment. This naturally motivates multi-view reasoning and cue aggregation, making the problem both societally important and technically challenging.

3) What is the current state-of-the-art?

There is limited direct work on *street parking presence inference* from street-level imagery. Most related work falls into three categories:

(A) Parking sign detection and recognition. A small set of works focuses specifically on detecting and recognizing street parking signs. Chau et al. study real-time parking sign detection and symbol detection, noting limited public data and the complexity of parking sign understanding [2]. Jiang presents an end-to-end pipeline (detection + recognition + trust system) for street parking signs in a thesis setting [3]. Zhong et al. target parking-sign text recognition using a dictionary-guided approach [4]. These works primarily emphasize *sign-centric* understanding and typically assume that the sign is observable in-frame.

(B) On-street parking space / occupancy localization. Other work detects parking spaces or occupancy rather than inferring whether parking is allowed/instrumented on a segment. For example, Garta et al. detect on-street parking spaces using YOLO models and recommend spaces using KD-tree search [5]. Morell et al. locate on-street parking spaces from low-quality public camera imagery using deep learning and geometric postprocessing [6]. These approaches address "where are available spaces?" rather than "does this street segment support street parking infrastructure/permission?"

(C) Curb monitoring and curbside analytics. Vision-based monitoring of curb usage and illegal parking impacts further supports the importance of curbside perception, but is typically framed around activity/impact estimation rather than segment-level parking presence inference [1].

Key gap. Prior work that directly addresses street parking focuses mostly on parking signs and often requires custom data collection [2, 3, 4]. There is no widely adopted benchmark dataset or standard end-to-end approach for *street-segment* parking presence inference that integrates multiple cues and multi-view aggregation (and this a very big problem - the lack of good datasets).

4) Re-implement an existing solution or propose a new approach?

We will do both:

1. **Baseline (re-implementation) -** We will implement a strong baseline for **parking sign presence** using a modern detector/classifier, aligning with sign-centric pipelines in the literature [2, 3].
2. **Proposed approach (novel extension) -** We will extend beyond sign-only inference by introducing:
    - **Multi-cue detection:** parking sign presence, and optionally parking meter/pay station presence; plus curb-related cues (curb presence/appearance; curb color if feasible).
    - **Geo/segment aggregation:** cluster images by geographic proximity (and/or capture sequences) and aggregate cue evidence to infer street-segment parking presence.

This reframes the task as multi-view evidence integration, closer to a multiple-instance inference problem: a segment can be positive even if many individual images do not contain visible cues.

5) Why are existing approaches insufficient? Why will ours work better?

Why insufficient. Parking sign approaches can fail when signs are occluded, distant, blurred, or simply not in frame [2, 3, 4]. Parking space localization work solves a different objective (slot/occupancy localization), often relying on fixed camera settings or gap/space cues rather than curbside infrastructure cues [5, 6]. Curb monitoring work demonstrates feasibility and importance but does not directly infer segment-level parking presence from distributed cues [1].

Why ours can work better. Street parking presence is fundamentally a *segment-level* property. Cue redundancy (signs, meters, curb cues) and multi-view aggregation should improve robustness: even if a single image lacks visible evidence, nearby images along the same segment may contain a clear sign or meter. Aggregation therefore targets the dominant failure mode of single-image sign-centric systems: missing evidence due to viewpoint and sparsity.

Feasibility note. Curb color is potentially helpful but may be dataset- and time-dependent; if we cannot obtain reliable supervision and visibility, we will treat curb color as a stretch goal and keep curb presence/appearance cues in-scope.

6) How will we evaluate performance? What results/comparisons will we show? Timeline?

Evaluation

Task A: Image-level cue detection. Binary classification/detection for:

- parking sign present vs. not present
- parking meter/pay station present vs. not present (if feasible)
- curb cue present vs. not present; curb color bucket (if feasible)

Metrics: accuracy, precision/recall/F1; and mAP for detection-based variants.

Task B: Segment-level street parking inference. Cluster images by GPS proximity (e.g., DBSCAN radius / fixed-distance threshold) and infer segment label. Metrics: accuracy, AUROC, precision/recall/F1.

Planned comparisons / ablations

- **Single-image vs. segment aggregation:** quantify gains from multi-view inference.
- **Cue ablations:** signs only; signs+curb; signs+meters (if feasible); full multi-cue model.
- **Error analysis:** occlusion, distance, night, motion blur, rare sign designs, worn/ambiguous curb paint.

Timeline

- **Weeks 1–2:** implement parking sign baseline; implement clustering/aggregation; build a small manually verified evaluation set.
- **Weeks 3–4 (midterm):** segment inference experiments; initial ablations (single-image vs aggregated).
- **Weeks 5–6:** integrate curb cues; integrate meter cue if feasible; run full ablations and error analysis.
- **Weeks 7–8:** finalize experiments; prepare project webpage and presentation materials; package code for reproducibility.

References

[1] J. Gao, F. Zuo, K. Ozbay, O. Hammami, and M. L. Barlas, "A new curb lane monitoring and illegal parking impact estimation approach based on queueing theory and computer vision for cameras with low resolution and low frame rate," *Transportation Research Part A: Policy and Practice*, vol. 162, pp. 137–154, Aug. 2022.

[2] H. Chau, Y. Jin, J. Li, J. Hu, and W. Cheng, "Real-time street parking sign detection and recognition," in *AI4AD Workshop at IJCAI-ECAI 2022*, 2022.

[3] Z. Jiang, "Street parking sign detection, recognition and trust system," Master's thesis, University of Washington, 2019.

[4] D. Zhong, J. Li, W. Cheng, and J. Hu, "Dictionary-guided text recognition for smart street parking," in *Proceedings of the British Machine Vision Conference (BMVC)*, 2023.

[5] I. Y. Garta, W. E. Manongga, S.-W. Huang, and R.-C. Chen, "On-street parking space detection using yolo models and recommendations based on kd-tree suitability search," *CMC-Computers, Materials & Continua*, 2025.

[6] J. Á. Morell, G. Luque, and E. Alba, "On-street parking space localization with deep learning using low-quality images from public cameras," *Internet of Things*, vol. 32, p. 101619, 2025.