

“Where Can I Park?” Understanding Human Perspectives and Scalably Detecting Disability Parking from Aerial Imagery

Jared Hwang
Allen School of Computer Science
University of Washington
Seattle, Washington, USA
jaredhwa@cs.washington.edu

Chu Li
Allen School of Computer Science
University of Washington
Seattle, Washington, USA
chuchuli@cs.washington.edu

Hanbyul Kang
Allen School of Computer Science
University of Washington
Seattle, Washington, USA
kahaby@cs.washington.edu

Maryam Hosseini
City and Regional Planning
University of California Berkeley
Berkeley, California, USA
maryamh@berkeley.edu

Jon E. Froehlich
Allen School of Computer Science
University of Washington
Seattle, Washington, USA
jonf@cs.washington.edu




Figure 1: We introduce *AccessParkCV*, a novel ML-based disability parking detection and characterization pipeline, composed of an object detection and oriented bounding box model, that (1) locates disability parking from aerial orthoimages and (2) estimates the accessible width of disability parking spaces. We motivate its design using results from a needfinding interview study, and present two case study applications utilizing the pipeline inferences.

Abstract

Accessible parking is critical for people with disabilities (PwDs), allowing equitable access to destinations, independent mobility, and community participation. Despite mandates, there has been no large-scale investigation of the quality or allocation of disability parking in the US nor significant research on PwD perspectives and uses of disability parking. In this paper, we first present a semi-structured interview study with 11 PwDs to advance understanding of disability parking uses, concerns, and relevant technology tools. We find that PwDs often adapt to disability parking challenges according to their personal mobility needs and value reliable, real-time accessibility information. Informed by these findings, we then introduce a new deep learning pipeline, called *AccessParkCV*, and parking dataset for automatically *detecting* disability parking and inferring *quality characteristics* (e.g., width) from orthorectified aerial

imagery. We achieve a micro- $F1=0.89$ and demonstrate how our pipeline can support new urban analytics and end-user tools. Together, we contribute new qualitative understandings of disability parking, a novel detection pipeline and open dataset, and design guidelines for future tools.

CCS Concepts

• **Human-centered computing** → **Accessibility systems and tools; Empirical studies in accessibility.**

Keywords

accessibility, disability parking, aerial imagery, computer vision, urban planning

ACM Reference Format:

Jared Hwang, Chu Li, Hanbyul Kang, Maryam Hosseini, and Jon E. Froehlich. 2025. “Where Can I Park?” Understanding Human Perspectives and Scalably Detecting Disability Parking from Aerial Imagery. In *The 27th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS ’25)*, October 26–29, 2025, Denver, CO, USA. ACM, New York, NY, USA, 20 pages. <https://doi.org/10.1145/3663547.3746377>



This work is licensed under a Creative Commons Attribution 4.0 International License. *ASSETS ’25, Denver, CO, USA*

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0676-9/25/10

<https://doi.org/10.1145/3663547.3746377>

1 Introduction

Accessible parking is critical for people with disabilities (PwDs), allowing equitable access to destinations, independent mobility, and community participation [43, 49, 70]. Although public transit remains a valuable option, it is not always available, equally distributed, or accessible [56, 65, 74]. Since the *Americans with Disabilities Act* (ADA) of 1990 [76], 4-8% of public parking spaces in the US are required to be reserved for drivers and passengers with disabilities [43]. However, this mandate includes no compliance or verification measures [76], and there has been no systematic large-scale study of the allocation and characteristics (e.g., size) of disability parking in the US. Further, there is limited qualitative research about how PwDs feel about and use disability parking [43, 60].

In this paper, we first present findings from semi-structured interviews with PwDs about their perspectives and uses of disability parking. We then introduce a novel computer vision (CV) pipeline for disability parking inference, called *AccessParkCV*, that not only *identifies* and *locates* disability parking slots from aerial imagery (e.g., plane flyovers) but also *characterizes* their widths. While our algorithms build on emerging work in deep learning to assess built infrastructure from aerial images (e.g., [32, 67]), including parking [8, 9, 21, 82, 91], we uniquely contribute new techniques for disability parking and an accompanying labeled dataset for others to build on our work. As mixed-methods research, our research questions (RQs) span both the formative and technical:

- **RQ1:** How do PwDs use disability parking, and how does it fit into their overall transportation profile?
- **RQ2:** How well can CV models locate and characterize disability parking from aerial imagery?
- **RQ3:** How could inferences from such a CV model be used to build technology tools for PwDs and policymakers?

To address these questions, we first recruited eleven PwDs for a three-part, semi-structured interview (60-90 minutes). In Part 1, we asked about general transit choice behavior and travel before specifically investigating disability parking in Part 2. In Part 3, we showed thirteen sketches of envisioned technology tools for disability parking, including route planning, *in situ* navigation, and the "out-of-car" experience once parked. From thematic analysis of study transcripts, we find that PwDs' priorities vary significantly with each individual's mobility needs and that such needs are not met by current disability parking policies or implementations. Consequently, PwDs must adapt dynamically when preparing to travel and when parking. We also find that PwDs value real-time and reliable information about accessible parking and places of interest (POIs) in advance of making trips.

Informed by these qualitative findings, we designed, built, and evaluated a novel disability parking inference pipeline, *AccessParkCV*, capable of identifying, locating, and characterizing the width of disability parking spots in orthorectified aerial imagery. To train and evaluate our pipeline, we annotated 11,762 parking spots from open-source aerial orthoimages of Seattle, WA [84], Washington, DC [28], and Spring Hill, TN [48]. Our open-source pipeline achieves state-of-the-art performance: a micro-F1 score [80] of 0.89 on detections and an average error of 5.40% ($SD = 17.52\%$) when estimating parking width. Further, to explore generalizability across datasets and the pipeline's susceptibility to error, we performed

two additional investigations: an evaluation on regions outside the training set including at different resolutions and a comparison of the pipeline's parking width error to the error between two human annotators. Moreover, to demonstrate the value of these automated extractions, we introduce two example applications: a geoanalytic disability parking map for policymakers and disability advocates, and an end-user personalized mobile app for seeing and filtering for disability parking.

In summary, we contribute **formative study findings** about the perceived benefits, uses, and drawbacks of disability parking in the US; a **new algorithmic pipeline** for detecting and locating disability parking spots from aerial imagery and for characterizing their quality; a **new disability parking dataset** for others to build on our work; and **two example applications** demonstrating the potential of our pipeline and dataset. To support open science and enable others to build on our work, both the annotated dataset¹ and *AccessParkCV* code² are open source.

2 Related Work

We provide background on the role and implementation of disability parking before contextualizing our work in parking-focused mapping tools and computer vision techniques.

2.1 Disability Parking

Disability parking is fundamental to ensuring equitable access and independent mobility for people with disabilities (PwDs) [43, 49]. In the US, its origins are intertwined with the broader disability rights movement and the push for societal inclusion throughout the latter half of the 20th century [78]. Early federal legislation, such as the *Architectural Barriers Act* of 1968 [1] and *Section 504* of the *Rehabilitation Act* of 1973 [2], began mandating accessibility in specific contexts, laying the groundwork for more comprehensive standards. However, it was the passage of the *Americans with Disabilities Act* (ADA) in 1990 [76] that established enforceable, nationwide requirements for accessible parking in public accommodations and commercial facilities.

The *ADA Standards for Accessible Design* [70] provide detailed specifications for disability parking spaces. Key requirements include specific minimum dimensions, adjacency to marked access aisles that provide critical space for ramp deployment or maneuvering mobility devices, and clear signage featuring the *International Symbol of Accessibility*. Van-accessible spaces mandate even greater width and vertical clearance, along with specific signage, to accommodate larger vehicles often equipped with lifts [87]. The required number of accessible spaces is calculated based on the total parking capacity of a lot, ensuring proportional availability [43, 70].

Despite these federal mandates, manual audits conducted across regions consistently reveal significant deficiencies in the implementation and maintenance of disability parking [36, 81]. For instance, a survey of 50 public accommodations in Maryland found that 26% failed to provide the legally required number of accessible spaces, and 32% lacked a necessary van-accessible space [81]. Common defects included improperly marked or absent access aisles and inadequate accessible routes connecting parking to building

¹<https://huggingface.co/datasets/makeabilitylab/disabilityparking>

²<https://github.com/makeabilitylab/AccessParkCV>

entrances [81]. Similarly, fieldwork at 36 locations in Melbourne, Australia—which has similar disability parking requirements to the US [3]—found that none fully complied with recommended best practices, citing issues like undersized spaces, poor line marking, and missing curb ramps [36].

These infrastructural shortcomings directly impact the lived experiences and mobility of PwDs [36, 58, 73]. For example, the aforementioned Melbourne study also surveyed 71 disability parking permit holders, revealing widespread feelings of insufficient accessible parking availability, particularly in busy activity centers, along with concerns about the safety of space placement relative to traffic [36]. Research in Saga, Japan, observed decreased vacancy rates after introducing a formal permit system, with subsequent surveys indicating user frustration regarding lack of availability [58]. Furthermore, ethnographic work with families managing childhood disability in Ontario, Canada, demonstrated how parking facilities could be "technically" accessible according to regulations but functionally inaccessible due to factors like timing, poor snow clearance, or temporary obstructions. This often forced families to develop their own coping strategies with little institutional support, highlighting the difference between compliance and true usability [73].

Finally, a large body of work investigates disability parking abuse or violations, including motivations [43] and potential enforcement or prevention strategies [23, 59, 79]. Though our research does not focus on violations, such misuse exacerbates the availability issues stemming from inadequate provision or poor design [79].

In sum, prior research establishes the critical importance of disability parking, details the design standards intended to ensure its utility, and documents common failures in implementation along with the negative consequences for users. Building on this foundation, our work seeks to deepen our understanding of how disability parking affects the planning and travel behavior of PwDs through qualitative inquiry. Further, we introduce a novel disability parking detection pipeline, moving beyond the limitations of localized manual audits highlighted in previous studies.

2.2 Accessibility and Parking Technology Tools

Many technology tools aim to provide PwDs a gateway to information about accessibility and infrastructure features in the physical world, including sidewalks [74], crosswalks [5], and POI access [11]—though disability parking itself has received less attention. To enable these tools, varying approaches are used to source data about the features in question, including crowdsourcing and specialized hardware and computer vision. We discuss both accessibility crowdsourcing and parking management systems as two of these information gateways.

Crowdsourcing is a commonly used approach to scalably collect information about the built environment, and many tools are built off of crowdsourced data [35, 57, 63, 74]. A user survey of the wheelchair accessibility crowdsourcing app, *Wemap* [57], found that people liked the idea of crowdsourced apps and said that they would use them. A non-accessibility focused survey on crowdsourcing found that people want reliable information that can suggest parking location and dimensions [61]. For accessibility specifically, tools such as *AXS Map* [35], *EasyWheel* [62], *WheelMap* [63], and *SmartBFA* [38] are all crowdsourcing platforms that allow users to

manually (or passively, in the case of *SmartBFA*) tag the accessibility of locations, with both *AXS Map* and *EasyWheel* including parking information. Although crowdsourcing is a powerful tool, it can suffer from data reliability and requires a user pool to maintain and update accessibility tags [12, 74]. Our contribution approaches the problem from a large-scale auditing perspective that does not depend on users to characterize accessible spaces, though crowdsourcing could be used for human review and quality control.

Another avenue for information access to the real-world is parking management systems that provide both real-time information and may help mitigate disability parking violations [26, 44, 75]. Proposed solutions range from sensors placed on the car or space [15, 24, 26, 44] to camera and machine learning-based parking monitoring [50, 85]. While such systems may allow the user to see availability in a particular lot ahead of time, installing sensors requires hardware and does not address parking that was not initially integrated in the parking management system. Moreover, our focus is on new automated techniques to identify disability parking spaces at scale (e.g., to enable new types of urban analytics), something bespoke infrastructure-sensing solutions cannot achieve.

2.3 Automatically Detecting Parking Spots

Our work introduces a novel CV pipeline for disability parking detection, building on literature in automated analyses of the built environment from aerial imagery [6, 32, 34, 40, 47, 66]. Given the many potential applications of automated parking detection in general—from autonomous driving to parking management—there is a wide array of approaches for detecting parking spaces and occupancy (e.g., [21], [33], [9]). However, no such work is dedicated to detecting and characterizing *disability parking* specifically, so we discuss general methods here.

Live occupancy attempts to detect if there is a vehicle in a given space at a given time [7]. A common approach involves a static video camera overseeing the parking lot, with parking space positions annotated *a priori* by a human; subsequently, an algorithm attempts to detect the presence of a vehicle within each designated space, including background subtraction [10, 21, 90], feature matching from drone footage [72], and variations of convolutional neural nets (CNNs) [7, 68]. Similar approaches are also used without having a static feed; [27] and [75] determine occupancy via self-supervised learning and drones to examine license plates with CNNs, respectively, without requiring a permanently installed camera. Although these approaches are generally effective at detecting if a car is present, they are not as relevant to our work given that (1) the spaces must already be identified, (2) we do not attempt to determine if a space is occupied in real time, and (3) we make use of aerial, bird's-eye imagery vs. infrastructure-mounted cameras.

Another common approach, particularly for automated urban analyses like zone use, is to detect entire parking lots from aerial images, including producing segmentation masks with variations on CNNs (e.g., [9], [31]), or just classifying if an image contains a parking lot [18]. However, we do not attempt to detect full parking lots, only individual spaces given an arbitrary location.

More relevant to our work are methods to detect if there is a parking space and, if so, geometric boundaries. One approach uses aerial




Figure 2: Three example design probes of the 13 total shared with participants: (a) a routing tool showing the path from a parking space to the entrance of a destination, estimating both driving time and time to traverse to the entrance, (b) information about a specific place of interest, including disability parking space counts, and (c) a settings page, where one can filter for disability parking based on personalized needs. See the supplementary materials for the full 13 probes.

imagery and breaks parking lots into their constituent structures (e.g., lines and boxes), then counts the boxes [8, 39, 77]—similarly, our characterization model also leverages the rectilinear structure of parking spaces. Other approaches to detecting disability parking use machine learning; Varghese *et al.* [88] use support vector machines and background subtraction to detect both delimited and non-delimited (no demarcating lines) parking spaces. Seo *et al.* [77] use self-supervised learning with ~70% accuracy. Lindblom *et al.* [54] use CNNs to detect parking spaces from aerial imagery, but they see a high false positive rate in cases where there are many lines in the street or box-like shapes. Most similarly to our work, Huang *et al.* [33] attempt to generate inclined, or oriented, bounding boxes around parking spaces from aerial footage (as we do to characterize the width of spaces). They achieve precision and recall in the high-90-percent range, depending on the test dataset. We build on the above work but differ in our focus on detecting and characterizing disability parking specifically, and use of a transformer-based architecture for our parking detection model.

3 Study 1: Interview Study of PwDs

We first present our semi-structured interview study of how PwDs use, feel about, and struggle with disability parking (Study 1) before introducing AccessParkCV and example applications partially informed by those formative findings (Study 2).

3.1 Methods

To investigate the effects of disability parking availability, quality, and characteristics on PwD travel planning, behavior, and attitudes, we conducted a remote interview study with 11 PwD participants.

3.1.1 Participants. Participants were recruited through word of mouth, social media, and disability organizations. Study advertisements linked to a screening survey that collected demographics, disability information, and parking eligibility duration. We selected 11 adult participants (18+) who either had disability parking eligibility themselves or were primary drivers for eligible individuals, all possessing disability parking placards. Because our study was advertised online by some partners, we used *reCAPTCHA* scores embedded in our Qualtrics screener to filter fraudulent responses. See Table 1 for detailed participant information.

3.1.2 Procedure. Our semi-structured interview consisted of three sections: (1) broad exploration of how PwDs use transportation and disability parking to identify transit choice patterns; (2) examination of disability parking experiences, including advantages, disadvantages, and potential improvements; and (3) assessment of

P	Age	Placard Reason	Mobility Aid	Travel Freq.
1	25-34	Temp. Disability	Crutches	> 4/week
2	45-64	Perm. Disability	Wheelchair	> 4/week
3	45-64	Perm. Disability	Wheelchair	> 4/week
4	25-34	Perm. Disability	Wheelchair, Leg Braces, Crutches	> 4/week
5	18-24	Perm. Disability	Wheelchair	> 4/week
6	45-64	Primary driver for PwD	Wheelchair	> 4/week
7	45-64	Perm. Disability	Wheelchair	> 4/week
8	45-64	Perm. Disability	Wheelchair	> 4/week
9	45-64	Perm. Disability	Wheelchair	2-4/week
10	45-64	Perm. Disability	Wheelchair	> 4/week
11	45-64	Perm. Disability	Wheelchair	> 4/week

Table 1: Participant demographics. Travel frequency includes commuting and personal trips.

participants' technology usage habits and feedback on 13 design probes representing potential disability-parking related tools across five thematic areas: route planning, place-of-interest investigation, in situ navigation, out-of-car experiences, and curbside assistance—see examples in Figure 2. Informed by [30], design probes were shown as PowerPoint slides.

Prior to the main study, two researchers conducted a pilot interview with one participant (P1), after which minor changes were made to the interview script to ensure clarity. For the subsequent sessions, two researchers led five interviews each. All interviews were conducted via Zoom, with each session lasting 60-90 minutes. Participants were compensated US\$25 per hour.

3.1.3 Analysis. We recorded and transcribed all interview sessions. For analysis, we used a combination of deductive and inductive coding [13]. To start, we created an initial codebook based on the interview protocol. Two primary researchers then engaged in an iterative process of coding and peer checking [20, 53, 64] to ensure the reliability and validity of the analysis. Of the 11 interviews, one researcher coded five and another coded six. The two researchers then met to discuss and resolve any disagreements, updating the codebook as necessary. The researchers then swapped their set of assigned interviews and spot-checked each others' coding.

3.2 Results

Overall, we found that PwDs approach and use disability parking in diverse and personalized ways—often as a function of their mobility disability, prior experiences, and vehicle type. Still, common patterns emerged around preferred disability parking designs, concerns and problems, and opportunities for improvement. We organize our findings around two key research questions: (1) How do PwDs use disability parking? (2) What kind of technology tools or information would PwDs find useful?

3.2.1 Experiences with Disability Parking. Accessible parking is a complex reality for PwDs, shaped by individual mobility, the need for self-devised solutions, and frustrating encounters with poorly designed built environments and public misuse. While regulations exist, their inconsistent implementation creates access barriers.

Personal mobility needs shape parking preferences. PwDs assess parking spaces based on a range of factors from mobility needs to vehicle size and door locations. P7 noted, for example, that "if the car can fit in and I can get in and out of the car, that's all I look for." Interestingly, the need to find accessible parking dominates, even compared to traversal distance to the final destination. P6 observed, "Let planners know it doesn't have to be really close to the entrance, right? It doesn't. That's not so important." P9 echoed this but also emphasized the need for accessible routes from the parking location to the destination: "As long as I can park someplace and use my equipment, I don't mind going further, but other times without sidewalks and curb cuts, there's been times where... I couldn't park anywhere." Others mentioned the design of spaces themselves: For example, P2 emphasized that the directional location of the aisle is also important, saying "Most are on the right side, why not put them on the left side? Doesn't make any sense to me, except for if somebody's in a manual chair." These diverse perspectives emphasize the subjectivity of a preferred disability parking space.



Figure 3: Participants placed decals, signs, and cones on or near their vehicles to warn other drivers not to obstruct the space needed for them or their wheelchair to enter and exit the vehicle. Photos from participants P7 and P2.

Still, P8, who works as a disability advocate, stated that "the perfectly designed parking spot is truly an ADA spot, and I don't say that about everything that the ADA does; but in terms of parking, it actually did it right." This indicates that *improper implementation* of disability parking requirements is a key issue, not the regulation itself.

Adaptive strategies. Given inconsistencies in accessible parking, PwDs develop a variety of strategies to ensure their specific requirements are met. A common approach involves relying on the assistance of others, as exemplified by P2, who shared: "[my husband] would drop me off in front and go park and pick me up." Similarly, P4 mentioned "I have two brothers... who have a cute little minivan because of me", and P6 serves as the primary driver for her son. Beyond human assistance, multiple participants carried a sign or a parking cone to delineate their space or warn others from parking too close, as illustrated in Figure 3. Some reluctantly resort to parking over multiple spaces, acknowledging its inconvenience and potential to frustrate others: P7 mentioned "I don't like taking

two spots at an angle because that pisses people off”, and P3 echoed “I’ll park out over the white line... and then you come back and you can see people who got really annoyed, and it’s like, I’m sorry, I can’t help it.” These strategies, while effective, highlight the persistent need for PwDs to take extra measures simply to achieve access.

Burden of “wrong” and mental calculus for access. Moreover, employing these adaptive strategies often creates a sense of doing something “wrong” just to gain access. P2 recounted: “There’s been quite a few times when I’ve had to either create my own space or illegally park in a bank that was closed.” P2 described relying on facility staff to “find a spot and put cones in the spot next to me... but it’s a pain in the ass and it really, I wish it wasn’t like that” (P11). This tension between access vs. effort was a persistent backdrop and stressor influencing travel decisions: “I do math in my head about...how badly do I need to go to the store today?” (P11).

Public misuse. Beyond structural and implementation issues, the challenges of disability parking were exacerbated by members of the general public. Nearly every participant expressed frustration at drivers without disability parking placards illegally occupying disabled parking spots, reducing availability. P8 lamented that these spots are “always full of people just parking there for two minutes”, while P5 noted that “people park on the loading spaces all the time which blocks me from getting in a car completely.” The problem is further compounded by concerns related to confronting individuals, as P10 starkly emphasized: “Number one, it’s not my job to educate the public about disability parking, and number two, people carry firearms. I don’t want to get shot.”

Alternative transit. While public transit can be a promising option—and several participants reported using public transit in combination with other travel forms—similar to prior work [56, 65], we found that public services are not always available or problematic in other ways. For example, P7 stated, “I’ve never taken public transit in Seattle because it’s just impractical... it’s unpredictable to take public transit.” In contrast, P8 found success with park-and-ride options, noting “I’ve found that we drive to [the park and ride] with lovely accessible spots... and it takes me 30 seconds to take the train and get dropped off.” However, P9 highlighted key limitations: “[Public transit] only runs 8 to 5 and doesn’t go out of city limits. It’s so limited that I’m forced to have a vehicle.” P6 reflected on the ease of getting on/off accessible buses but pointed to surrounding challenges: “Being on the bus and getting on the bus is easy, but there’s all the stuff around it” (P6), referring to other factors like distance to the bus stop, uphill climbs, and weather conditions.

3.2.2 Desired Data and Technology Tool Needs. When asked about what information would help participants find and navigate to accessible spaces as well as opportunities for technology support tools to help (as part of our design probes), participants commonly emphasized detailed and up-to-date information about parking occupancy, location, or existence.

Comprehensive routing. Participants desired accessibility-aware routing not just to a disability parking space but an accessible route from their car to the entrance of their destination. P4 noted, “Sometimes you want to go somewhere and have no idea how you’re going to get there. That’s what I like about routing tools, they may not give you the exact [details], but they give you something to work with.” This was highlighted in participants’ responses to our probes (e.g.,

probes #2, #3, and #5): “that information [showing three spaces, one van space] is super helpful” (P6), “seeing [the path from the parking space] is helpful, because when I get out of the car, I don’t necessarily know whether there’s a flight of stairs or the accessible ramp is that way... just having that extra information, it can’t not help” (P8); “oh my gosh, these would be so incredibly helpful. It’s kind of... knowing what you’re up for” (P10).

Value of dynamic data. Despite the existence of some current tools that provide static information (e.g., Google Maps shows accessible entrances, seating, and restrooms [11], though often incomplete), participants emphasized that dynamic, real-time data was paramount. P3 articulated this need clearly: “If you could say ‘All these are full; the next closest ADA spot is here,’ that would be pretty cool”. Similarly, P6 expressed a desire to “to know how full the parking lot is;” with P8 echoing, “that parking garage says it’s full, but there’s still four accessible spots left. Awesome, I would definitely use an app like that.” Participants also mentioned learning from non-traditional information sources, with one explaining, “I have received more help from disability influencers who want to take trips than probably any typical brochure” (P4). These desires for accurate, real-time information were particularly visible as positive reception to probes #3, #6, #9, and #13.

Accuracy and trust concerns. However, there are challenges to real-time information access, both in determining if a parking space exists or if it is full. Some participants expressed skepticism about the reliability of information systems, “[I don’t want] somebody else’s computer program getting me the information that I want when I can just do it myself” (P7), in response to probes #4, #5, and #6; “I don’t think I would believe that those numbers are up to date” (P1); “I don’t trust [a tool that shows allocation and location of disability parking]” (P7). The issue of parking violations further complicates the usefulness of information systems since violations decrease the reliability of information about available spaces. As P8 observed, “The Uber drivers will pull right into the two white stripes because they know they’re just going to be there for a minute.”

3.2.3 Design Recommendations for Disability Parking. Drawing on our formative study and ADA guidelines [87], we synthesize the following design recommendations for disability parking. A key overarching goal of this paper is to develop sensing solutions capable of tracking and ensuring compliance with these recommendations, which AccessParkCV in Study 2 partially addresses.

- (1) **Spacing:** Provide adequate, ADA mandated [87] clearance on one or both sides of the disability parking spot (including at least one access aisle).
- (2) **Signage:** Ensure there is visible signage for each space (as mandated by the ADA [87]), indicating disability parking designation and van accessibility. Signage should emphasize that access aisles are *not* for parking and the violation consequences.
- (3) **Maintenance:** Keep parking spaces well maintained, including refreshing fading paint, to ensure that their boundaries and accessible zones are clear and visible.
- (4) **Location:** Avoid placing disability parking on sloped areas or directly adjacent to high traffic zones.




Figure 4: Our dataset contains 11,762 objects in seven categories: *access_aisle*, *curbside*, *dp_no_aisle*, *dp_one_aisle*, *dp_two_aisle*, *one_aisle*, and *two_aisle*. The images show examples for each parking category, with *access_aisle* implicitly included in the latter four. We include typical and atypical examples, and examples of the resulting detections from our pipeline (with parking detections in cyan and aisle detections in orange).

- (5) **Connectivity:** Ensure that all disability parking has an accessible path to the entrance of the destination the lot is intended to serve.
- (6) **Compliance:** Enact policy and auditing steps to ensure that disability parking is allocated according to the ADA required minimums. This is an especially important consideration for public infrastructure projects that remove parking since for some, vehicle travel is the only viable option.

4 Study 2: Detecting and Characterizing Disability Parking from Aerial Images

Building on our Study 1 findings and established ADA guidelines [87], Study 2 introduces AccessParkCV, a novel computer vision (CV) pipeline designed for the scalable identification and characterization of disability parking spaces from aerial imagery. Our pipeline directly addresses the need for robust tools to assess compliance with accessibility standards and to provide actionable data, as highlighted by our participants. In particular, *location*, *spacing*, and *compliance*, from the six design recommendations above; however, our ability to track real-time occupancy or maintenance status is limited given aerial imagery update frequency (*i.e.*, several years between surveys), which we expand on in the Limitations of our work (Section 6.5).

Below, we begin by detailing our dataset and labeling methodology, followed by an explanation of AccessParkCV—composed of two key pieces, a *locator* and a *characterizer*—and three empirical evaluations: an overall performance evaluation across test sets from

three US regions, an examination of our width characterizer, and an investigation of model performance on regions outside our dataset.

4.1 Dataset

We present and describe our dataset, including orthorectification, the collection and labeling methodology, and the seven inference classes, which allow us to not just detect disability parking spaces but also categorize them into types (*e.g.*, one vs. two aisle spaces).

4.1.1 Orthorectification. AccessParkCV uses aerial images (*e.g.*, from a satellite or a plane flying overhead) with sufficiently high resolution ($\sim 7.6\text{cm/px}$) that parking slots are visually discernible. These images are typically collected by local or federal governments and released openly (*e.g.*, on ArcGIS Online) but can be distorted due to sensor hardware, terrain structure/elevation, and the curvature of the earth. Orthorectification removes these distortions [86, 92]. The orthoimage tile system [89] defines each tile by an *x*, *y*, and *z* axis (with *z* referring to zoom level, a value expressing the scale at which the Earth is rendered into tiles), letting us map each pixel in the image to a precise geospatial coordinate on the Earth’s surface.

4.1.2 Data Sampling and Preprocessing. We source and preprocess the images used for our final labeled dataset in three steps: (1) downloading the aerial imagery as $256 \times 256\text{px}$ tiles from three US cities, (2) stitching the tiles together to $512 \times 512\text{px}$, and (3) sampling the resultant 512×512 tiles using a small object detection model (separate from our main detection model in Section 4.2) to identify candidate parking areas. For the latter, we use a small seed dataset from a fourth city (Denver). We describe each step below.

In Step 1, we use *Tile2Net* [32] to download our open-source aerial imagery datasets formatted in the orthoimage tile system. To ensure geographic diversity and varied urban composition, we use three US areas: Seattle, WA [84], Washington D.C. [28], and Spring Hill, TN [48]. For Seattle and Spring Hill, the tiles are natively downloaded from the source as 256×256 images, and we adopt zoom level 20 (~15 cm/pixel), balancing a high enough resolution to visually resolve individual parking spaces and the availability of such data. For DC, the tiles are 512×512 natively, so for objects in the image to achieve the same scale as the other two cities, we use zoom level 19. Then, to unify the image sizes across all downloaded tiles, we resize the DC tiles from 512×512 to 256×256 using the *Python Pillow* library’s *resize* function [19] with the *Lanczos* algorithm [45]. For the bounding coordinates of our downloaded regions, see Table 2.

In Step 2, we stitch together the 256×256 tiles in a 2×2 grid, making 512×512 images, so each image covers a large enough land area to capture several parking spots in one image without being cut off, while balancing resource constraints. The output of Step 2 is over two million contiguous 512×512 aerial images across our three cities; however, this is too many to manually review and label.

Thus, in Step 3, we created a small seed dataset and trained an initial object detection model (a “hint” model) to guide us towards images that may contain parking. For hint model training, we use a separate open-source aerial image dataset, Denver [17], to avoid bias towards one of our in-dataset cities. We first manually reviewed Denver’s aerial imagery and identified 70 images with parking lots, which we label for disability parking using the same methodology described in section 4.1.3. These images are 1024×1024, rather than 512×512, as we processed this dataset before defining our broader



Figure 5: Examples of potential disability parking spaces that we decided *not* to label due to overly narrow aisles, spaces being blocked due to garbage containers or other obstructions, occlusions from trees or shadows, or generally poor definition images.



Figure 6: Examples where the research team struggled to determine disability parking and/or their bounds. (a) A space with a logo, which is completely blocked by a temporary structure. (b) A space with non-visible lines. (c) An atypical striped zone painted on top of what may be a parking space. (d) The only case where striping overlaps the logo, creating ambiguity about whether it is a parking space.

Region	Bounding Coordinates	Source Resolution	# images in dataset
Seattle	(47.9572, -122.4489) (47.4091, -122.1551)	7.62 cm/pixel	2,790
Wash. D.C.	(38.9979, -77.1179) (38.7962, -76.9008)	7.62 cm/pixel	1,801
Spring Hill	(35.7943, -87.0034) (35.6489, -86.8447)	Unknown	534
Total			5,125

Table 2: Each region represented in our dataset and (1) their bounding coordinates (top is top left coordinate, bottom is bottom right coordinate), (2) their source resolutions, and (3) the number of 512x512 images (tiles).

methodology. We kept this size as they are only used for the hint model, where accuracy requirements are lower.

From this small seed dataset, we train the hint model using YOLOv8 [37, 71] for 113 epochs with default settings, which we then run on all two million images from our main cities. Those 512×512 images with a detection confidence higher than 0.3 are marked “*may contain parking*”; all others as “*may not contain parking*.” We choose 0.3 experimentally: it enables enough flexibility for the hint model to provide false positives, thereby avoiding overfitting.

Finally, we randomly sample 15,500 images from these two pools (“*may contain parking*”, “*may not contain parking*”) to again avoid overfitting. From the former, we sample 9,500 images: 6,000 from Seattle, 2,750 DC, and 750 Spring Hill. From the latter, we sample 6,000 images: 4,000 from Seattle, 1,750 DC, and 250 Spring Hill. As the 15,500 images are still composed primarily of null images (images without any disability parking spaces), we further remove a proportion via manual review, so they do not dominate the final dataset. We randomly remove all but ~40% of the null images from the “*may contain parking*” pool, as these are the images that “tricked” the hint model, thereby increasing the robustness of our final dataset. In total, our final dataset includes 5,125 images: 3,065 null and 2,060 with at least one instance of disability parking.

4.1.3 Classes and Labeling. Inferring a disability parking spot requires an explicit definition of what constitutes disability parking. Though the ADA clearly defines *what* disability parking is and *how* it should look [70], the only required visual indicator is a vertical sign with the international symbol of access. However, such signage is *not* visible from aerial imagery. Although not required, many disability parking spaces also have the access symbol directly painted on the ground, which is discernible using CV methods and aerial imagery. If absent, we must rely on the other visible ADA requirements (e.g., access aisle adjacent to the parking space [87]). Thus, in our work, we optimistically consider every space with a visible access aisle as a disability parking candidate, which is also corroborated by our Study 1 participants who emphasized the *utility of an accessible space* whether or not explicitly marked. This biases our dataset towards false positives, which we return to in our Discussion (Section 6.2).

To capture such nuanced definitions of disability parking and to support diverse downstream applications, we define seven classes




Figure 7: Our definition of width, and a comparison of that definition for ground truth boxes (that can be arbitrary polygons) and predicted boxes (that are restricted to being angled rectangles). Note how the overlap between parking space and access aisle bounding boxes is not double counted.

in our dataset (Figure 4): (1) access aisle (*access_aisle*), (2) curbside (*curbside*), (3) disability parking with no aisles (*dp_no_aisle*), (4) disability parking with one aisle (*dp_one_aisle*), (5) disability parking with two aisles (*dp_two_aisle*), (6) spaces with one aisle (*one_aisle*), and (7) spaces with two aisles (*two_aisle*). *Access aisle* refers to visible non-parking zones adjacent to a parking space. *Curbside parking* denotes spaces that are along the curb of a street and are visibly marked as disability parking. Three classes (*dp_no_aisle*, *dp_one_aisle*, *dp_two_aisle*) are spaces that are visibly distinguishable as disability parking, generally via a painted logo, with zero, one, and two access aisles, respectively. The remaining classes (*one_aisle*, *two_aisle*) are spaces without any obvious indication of being disability parking but with adjacent access aisles.

We annotate each object in each image with polygons. Although parking spaces are *generally* consistently sized quadrilaterals, access aisles can be more variable in shape and size, such as triangles, semicircles, or arbitrary polygons. Therefore, to be considered an access aisle (and not simply a striped zone adjacent to a parking space), the aisle must be sufficiently wide and extend along the side of the parking space to a degree where it could aid access in and out of a vehicle. While this typically means along at least ~30% of the parking space edge, we deferred to the labelers’ visual judgment. We include non-labeled examples in Figure 5. Additionally, definitional gaps remain in instances of distortion or atypical parking implementation. In these cases, we again deferred to the labelers’ judgment (see Figure 6 for examples).

Two members of the research team manually labeled all selected images. Labelers met before, during, and after the labeling process to define and clarify the labeling procedure. Three quarters of the dataset were labeled by researcher A, with the remaining quarter by researcher B. Researcher A then quality checked and adjusted, if necessary, researcher B’s labels to ensure consistency. We did not perform an inter-rater reliability test given that all labels were checked for consistency by researcher A.

4.1.4 Summary. Our final dataset comprises 5,125 images, with 2,060 containing a total of 11,762 labeled objects across all classes (see Table 3 for a breakdown per class). The dataset exhibits significant class imbalance, with one-aisle classes dominating all others, which we discuss further in Section 6.

Class	Quantity in dataset	Percentage
<i>access_aisle</i>	4,693	39.90%
<i>curbside</i>	36	0.31%
<i>dp_no_aisle</i>	300	2.55%
<i>dp_one_aisle</i>	2,790	23.72%
<i>dp_two_aisle</i>	402	3.42%
<i>one_aisle</i>	3,424	29.11%
<i>two_aisle</i>	117	0.99%
Total	11,762	100.00%

Table 3: The class composition of our dataset.

4.2 Locating Disability Parking

We first describe our approach to *locate* disability parking spots followed by *characterizing* its physical dimensions. To detect and locate disability parking spaces in aerial imagery, we employ *multi-class object detection*. Given a 512×512 orthorectified aerial tile, the goal is to predict bounding boxes corresponding to one of six parking-related classes, excluding *access_aisle*, which we remove during preprocessing because, by definition, this class is associated with at least one adjacent parking space, and its presence is implicitly captured in the class definitions for adjacent labeled spaces, thus not requiring independent prediction.

To examine performance across models, we first benchmarked three candidate object detection models—YOLOv11 *large* [37], DINO with a ResNet50 backbone [16], and CoDETR with a SWIN-L backbone [93]—the last representing a state-of-the-art transformer-based detection architecture. All models were finetuned on our dataset using a 70/15/15 train/valid/test split. Training was performed on a high performance computing cluster (16 CPU cores, 128 GB RAM, NVIDIA L40S GPU with 48 GB VRAM). Batch sizes were chosen based on model size and hardware constraints: YOLOv11 was trained with a batch size of 16, while the other models used a batch size of 2. Fine-tuning was initialized from publicly released pre-trained weights from each model’s official repository (also available in our Hugging Face repository³). All other hyperparameters, including data augmentation settings, followed default configurations in their respective repositories. Full details, including training time, learning rates, and augmentation strategies, are available in

³<https://huggingface.co/datasets/makeabilitylab/disabilityparking>




Figure 8: Results from the object detection model on the test set, shown as a confusion matrix, and common examples of the four largest misclassification categories. Predicted bounding boxes are highlighted in cyan, and ground truth bounding boxes (for false negatives) in green. We note certain trends in the misclassifications, including occlusion, smudges, or atypical spaces.

the appendix (Section A.4) and our GitHub repository⁴. For each model, we selected the best performing checkpoint based on validation set performance, prioritizing high recall to minimize false negatives. This decision reflects the application context, *i.e.*, false positives are easier to manually validate whereas false negatives result in lost detections that are more costly to recover *post hoc*.

We then evaluate each candidate model on the test set. CoDETR demonstrated the best performance, achieving the highest micro-F1-score (0.89) and recall (0.94)—see Section A.4 for detailed metrics for each model. We thus selected CoDETR for all subsequent analyses. We reuse the results of this experiment for the in-depth evaluation of our locator model’s performance (see subsection 4.5.1), as the application context is the same.

4.3 Characterizing Disability Parking

In addition to identifying and locating a disability parking space, we also aim to infer key characteristics such as size. Drawing on our Study 1 results, we found that access aisles and the total number of accessible spaces were key priorities for PwDs. For geometric characterization, we fit *oriented bounding boxes* (OBBs) to each inferred object instance. The OBBs are aligned with the dominant axis of the space, enabling accurate estimations of physical dimensions—particularly width—regardless of the space’s orientation in the image. Given their rectilinear structure, parking spaces do not require the granularity of segmentation masks, making OBBs a more efficient and appropriate choice.

To train the OBB model, we again make use of our labeled dataset, however with some preprocessing to align the input data with the OBB model’s task. Here, we generate 100×100 crops centered

around the centroid of each parking object. In each crop, we extract labels for all objects in the image, adjusting coordinate positions and shapes accordingly (*i.e.*, aligning the coordinates with the new, 100×100 image positions from the old 512×512 image positions). We maintain the same train/valid/test splits as above: any object used for training the locator model was used for training the OBB model. For the test set, we remove any objects close to the image’s edge (*i.e.*, that would require padding to reach 100×100), which ensures that the model is tested only on objects where the parking space is whole and not cut-off (for rationale and validity justification, see Section 4.4 and Section A.1). Our final object-wise train/valid/test splits are 4983/1015/706.

For the OBB model, we use YOLOv11x [37]. We did not trial other infrastructures, unlike for the locator model, due to YOLOv11x showing positive results at the outset (see subsection 4.5.2). We train with the same compute as above, for 200 epochs, with a batch size of 16, a learning rate of 0.01 with 0.937 momentum and a weight decay of 5×10^{-4} . We use all default settings for augmentation, including hue scaling, translation, and horizontal flipping. Details are in the *Ultralytics* documentation.⁵ Further, all specific training values are in our GitHub repository.

4.4 AccessParkCV: An End-to-end CV Pipeline

Finally, we compose the preceding two models—the locator and the characterizer—into an end-to-end pipeline called *AccessParkCV* (Figure 1). Given input bounding coordinates for a region, AccessParkCV locates, classifies, and calculates the width of disability parking spaces within that region. The two models function in the same manner as described in Sections 4.2 and 4.3, with the locator

⁴<https://github.com/makeabilitylab/AccessParkCV>

⁵<https://docs.ultralytics.com/usage/cfg/#augmentation-settings>




Figure 9: Results from the OBB model on the test set, shown as a confusion matrix, and examples of access aisle false negatives and positives. Predicted OBBs are highlighted in cyan for parking spaces and in orange for access aisles. Ground truth polygons are not shown for false negatives. Once again, we note misclassification trends for occluded and atypical parking spaces.

	Cnt	Researcher A vs. Researcher B				Researcher A vs. Model			
		Mean Difference		SD		Mean Difference		SD	
		Pixels	%	Pixels	%	Pixels	%	Pixels	%
<i>null</i>	10	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
<i>curbside</i>	4	-1.03	-3.92	1.33	5.75	3.44	15.75	4.25	20.11
<i>dp_no_aisle</i>	12	1.94	7.37	4.84	18.22	2.47	9.11	8.63	30.74
<i>dp_one_aisle</i>	174	1.03	2.28	4.13	9.57	0.81	2.30	5.98	13.01
<i>dp_two_aisle</i>	24	-0.03	0.73	7.26	9.43	0.5	1.05	2.67	4.24
<i>one_aisle</i>	139	0.16	0.91	6.80	14.47	2.98	7.65	7.62	19.36
<i>two_aisle</i>	3	1.85	2.01	4.40	6.21	1.18	1.80	1.72	2.56
Total	366	0.65	1.77	5.57	12.05	2.01	5.40	6.91	17.52

Table 4: Results of the error in estimation of width of parking spaces. We compare the difference between two human annotators (Researcher A vs. Researcher B), and the difference between a human annotator and the model (Researcher A vs. Model). Positive indicates overestimation (compared to Researcher A), and negative underestimation. The model sees both larger mean error, and standard deviation of error. Colors indicate magnitude of values using ColorBrewer sequential blue palette.

model passing 100×100 crops around the centroid of each detected parking object to the OBB model. Both models use a confidence threshold of 0.3, which were determined experimentally.

To form the final pipeline output, we apply additional logic to the resulting characterizer inferences. Firstly, the OBB model detects all instances of parking spaces and access aisles in an image, however, we are interested in only the center space and its neighboring aisles (as the other spaces have their own separate, independent crops). Thus, we select only the parking space with the highest confidence score that contains the center point of the 100×100 crop. Then, we select aisles that are along at least 40% of the longest edges of the parking space (*i.e.*, the sides), defined by being within 20 pixels of each other. Second, we derive the width from the resultant OBB detections of the space and its aisles. To calculate the width of a rectangular parking space from its OBB is trivial; however, as an access aisle's OBB's axis can be at an angle relative to the parking space's, what the accessible width of the aisle is is less clear. Therefore, we define the access aisle width as the length of a vector,

starting at the midpoint and perpendicular to the parking space edge, pointing away from the centroid, to the farthest intersection point to an access aisle. If there is no intersection point, the width is considered zero. Finally, we calculate the total accessible width for a space by summing the width of the parking space and the width of any aisles. Since the width of the aisle is defined relative to the parking space, any overlap is *not* double counted (see Figure 7).

To format the output, each detected object is georeferenced to a position on Earth, and the results are output to both a JSON file and a shapefile for GIS mapping software.

4.5 Evaluation and Results

We individually evaluate both parts of the pipeline, the *locator* and *characterizer*, examining their effectiveness on out-of-sample regions and diverse resolutions, and compare the pipeline's parking width error to the discrepancy between human annotations.

Region	Archetype	True Positives	Ground Truth	Recall (%)	FN	FP	Mean Width Error (m)
Roosevelt (Seattle)	Residential	43	49	87.8	6	4	-0.20 (SD=1.63)
Northgate (Seattle)	Transit Center	164	174	94.3	10	39	
USPS Main Office (DC)	Urban Commercial	117	122	95.9	5	11	
Audi Field (DC)	Sports Stadium	41	43	95.3	2	8	
Koreatown (Los Angeles)	Commercial/Residential	79	123	64.2	44	23	-0.29 (SD=1.63)
Del Amo Fashion Center (Torrance, LA)	Urban Commercial	163	227	71.8	64	30	
Shopping Mall (Waltham, MA)	Suburban Commercial	132	164	80.5	32	20	0.53 (SD=1.16)
Total	All Regions	739	902	81.9	163	135	-0.10 (SD=1.58)

Table 5: The seven one km² regions we evaluated our pipeline on. We do not distinguish between categories for parking spaces detected, only if they were detected at all. Mean Width Error is shown, in meters, aggregated across three categories: regions represented in our training data (Seattle, D.C.), Los Angeles, and Massachusetts. Again, positive number indicates overestimation, and negative underestimation. We see that the model consistently detects a higher proportion of objects, and with a lower error for width, for the regions represented in the training set.

4.5.1 Detection Model. To examine overall performance, we evaluated the CoDETR disability parking locator model on the test set (717 images across Seattle, DC, and Spring Hill). We matched detected objects to ground truth using the *Hungarian* method for bipartite matching [41], with an *Intersection Over Union* (IOU) threshold exceeding 0.5 to ensure optimal correspondence. For all aggregated performance statistics (precision, recall, and F1 score), we use the micro-average [80] to capture the holistic performance of the model across classes. However, given the dataset imbalance, we encourage readers to observe the per-class statistics for a comprehensive view. Overall, the CoDETR disability parking locator model performs well with an overall precision of 0.85, recall of 0.94, and F1 score of 0.89—see class-specific statistics in Table 6.

Qualitatively examining errors. To better understand model performance, we manually examined all misclassifications, false negatives (where a parking space that exists is not detected), and false positives (where null pixels are detected as a parking space)—Figure 8. For misclassifications ($N=31$), the most common error is *one_aisle* misclassified as *dp_one_aisle* or vice versa, perhaps since the only difference from the aerial view is a visual indicator (e.g., painted logo). In fact, common cases of *dp_one_aisle* being predicted as *one_aisle* are when the logo is faded, smudged, or partially occluded. Inversely, in many cases, a space that is occluded or has some mark near the bottom, is misclassified as *dp_one_aisle*.

In examining *false negatives* ($N=66$), common patterns are paint fade (thus lower contrast with the pavement), occlusion (e.g., from trees or shadows), or an access aisle not being detected (thus, not identifying the space as a *one_aisle*). For *false positives* ($N=130$), common patterns are access aisles that were blocked by an object, spaces or aisles that were not labeled (due to occlusion or distortion), or spaces that looked similar to a parking space with a logo (e.g., an empty area with an adjacent striped zone, or a number at the bottom). Interestingly, in the former two cases, perhaps a more lenient labeling scheme may have identified the spaces, making the false positives correct detections.

Finally, single parking spaces detected multiple times with different classifications (with one being correct) contributes to the error rate. This is avoided for the full pipeline by filtering overlapping detections by confidence score.

4.5.2 OBB Model. Beyond detections, we also examine the other key component of our pipeline, the *characterizer* model. Here, we evaluate it as we did the detection model: on the (100×100) test set, matching objects with the Hungarian method and an IOU of 0.5. However, we compute only whether the space was detected as an access aisle or parking space and the presence of false negatives/positives, because the OBB model is solely used for width estimation, with class prediction handled by the detection model.

Overall, similar to our detection model, we see high performance, with only 29 false negatives and 31 false positives, of 1,426 detected objects—see Figure 9. Our space characterizer perfectly distinguishes access aisles from parking spaces. However, access aisles experience some false negatives (3.5% not detected), often due to occluded, faded, or smudged aisles, and false positives (4.3% of total access aisle detections) that commonly result from duplicate detections of non-rectangular aisles, or the misidentification of neighboring spaces or striped zones.

Class	# in test set	Precision	Recall	F1
<i>curbside</i>	10	0.50	0.40	0.44
<i>dp_no_aisle</i>	45	0.87	0.87	0.87
<i>dp_one_aisle</i>	367	0.89	0.96	0.93
<i>dp_two_aisle</i>	62	0.89	0.92	0.90
<i>one_aisle</i>	553	0.82	0.88	0.85
<i>two_aisle</i>	10	0.75	0.55	0.64
Total	1,047	0.85	0.94	0.89

Table 6: Results of our locator model, by class. Note that for the Total results, misclassifications, or instances where an object of one class is predicted as another, are considered false positives for the statistics calculation. Performance is high for classes with high representation, but lower for classes that are underrepresented (*curbside*, *two_aisle*). Total, aggregated statistics are formed from the micro-average of the class level statistics.



Figure 10: Example FNs and FPs for each evaluation region with ground truth annotations and predictions shown. We see similar error patterns as before; however, LA disability parking demarcations are often blue (rather than white or yellow) and more faded, resulting in lower contrast with the pavement. Waltham data is at a much lower resolution, resulting in blurriness.

Evaluating width characterization. In detecting parking spots and their aisles in a cropped image as an OBB, the key responsibility of the characterizer is to quantify parking space width (including access aisles), via the dimensions of the OBBs. To evaluate width characterization, we compared our model’s error to the error between two human researchers’ labels. Researchers A and B labeled the same, randomly sampled 356 object subset of the test set (plus 10 randomly sampled false positives from the detection model evaluation, as null trials), and we calculated the differences in width from their respective labels. As location data (which is required to derive real-world distance) was lost during the test set preprocessing, all measurements for evaluation are in pixels, not real-world units.

Overall, we find that the model performs well: only slightly worse than the humans with an average error of 5.40% vs. 1.77% and overall standard deviation of 17.52% vs. 12.05%, though performance depends on class—see Table 4 and Section A.3.

4.5.3 Cross City Performance. Finally, to investigate how well AccessParkCV generalizes to regions outside of those used to train the models, we selected seven additional one km² regions, four from our original cities that were not in our training set (two from Seattle; two from DC) as well as two neighborhoods in Los Angeles [46] and one from Massachusetts [69]. These selections were based on geographic and zone type variation in addition to data availability in Tile2Net. Aerial imagery from Massachusetts has a much lower resolution (15 cm/px vs. 7.6 cm/px), which also enabled us to examine performance on lower resolutions. For ground truth, one member of the research team manually labeled all parking spaces and aisles with the same seven classes as before, and we release these additional seven regions and labels to our HuggingFace. We evaluate the entire pipeline: the ability to detect parking spaces, and the accuracy of the pipeline’s estimation of their widths. Each detected parking space was matched with a ground truth label for evaluation, again with the Hungarian bipartite matching method based on an IOU of 0.3. We report results for *total objects detected*,

without distinguishing between parking classes, due to the low cross-class sample count.

Overall, as expected, detector model performance drops but still performs fairly well, particularly on the neighborhoods drawn from the same cities as the training set (Seattle and DC) with average recall above 90%; however, performance in LA and Waltham, MA drops significantly: to ~69% in LA and 81% in MA, which suggests some amount of per-city training data may be necessary. To advance understanding of model performance, we once again qualitatively analyzed errors in each region (Figure 10); we find similar error patterns as reported in Sections 4.5.1 and 4.5.2. However, we note two notable distinctions with LA parking data: a prevalence of blue paint for disability parking (vs. yellow or white), resulting in reduced contrast with the surrounding pavement, and more frequent and severe paint fading, possibly attributable to intense sunlight. The human labeler also found LA data more difficult to label than other regions given these reasons, which may have contributed to the overall error rate (given that the labels themselves may not be 100% accurate). These patterns may also have contributed to the width error rate, as parking spaces where the borders are not clearly defined are more imprecisely detected.

5 Case Study Applications

To help demonstrate the utility of AccessParkCV as well as our findings from Study 1, we present two case study applications: a personalized disability parking app (a lo-fi mockup) and an interactive urban analytics visualization (a mid-fi prototype).

5.1 Personalized Parking Search and Filter

As highlighted in Study 1, individuals with disabilities (PwDs) highly value pre-trip information regarding location accessibility, as it significantly reduces travel time and anxiety. Regarding parking, the availability and size of spaces were frequently cited as crucial factors. Leveraging classifications and width data from




Figure 11: Two applications utilizing the results from our detection tool: (a) a personalized parking search and filter application, and (b) an urban analytics visualization displaying number of disability parking spaces per census tract in Seattle.

AccessParkCV, we designed a lo-fi mapping application (Figure 11a) that identifies accessible parking options within a given area. The tool allows users to refine their search based on: the number of adjacent aisles, the total width of the parking space, and its proximity to their destination. For example, a user operating a ramp-equipped van, who prioritizes space over immediate proximity to an entrance, can effectively filter for suitable parking spots and plan their travel accordingly. While the application does not provide real-time occupancy data (a limitation of our pipeline), it empowers users to locate appropriately sized parking spaces that meet their specific vehicle and accessibility requirements, thereby facilitating the identification of suitable alternatives if a preferred space is occupied.

5.2 Disability Parking Analytics

As stipulated by the Americans with Disabilities Act (ADA), 4-8% of total public parking must be designated as accessible parking [43]. To help disability advocates, government officials, and even end-users themselves understand the distribution and density of disability parking in their communities, we designed and developed a mid-fi interactive disability parking visual analytics tool. Illustrated in Figure 11b, the tool enables users to conduct in-depth investigations into disability parking trends, which can be further disaggregated by community, zoning regulations, and business concentration. Such visual analytics provide a data-driven foundation for local governments to formulate and revise parking and zoning policies. Furthermore, these visualizations could be integrated into other regional accessibility metrics, such as *AccessScore* [51], empowering individuals to make informed decisions regarding travel and residential choices based on accessibility considerations.

6 Discussion

In this paper, we explored the nuances of disability parking in the US, sharing formative findings on its perceived benefits, uses, and drawbacks. We also developed a new algorithmic pipeline to detect and assess disability parking spots from aerial imagery, created a novel disability parking dataset, and showcased two practical applications of our methods. Below, we reflect the implications of our work, our dataset contributions, the future of CV applied to disability parking, and key limitations and future work.

6.1 Implications for Design

Drawing on our study findings, we synthesize five key recommendations for the design of interactive apps for disability parking.

- **Beyond location.** While showing the location of disability parking spots on a map is a critical accessibility advancement, it is insufficient to support diverse mobility needs, vehicle types, and the various requirements for access aisles. Future apps should leverage AccessParkCV’s ability to infer characteristics like the presence and width of aisles and surface related metadata to the user (similar to Figure 11a).
- **Personalized filtering.** Because preferences for disability parking vary significantly—as a function of ambulation ability, use of mobility aid, vehicle type—apps should empower users to filter for their specific accessibility criteria (e.g., minimum width, presence and orientation of access aisles).
- **Comprehensive routing.** While finding an accessible parking spot is critical so too is the journey from the parking to the building entrance. Future tools should integrate assessments of the “out-of-car” experience, including the presence of sidewalks, pathway accessibility, and distances.

- **Trust and transparency.** Our participants expressed skepticism about the reliability of disability parking information, partly due to issues like illegal occupancy and inconsistent implementation of regulations. To build trust, apps should surface information sources and, if possible, allow users to actually look at raw image sources (for verification).
- **Real-time occupancy.** The ideal app would also provide real-time availability data, which is beyond the capabilities of our techniques. Future work should explore hybrid solutions connected to real-time sensors.
- **Human-in-the-loop.** While our CV techniques perform well, they are imperfect. Apps should empower users to provide *in situ* insights on the quality of a parking spot and provide feedback on CV detections via their lived experience.

6.2 Dataset Contribution

A key contribution of our work is the first open-source, human-labeled dataset of disability parking and access aisles from aerial orthoimages. This dataset comprises 11,762 labeled objects across 5,125 images from three distinct US regions: Seattle, Washington D.C., and Spring Hill, TN. By stratifying "accessible parking" into seven distinct classes, including those with varying numbers of aisles and different visual indicators (*e.g.*, painted logos), we enable fine-grained analysis and customization for both macro-scale urban planning and micro-scale personalization. Our dataset includes a significant number of null examples (3,065 images) from the same regions, allowing researchers to filter null proportions appropriate for their desired use case. We publish the dataset in the commonly used COCO format [52] for convenience.

While the dataset exhibits class imbalance, with one-aisle classes dominating, this imbalance likely reflects real-world conditions, as ADA guidelines do not mandate two aisles for all accessible spaces: only "van accessible" spaces are required to have two aisles and curbside disability parking is not mentioned [87]. In our locator model, we find that this imbalance does not significantly affect results (beyond the curbside class, which only has 36 samples). However, future dataset expansions should include efforts to reduce the imbalance.

6.3 Disability Parking CV: A New Benchmark

Our novel computer vision pipeline, composed of a locator and a characterizer model, sets a new benchmark for scalably detecting and characterizing disability parking from aerial imagery. The pipeline achieves a micro-F1 score of 0.89 for detection, demonstrating high overall performance, particularly for classes with strong representation in the dataset. Importantly, our characterizer model estimates the width of parking spaces with an average error of only 5.4%, which directly helps assess compliance with accessibility standards and to provide actionable data to PwDs, a need strongly voiced by our interview participants.

The open-sourcing of our pipeline code and experiments alongside our dataset helps advance research in CV-based analyses of accessible urban infrastructure [22, 32, 42, 47, 55] and provides foundational benchmarks for this emerging area. While our pipeline excels at identifying clearly defined and visible disability parking

spaces, it struggles with "atypical" parking spaces—just as our human annotators did—due to factors such as faded paint, blurriness, and unexpected design variation. Our decision to prioritize lowering false negatives over false positives reflects the application context, where undetected accessible parking is more costly than manually verifying a false positive. This approach aligns with the principle that humans, particularly those with lived experience, should remain in the loop for validating and verifying algorithmic outputs related to accessible urban infrastructure.

6.4 Scale and Generalizability

Our longterm, overarching goal is to accurately detect, characterize, and track every disability parking spot in the US, if not the world. While our experiments demonstrate high performance, particularly for regions within our training set, there are two key limiting factors: (1) in subsection 4.5.3, we show that performance drops without some per-city training—future work should more comprehensively examine performance as a function of training set size and diversity; (2) the availability of high-resolution aerial orthoimagery, which are often created by public agencies like the US Geological Survey [83] or local governments (the Seattle [84], D.C. [28], Massachusetts [69], and LA [46] datasets are all provided on open gov city websites). Continued expansion of the dataset across a higher diversity of regions, resolutions, and zoom levels would bolster model generalizability and potentially reduce the need for extensive fine-tuning on deployment regions.

6.5 Limitations

Our work has several primary limitations. First, in Study 1, our participant pool was almost entirely composed of wheelchair users. Greater diversity of disability type and mobility aid would enhance the generalizability of our qualitative findings. Second, aerial imagery inherently limits the scope of detectable parking, as indoor parking facilities (*e.g.*, in parking garages) are not visible. Third and relatedly, not all disability parking can be identified from aerial imagery, including in instances where a space lacks an access aisle and is only visually denoted by a vertical sign. Furthermore, environmental factors like occlusion from trees or shadows can impact accuracy. Fourth, our approach is reliant on the availability and quality of high-resolution orthoimagery. Even in cases where such data is available, aberrations within the dataset can affect the results (*e.g.*, in DC, the entire National Mall is blurred, presumably for security reasons). Similarly, real-time applications are also not possible using only orthoimage datasets, given they are typically only updated after several years, if ever. Finally, our definitions of "access aisle" and "width" may not be appropriate in all scenarios. For example, in our dataset we choose to label striped, triangular wedges between two parking spaces as an access aisle; however, if two long, parked cars were to park in the adjacent spaces, there would be no path out from the aisle.

6.6 Future Work

Future work should address the limitations outlined above. Employing multiple aerial orthoimage datasets of the same region, captured at different times, could help mitigate issues caused by

parked cars occluding disability parking features. Integrating non-aerial imagery sources, such as Google Street View, could enable the identification of spaces marked only by vertical signage and potentially provide visibility into indoor parking areas or spaces entirely occluded from a bird's-eye view. Further expansion of the dataset to include a higher diversity of regions, resolutions, and zoom levels would likely bolster model generalizability and may reduce the need for extensive fine-tuning on new deployment regions.

Beyond addressing limitations, several potential extensions could improve AccessParkCV. First, additional classification models or ensemble voting approaches [25] could be used to reduce the number of false positives. Second, additional characterizer models could be added along the pipeline that accept the same image input shape as the OBB model, such as quantifying a parking space's degree of paint fade or if it is angled. Finally, incorporation of contextual information, such as obstructions on the path to the entrance or distance from traffic, would allow us to make more holistic claims on the accessibility of a given parking space.

We envisioned two, non-exhaustive use cases for the data produced by our pipeline. More applications and tools could also benefit from the disability parking detections. Disability parking data could be combined with sidewalk data [51, 74], which could then be integrated into a personalized routing app, or to generate a community level accessibility score. Further, a tool could be created that audits and notifies individual businesses if they are compliant with ADA parking guidelines [87] based on quantity and width.

7 Conclusion

Through mixed-methods research, including an interview study of 11 PwDs and the development and evaluation of a CV pipeline that can locate and characterize disability parking spaces, we contribute:

- (1) Design and policy guidelines for the implementation and maintenance of disability parking.
- (2) An open-source dataset of 11,762 labeled disability parking spaces and access aisles that can be used for classification, detection, segmentation, or other analysis tasks.
- (3) A CV tool that locates, classifies, and characterizes the width of disability parking from aerial orthoimagery.
- (4) Two envisioned applications that leverage the preceding contributions to enable large-scale disability parking auditing and individual-scale accessibility filtering.

Our work highlights the potential for large-scale, computer vision disability parking detection and characterization tools to improve both the implementation of disability parking at large and for the data from such tools to bolster government-level analysis and individual-level customization. Our work showcases both the need for such investigative and personalized tools and the ability for computer vision to serve as a vital step towards that end.

Acknowledgments

We thank all study participants as well as Adela Mu for urban planning and GIS guidance, Daniel Campos Zamora and Charles DeLorey for their local expertise on our test regions, academic writing advisor Sandy Kaplan, UW CREATE, and CREATE community engagement manager Kathleen Quin Voss. This work was supported

by NSF grant #2411222 and the US DOT National Center for Understanding Future Travel Behavior and Demand (#69A3552344815 and #69A3552348320). Figure 1 icons by [14, 29].

References

- [1] 1968. Architectural barriers act of 1968. <https://www.access-board.gov/aba/>
- [2] 1973. Rehabilitation act of 1973, section 504. <https://www.dol.gov/agencies/oasam/centers-offices/civil-rights-center/statutes/section-504-rehabilitation-act-of-1973>
- [3] 2024. Disability parking scheme | Department of Social Services. <https://www.dss.gov.au/disability-support-services/disability-parking-scheme>
- [4] National Geospatial-Intelligence Agency. 2017. *Map Projections for GEOINT Content, Products, and Applications*. Technical Report. <https://nsgreg.nga.mil/doc/view?i=4478>
- [5] Dragan Ahmetovic, Roberto Manduchi, James M. Coughlan, and Sergio Mascetti. 2017. Mind Your Crossings: Mining GIS Imagery for Crosswalk Localization. *ACM Trans. Access. Comput.* 9, 4 (April 2017), 11:1–11:25. doi:10.1145/3046790
- [6] Adrian Albert, Jasleen Kaur, and Marta C. Gonzalez. 2017. Using Convolutional Networks and Satellite Imagery to Identify Patterns in Urban Environments at a Large Scale. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '17)*. Association for Computing Machinery, New York, NY, USA, 528–531. doi:10.1145/3097983.3098070
- [7] Giuseppe Amato, Fabio Carrara, Fabrizio Falchi, Claudio Gennaro, Carlo Meghini, and Claudio Vairo. 2017. Deep learning for decentralized parking lot occupancy detection. *Expert Systems with Applications* 72 (April 2017), 327–334. doi:10.1016/j.eswa.2016.10.055
- [8] Mujahid I Ashqer and Marwan Bikkdash. 2019. Parking Lot Space Detection Based On Image Processing. In *2019 SoutheastCon*. 1–6. doi:10.1109/SoutheastCon42311.2019.9020584 ISSN: 1558-058X.
- [9] Tessa Berry, Nicholas Dronen, Brett Jackson, and Ian Endres. 2019. Parking Lot Instance Segmentation from Satellite Imagery through Associative Embeddings. In *Proceedings of the 27th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (SIGSPATIAL '19)*. Association for Computing Machinery, New York, NY, USA, 528–531. doi:10.1145/3347146.3359364
- [10] Nazia Bibi, Muhammad Nadeem Majid, Hassan Dawood, and Ping Guo. 2017. Automatic Parking Space Detection System. In *2017 2nd International Conference on Multimedia and Image Processing (ICMIP)*. 11–15. doi:10.1109/ICMIP.2017.4
- [11] Sasha Blair-Goldensohn. 2020. Find wheelchair accessible places with Google Maps. <https://blog.google/products/maps/wheelchair-accessible-places-google-maps/>
- [12] Erin Brady and Jeffrey P. Bigham. 2015. Crowdsourcing Accessibility: Human-Powered Access Technologies. *Foundations and Trends® in Human-Computer Interaction* 8, 4 (Nov. 2015), 273–372. doi:10.1561/11000000050 Publisher: Now Publishers, Inc..
- [13] Virginia Braun, , and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (Jan. 2006), 77–101. doi:10.1191/1478088706qp0630a Publisher: Routledge _eprint: <https://www.tandfonline.com/doi/pdf/10.1191/1478088706qp0630a>
- [14] bsd studio. 2024. Persons with disabilities Icon Set. <https://thenounproject.com/icon/person-with-an-amputation-6613450/>
- [15] Pietro Carnelli, Joy Yeh, Mahesh Sooriyabandara, and Aftab Khan. 2017. ParkUs: A Novel Vehicle Parking Detection System. *Proceedings of the AAAI Conference on Artificial Intelligence* 31, 2 (Feb. 2017), 4650–4656. doi:10.1609/aaai.v31i2.19090 Number: 2.
- [16] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. 2021. Emerging Properties in Self-Supervised Vision Transformers. 9650–9660. https://openaccess.thecvf.com/content/ICCV2021/html/Caron_Emerging_Properties_in_Self-Supervised_Vision_Transformers_ICCV_2021_paper
- [17] DRCOG Regional Data Catalog. 2021. Denver Regional Aerial Photography Project Tiles 2020. <https://data.drcog.org/dataset/denver-regional-aerial-photography-project-tiles-2020>
- [18] Daniel Cisek, M. Mahajan, Jedidiah Dale, Susan Pepper, Yuewei Lin, and Shinjae Yoo. 2017. A transfer learning approach to parking lot classification in aerial imagery. In *2017 New York Scientific Data Summit (NYSDS)*. 1–5. doi:10.1109/NYSDS.2017.8085049
- [19] Alex Clark. 2025. Pillow (PIL Fork) Documentation. <https://buildmedia.readthedocs.org/media/pdf/pillow/latest/pillow.pdf>
- [20] John W Creswell and Cheryl N Poth. 2016. *Qualitative inquiry and research design: Choosing among five approaches*. Sage publications.
- [21] Sébastien Drouyer. 2020. Parking Occupancy Estimation on PlanetScope Satellite Images. In *IGARSS 2020 - 2020 IEEE International Geoscience and Remote Sensing Symposium*. 1098–1101. doi:10.1109/IGARSS39084.2020.9323104 ISSN: 2153-7003.
- [22] Michael Duan, Sho Kiami, Logan Milandin, Johnson Kuang, Mikey Saugstad, Maryam Hosseini, and Jon E. Froehlich. 2022. Scaling crowd+AI sidewalk accessibility assessments: Initial experiments examining label quality and cross-city

- training on performance. In *Proceedings of the 24th international ACM SIGACCESS conference on computers and accessibility*. Athens, Greece. doi:10.1145/3517428.3550381 Number of pages: 6 tex.booktitleshort: Proceedings of ASSETS'22 2022.
- [23] V. Eleftheriou, E. Bakogiannis, I. Papadopoulos, A.Vassi, and C. Kyriakidis. 2025. Disabled People and "Mobility". Proposal for Infrastructures Intended to Prevent Parking in Disabled Spaces and Their Pilot Application in the City. In *Climate Crisis and Resilient Transportation Systems*, Eftihia G. Nathanail, Nikolaos Gavanis, and Evangelos Adamos (Eds.). Springer Nature Switzerland, Cham, 199–211. doi:10.1007/978-3-031-82714-3_14
- [24] Rifqi Muhammad Fikri and Mintae Hwang. 2019. Smart Parking Area Management System for the Disabled using IoT and Mobile Application. In *2019 IEEE International Conference on Internet of Things and Intelligence System (IoTals)*. 172–176. doi:10.1109/IoTals47347.2019.8980454
- [25] M. A. Ganaie, Minghui Hu, A. K. Malik, M. Tanveer, and P. N. Suganthan. 2022. Ensemble deep learning: A review. *Engineering Applications of Artificial Intelligence* 115 (Oct. 2022), 105151. doi:10.1016/j.engappai.2022.105151
- [26] R A J M Gining, S S M. Fauzi, A F A Hadi, M Z A Razak, and M N F Jamaluddin. 2018. Design and Development of Disabled Parking System for Smart City. *Journal of Physics: Conference Series* 1019, 1 (June 2018), 012016. doi:10.1088/1742-6596/1019/1/012016 Publisher: IOP Publishing.
- [27] Ratko Grbić and Brando Koch. 2023. Automatic vision-based parking slot detection and occupancy classification. *Expert Systems with Applications* 225 (Sept. 2023), 120147. doi:10.1016/j.eswa.2023.120147
- [28] DC GIS Group. 2021. Aerial Photography (Orthophoto) - 2021. <https://opendata.dc.gov/datasets/DCGIS::aerial-photography-orthophoto-2021/about>
- [29] Hadi. 2016. People Who Use a Wheelchair Icon Set. <https://thenounproject.com/browse/collection-icon/people-who-use-a-wheelchair-18080/>
- [30] Kotaro Hara, Christine Chan, and Jon E. Froehlich. 2016. The Design of Assistive Location-based Technologies for People with Ambulatory Disabilities: A Formative Study. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (CHI '16)*. Association for Computing Machinery, New York, NY, USA, 1757–1768. doi:10.1145/2858036.2858315
- [31] C. Henry, J. Hellekes, N. Merkle, S. M. Azimi, and F. Kurz. 2021. CITYWIDE ESTIMATION OF PARKING SPACE USING AERIAL IMAGERY AND OSM DATA FUSION WITH DEEP LEARNING AND FINE-GRAINED ANNOTATION. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences XLIII-B2-2021* (June 2021), 479–485. doi:10.5194/isprs-archives-XLIII-B2-2021-479-2021 Conference Name: XXIV ISPRS Congress <q>Imaging today, foreseeing tomorrow</q>. Commission II - 2021 edition, 5–9 July 2021 Publisher: Copernicus GmbH.
- [32] Maryam Hosseini, Andres Sevtsuk, Fabio Miranda, Roberto M. Cesar, and Claudio T. Silva. 2023. Mapping the walk: A scalable computer vision approach for generating sidewalk network datasets from aerial imagery. *Computers, Environment and Urban Systems* 101 (April 2023), 101950. doi:10.1016/j.compenvurbsys.2023.101950
- [33] Guang-Li Huang, Jing He, and Rui Wang. 2020. An Inclined Bounding Box Method for Detecting Parking Spaces in Open Parking Lots using Aerial Images. In *2020 IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT)*. 17–24. doi:10.1109/WIIAT50758.2020.00008
- [34] Roberto E. Huerta, Fabiola D. Yépez, Diego F. Lozano-García, Víctor H. Guerra Cobán, Adrián L. Ferriño Fierro, Héctor de León Gómez, Ricardo A. Cavazos González, and Adriana Vargas-Martínez. 2021. Mapping Urban Green Spaces at the Metropolitan Level Using Very High Resolution Satellite Imagery and Deep Learning Techniques for Semantic Segmentation. *Remote Sensing* 13, 11 (Jan. 2021), 2031. doi:10.3390/rs13112031 Number: 11 Publisher: Multidisciplinary Digital Publishing Institute.
- [35] AXS Lab Inc. 2025. AXSMap. <https://www.axsmap.com/>
- [36] Mary Ann Jackson and Ralph Green. 2015. Accessible Parking: are users voices heard within the Built Environment sphere?. In *49th International Conference of the Architectural Science Association 2015*.
- [37] Glenn Jocher, Ayush Chaurasia, and Jing Qiu. 2023. Ultralytics YOLOv8. <https://github.com/ultralytics/ultralytics>
- [38] Nazir KAMALDIN, Susan KEE, Songwei KONG, Chengkai LEE, Huiguang LIANG, Alisha SAINI, Hwee-Pink TAN, and Hwee-Xian TAN. 2019. SmartBFA: A Passive Crowdsourcing System for Point-to-Point Barrier-Free Access. In *2019 IEEE 44th Conference on Local Computer Networks (LCN)*. 34–41. doi:10.1109/LCN44214.2019.8990857 ISSN: 0742-1303.
- [39] Gou Koutaki, Takamochi Minamoto, and Keiichi Uchimura. 2016. Extraction of parking lot structure from aerial image in urban areas. 12 (Jan. 2016), 371–383.
- [40] Nikola Kranjčić, Damir Medak, Robert Župan, and Milan Rezo. 2019. Machine Learning Methods for Classification of the Green Infrastructure in City Areas. *ISPRS International Journal of Geo-Information* 8, 10 (Oct. 2019), 463. doi:10.3390/ijgi8100463 Number: 10 Publisher: Multidisciplinary Digital Publishing Institute.
- [41] H. W. Kuhn. 1955. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly* 2, 1-2 (1955), 83–97. doi:10.1002/nav.3800020109_eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/nav.3800020109>
- [42] Minchu Kulkarni, Chu Li, Jaye Jungmin Ahn, Katrina Oi Yau Ma, Zhihan Zhang, Michael Saugstad, Kevin Wu, Yochai Eisenberg, Valerie Novack, Brent Chamberlain, and Jon E. Froehlich. 2023. BusStopCV: A Real-time AI Assistant for Labeling Bus Stop Accessibility Features in Streetscape Imagery. In *Proceedings of the 25th International ACM SIGACCESS Conference on Computers and Accessibility (ASSETS '23)*. Association for Computing Machinery, New York, NY, USA, 1–6. doi:10.1145/3597638.3614481
- [43] Myron M. LaBan and Matthew Grimm. 2010. Handicapped Parking: A Privilege or a Right? *American Journal of Physical Medicine & Rehabilitation* 89, 4 (April 2010), 345. doi:10.1097/PHM.0b013e3181c56a53
- [44] Lambros Lambrinos and Aristotelis Dosis. 2013. DisAssist: An internet of things and mobile communications platform for disabled parking space management. In *2013 IEEE Global Communications Conference (GLOBECOM)*. 2810–2815. doi:10.1109/GLOCOM.2013.6831500 ISSN: 1930-529X.
- [45] C. Lanczos. 1950. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *J. Res. Nat. Bur. Standards* 45, 4 (Oct. 1950), 255. doi:10.6028/jres.045.026
- [46] LARIAC4. 2015. LACounty_Cache/LACounty_Aerial_2014 (MapServer). https://cache.gis.lacounty.gov/cache/rest/services/LACounty_Cache/LACounty_Aerial_2014/MapServer
- [47] Andrew Larkin, Xiang Gu, Lizhong Chen, and Perry Hystad. 2021. Predicting perceptions of the built environment using GIS, satellite and street view image approaches. *Landscape and Urban Planning* 216 (Dec. 2021), 104257. doi:10.1016/j.landurbplan.2021.104257
- [48] Jamie Lau and Austin Doeze. 2020. Services Directory - Spring_Hill_Imagery_WGS84(MapServer). https://tiles.arcgis.com/tiles/tf0XsRR9ptiKNVW2/arcgis/rest/services/Spring_Hill_Imagery_WGS84/MapServer
- [49] Chang Dae Lee, Alicia M Koontz, Rosemarie Cooper, Shivashankar Sivakanthan, William Chernicoff, Allison Brunswick, Nikitha Deepak, Hailee R. Kulich, Justin LaFerrier, Celia R Lopes, Nicole L. Collins, Brad E. Dicianno, and Rory A Cooper. 2024. Understanding Travel Considerations and Barriers for People with Disabilities to Using Current Modes of Transportation Through Journey Mapping. *Transportation Research Record* 2678, 5 (May 2024), 271–287. doi:10.1177/03611981231188730 Publisher: SAGE Publications Inc.
- [50] Jiwon Lee, Dongjin Lee, Jongwook Jang, and Sungjin Jang. 2021. Development of Disabled Parking System Using Deep Learning Model. *Proceedings of the Korean Institute of Information and Communication Sciences Conference* (2021), 175–177. <https://koreascience.kr/article/CFKO202121751202237.page> Publisher: The Korea Institute of Information and Communication Engineering.
- [51] Chu Li, Rock Yuren Pang, Delphine Labbé, Yochai Eisenberg, Maryam Hosseini, and Jon E. Froehlich. 2025. Accessibility for Whom? Perceptions of Sidewalk Barriers Across Disability Groups and Implications for Designing Personalized Maps. Yokohama, Japan. doi:10.1145/3706598.3713421 arXiv:2502.19888 [cs].
- [52] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. Microsoft COCO: Common Objects in Context. In *Computer Vision – ECCV 2014*, David Fleet, Tomas Pajdla, Bernt Schiele, and Tinne Tuytelaars (Eds.). Springer International Publishing, Cham, 740–755. doi:10.1007/978-3-319-10602-1_48
- [53] Yvonna S. Lincoln and Egon G. Guba. 1985. *Naturalistic inquiry*. SAGE Publications, Thousand Oaks, CA.
- [54] Max Lindblom and Carl Nilsson. 2019. Evaluation of Parking Space Detection From Aerial Imagery Using Convolutional Neural Networks. (2019). <http://lup.lub.lu.se/student-papers/record/8971684>
- [55] Alex Liu, Kevin Wu, Minchu Kulkarni, Mikey Saugstad, Peyton Rapo, Jeremy Freiburger, Maryam Hosseini, Chu Li, and Jon E. Froehlich. 2024. Towards fine-grained sidewalk accessibility assessment with deep learning: Initial benchmarks and an open dataset. In *Extended abstract proceedings of the 26th international ACM SIGACCESS conference on computers and accessibility*. St. John's, Newfoundland and Labrador. doi:10.1145/3663548.3688531 Number of pages: 11 tex.booktitleshort: Extended Abstract Proceedings of ASSETS 2024.
- [56] Luyu Liu, Armita Kar, Ahmad Ilderim Tokay, Huyen T. K. Le, and Harvey J. Miller. 2023. Disparities in public transit accessibility and usage by people with mobility disabilities: An evaluation using high-resolution transit data. *Journal of Transport Geography* 109 (May 2023), 103589. doi:10.1016/j.jtrangeo.2023.103589
- [57] Zhan Liu, Nicole Glassey, Maria Sokhn, and Eline de Gaspary. 2017. Crowdsourcing-Based Mobile Application for Wheelchair Accessibility. *Journal on Technology & Persons with Disabilities* 5 (April 2017), 1–15.
- [58] Weite Lu, Upali Vandebona, and Masaru Kiyota. 2014. Analysis of experience with formalizing handicapped parking system. *Transportation Research Part F: Traffic Psychology and Behaviour* 26 (Sept. 2014), 62–71. doi:10.1016/j.trf.2014.06.011
- [59] Weite Lu, Upali Vandebona, Masaru Kiyota, and Haiqiang Liu. 2019. Analysis of Signage Strategies to Reduce Misuse of Disabled Parking at a University Campus. *Journal of Applied Behavior Analysis* (2019).
- [60] Weite Lu, Chunqin Zhang, and Xunyou Ni. 2019. Research on the rationale of width of disabled parking spaces and correlation of selection by wheelchair users and others. *Journal of Transport & Health* 12 (March 2019), 130–141. doi:10.1016/j.jth.2019.01.006

- [61] Martin Margreiter, Philipp Mayer, and Foteini Orfanou. 2015. A Concept for Crowdsourcing of In-Vehicle Data to Improve Urban On-Street Parking.
- [62] Christian Menkens, Julian Sussmann, Miriam Al-Ali, Emanuel Breitsameter, Jelena Frtunik, Tobias Nendel, and Tobias Schneiderbauer. 2011. EasyWheel - A Mobile Social Navigation and Support System for Wheelchair Users. In *2011 Eighth International Conference on Information Technology: New Generations*. 859–866. doi:10.1109/ITNG.2011.149
- [63] Amin Mobasheri, Jonas Deister, and Holger Dieterich. 2017. Wheelmap: the wheelchair accessibility crowdsourcing platform. *Open Geospatial Data, Software and Standards* 2, 1 (Nov. 2017), 27. doi:10.1186/s40965-017-0040-5
- [64] Janice M. Morse. 2015. Critical Analysis of Strategies for Determining Rigor in Qualitative Inquiry. *Qualitative Health Research* 25, 9 (Sept. 2015), 1212–1222. doi:10.1177/1049732315588501
- [65] C. R. Mwaka, K. L. Best, C. Cunningham, M. Gagnon, and F. Routhier. 2024. Barriers and facilitators of public transport use among people with disabilities: a scoping review. *Frontiers in Rehabilitation Sciences* 4 (Jan. 2024). doi:10.3389/fresc.2023.1336514 Publisher: Frontiers.
- [66] Bipul Neupane, Teerayut Horanont, and Jagannath Aryal. 2021. Deep Learning-Based Semantic Segmentation of Urban Features in Satellite Images: A Review and Meta-Analysis. *Remote Sensing* 13, 4 (Jan. 2021), 808. doi:10.3390/rs13040808 Number: 4 Publisher: Multidisciplinary Digital Publishing Institute.
- [67] Huan Ning, Xinyue Ye, Zhihui Chen, Tao Liu, and Tianzhi Cao. 2022. Sidewalk extraction using aerial and street view images. *Environment and Planning B* 49, 1 (Jan. 2022), 7–22. doi:10.1177/2399808321995817 Publisher: SAGE Publications Ltd STM.
- [68] Julien Nyambal and Richard Klein. 2017. Automated parking space detection using convolutional neural networks. In *2017 Pattern Recognition Association of South Africa and Robotics and Mechatronics (PRASA-RobMech)*. 1–6. doi:10.1109/RoboMech.2017.8261114
- [69] MassGIS (Bureau of Geographic Information). 2019. MassGIS Data: 2019 Aerial Imagery | Mass.gov. <https://www.mass.gov/info-details/massgis-data-2019-aerial-imagery>
- [70] U.S. Department of Justice. 2010. 2010 ADA Standards for Accessible Design. <https://www.ada.gov/law-and-regs/design-standards/2010-stds/>
- [71] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. 2016. You Only Look Once: Unified, Real-Time Object Detection. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Las Vegas, NV, USA, 779–788. doi:10.1109/CVPR.2016.91
- [72] Andrew Regester and Vamsi Paruchuri. 2020. Using Computer Vision Techniques for Parking Space Detection in Aerial Imagery. In *Advances in Computer Vision*, Kohel Arai and Supriya Kapoor (Eds.). Springer International Publishing, Cham, 190–204. doi:10.1007/978-3-030-17798-0_17
- [73] Timothy Ross and Ronald Buliung. 2019. Access work: Experiences of parking at school for families living with childhood disability. *Transportation Research Part A: Policy and Practice* 130 (Dec. 2019), 289–299. doi:10.1016/j.tra.2019.08.016
- [74] Manaswi Saha, Michael Saugstad, Hanuma Teja Maddali, Aileen Zeng, Ryan Holland, Steven Bower, Aditya Dash, Sage Chen, Anthony Li, Kotaro Hara, and Jon Froehlich. 2019. Project Sidewalk: A Web-based Crowdsourcing Tool for Collecting Sidewalk Accessibility Data At Scale. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY, USA, 1–14. doi:10.1145/3290605.3300292
- [75] Sayani Sarkar, Michael W. Totaro, and Khalid Elgazzar. 2019. Intelligent drone-based surveillance: application to parking lot monitoring and detection. In *Unmanned Systems Technology XXI*, Vol. 11021. SPIE, 13–19. doi:10.1117/12.2518320
- [76] Tom [D-]IA Sen. Harkin. 1990. S.933 - 101st Congress (1989-1990): Americans with Disabilities Act of 1990. <https://www.congress.gov/bill/101st-congress/senate-bill/933> Archive Location: 1989-05-09.
- [77] Young-Woo Seo and Chris Urmson. 2011. A Hierarchical Image Analysis for Extracting Parking Lot Structures from Aerial Images. (April 2011).
- [78] Joseph P. Shapiro. 1994. *No pity: People with disabilities forging a new civil rights movement*. Times Books, New York.
- [79] Donald Shoup. 2018. Ending the Abuse of Disabled Parking Placards. In *Parking and the City*. Routledge. Num Pages: 3.
- [80] Marina Sokolova and Guy Lapalme. 2009. A systematic analysis of performance measures for classification tasks. *Information Processing & Management* 45, 4 (July 2009), 427–437. doi:10.1016/j.ipm.2009.03.002
- [81] Donald H. Stone. 2007. You Take My Space, I Take Your Air: An Empirical Study of Disabled Parking and Motor Vehicle Laws for Persons with Disabilities. *Ohio Northern University Law Review* 33, 2 (2007), 665–714. <https://heinonline.org/HOL/P?h=hein.journals/onu33&i=445>
- [82] Renato Stopic, Eduardo Dias, Maurice de Kleijn, and Eric Koomen. 2023. Satellite parking: a new method for measuring parking occupancy. *AGILE: GIScience Series* 4 (June 2023), 1–5. doi:10.5194/agile-giss-4-44-2023 Publisher: Copernicus GmbH.
- [83] U.S. Geological Survey. 2018. USGS EROS Archive - Aerial Photography - High Resolution Orthoimagery (HRO). <https://www.usgs.gov/centers/eros/science/usgs-eros-archive-aerial-photography-high-resolution-orthoimagery-hro>
- [84] EagleView Technologies. 2021. BaseMaps/KingCo_Aerial_2021 (MapServer). https://gismaps.kingcounty.gov/arcgis/rest/services/BaseMaps/KingCo_Aerial_2021/MapServer
- [85] Marc Tschentscher and Marcel Neuhausen. 2012. Video-based parking space detection. (2012). https://www.ini.rub.de/upload/file/1470692848_232b7133d00f0473cd4/tschentscherneuhausen_parking_space_fbi2012.pdf
- [86] Compton J. Tucker, Denelle M. Grant, and Jon D. Dykstra. 2004. NASA's Global Orthorectified Landsat Data Set. *Photogrammetric Engineering & Remote Sensing* 70, 3 (March 2004), 313–322. doi:10.14358/PERS.70.3.313
- [87] U.S. Department of Justice Civil Rights Division. 2025. Accessible Parking Spaces. <https://www.ada.gov/topics/parking/>
- [88] Arun Varghese and G. Sreelekha. 2020. An Efficient Algorithm for Detection of Vacant Spaces in Delimited and Non-Delimited Parking Lots. *IEEE Transactions on Intelligent Transportation Systems* 21, 10 (Oct. 2020), 4052–4062. doi:10.1109/TITS.2019.2934574 Conference Name: IEEE Transactions on Intelligent Transportation Systems.
- [89] OpenStreetMap Wiki. 2022. Slippy map - OpenStreetMap Wiki. https://wiki.openstreetmap.org/wiki/Slippy_map
- [90] Keiichi Yamada and Morimichi Mizuno. 2001. A vehicle parking detection method using image segmentation. *Electronics and Communications in Japan (Part III: Fundamental Electronic Science)* 84, 10 (2001), 25–34. doi:10.1002/ecjc.1039. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/ecjc.1039>.
- [91] Yifang Yin, Wenniao Hu, An Tran, Hannes Kruppa, Roger Zimmermann, and See-Kiong Ng. 2022. A Context-enriched Satellite Imagery Dataset and an Approach for Parking Lot Detection. In *2022 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. IEEE, Waikoloa, HI, USA, 1391–1400. doi:10.1109/WACV51458.2022.00146
- [92] Guoqing Zhou, Weirong Chen, J.A. Kelmelis, and Deyan Zhang. 2005. A comprehensive study on urban true orthorectification. *IEEE Transactions on Geoscience and Remote Sensing* 43, 9 (Sept. 2005), 2138–2147. doi:10.1109/TGRS.2005.848417 Conference Name: IEEE Transactions on Geoscience and Remote Sensing.
- [93] Zhuofan Zong, Guanglu Song, and Yu Liu. 2023. DETRs with Collaborative Hybrid Assignments Training. doi:10.48550/arXiv.2211.12860 arXiv:2211.12860 [cs].

A Appendix

A.1 Detecting Objects Split Across Tiles

Algorithm 1 Detect objects in a 512×512 Tile (refer to Figure 12)

Time Cost: $O(4n)$.

Assume: Objects are no wider than 100 pixels.

Input: Top left tile position $\{a\}$.

Step 1: Detect in $\{a, b, d, e\}$. Ignore objects with a centroid <50 px from border.

Step 2: Detect in $\{b, c, e, f\}$. Ignore objects with a centroid <50 px from border and >50 px from the vertical middle.

Step 3: Detect in $\{d, e, g, h\}$. Ignore objects with a centroid <50 px from border and >50 px from the horizontal middle.

Step 4: Detect in $\{e, f, h, i\}$. Only keep objects with a centroid within a 100 px square around the center point.

Return: All non-ignored objects and their positions.

Because of tiling, some objects will be split across tiles and potentially detected multiple times. We consider two possible solutions to this problem: (1) to combine detected objects across borders into one based on their geometry (the approach used by Tile2Net [32]), or (2) to create a new image, encapsulating the area at the border, and to detect within this new image, thereby detecting one whole object rather than multiple fragments of an object.

We opted for Method 2 as Method 1 would raise other detection problems to consider. For example, parking spaces are quite small relative to the size of the image; what if, by being split across tiles, a space is only detected in one tile or in no tiles at all? Furthermore, a parking space's classification significantly depends on its features




Figure 12: Illustration for Algorithm 1. We employ a sliding-window algorithm to detect objects split between tiles. In this example, the input to the pipeline is the top left tile position $\{a\}$, from which it will stitch together and detect objects in the 2×2 square $\{a, b, d, e\}$, denoted in white. The parking spaces in blue are half in and half out of the square. As the algorithm proceeds, the parking spaces will be detected in their entirety in the squares $\{d, e, g, h\}$ and $\{e, f, h, i\}$ and be included in the results for the square $\{a, b, d, e\}$.

(e.g., the number of aisles it has, if it has a painted logo, etc.). What if the space is split so that the logo or an aisle is visible in one tile but not the other? How would the class be determined? Finally, parking spaces tend to be both small and adjacent to one another. The object combination would need to be very robust to avoid combining two neighboring objects into one.

We avoid all these issues by using Method 2, ensuring we detect only *whole* objects, but at the cost of computation. To do this, we assume that a parking space (with aisles) is no wider than 100 pixels in a given dimension, which, at our zoom level, holds for all but the most extreme of scenarios (such as semi-truck parking). In this way, the locator model is only tasked with detecting entire parking spaces and passes only unpadded, 100×100 crops to the OBB model. For a detailed breakdown of the approach and algorithm, including computation cost, see Figure 12 and Algorithm 1.

A.2 Georeferencing

For a pixel in a given tile whose x and y positions are in the orthoimage tile system, we can obtain that pixel's location on Earth (given zoom level 20). Orthoimage tiles are representations of the common Web Mercator projection [4]. We discern a pixel's latitude and longitude (in the commonly used WGS84⁶) with the following:

$$longitude = \frac{x}{2^{20}} \cdot 360 - 180$$

⁶<https://epsg.io/4326>

$$latitude = \frac{180}{\pi} \cdot \arctan(\sinh(\pi - 2\pi \frac{y}{2^{20}}))$$

Here, x and y are the fractional values of the pixel in relation to the tile in which it belongs. For example, in a 256×256 orthoimage tile with $x = 168046, y = 366004$, if a pixel is at the 128th row and column, the values used in the preceding equations are $x = 168046.5, y = 366004.5$. Then, to calculate distances, and therefore widths, the latitude and longitude are converted into points in a projected coordinate system, EPSG:3857⁷ (i.e., pseudo-mercator with units in meters), and simple Euclidean distance is used.

⁷<https://epsg.io/3857>

A.3 Width Characterization Histograms




Figure 13: Results of the pipeline’s derivation of width on the test set, for the classes with over 30 samples: *dp_one_aisle*, *dp_two_aisle*, and *one_aisle*. Other classes are not shown, as histogram patterns are less visible and robust with the less common classes. Differences in pixels above are shown in cyan, and percentage difference below in green. Differences are predicted minus ground truth; therefore, a positive value indicates overestimation. The model sees a large peak around zero (*i.e.*, accurate estimation), with wider variation depending on class.

A.4 Locator Models and Results

	YOLOv11 large	DINO (ResNet50 backbone)	Co-DETR (Swin-L backbone)
Training Epochs	200	50	50
Selected Epoch	126	27	23
Batch Size	16	2	2
Learning Rate	10^{-2}	10^{-4}	2×10^{-4}
Weight Decay	5×10^{-4}	10^{-4}	10^{-4}
Augmentations	Hue		
	Saturation		
	Brightness		
	Translation	Resizing	Resizing
	Resizing	Random cropping	Random flipping
	Flipping (left right only)		Random cropping
	Mosaic		
	Erasing		
Precision/Recall/F1	0.76 / 0.92 / 0.83	0.87 / 0.86 / 0.86	0.85 / 0.94 / 0.89

Table 7: Training parameters and results for the three models tested for the object detection component of the pipeline. Co-DETR sees the best performance in Recall and F1 score.