Truck Parking Information and Management System (TPIMS)

WA-RD 956.1

Yinhai Wang
Kehua Chen
Nutvara Jantarathaneewat

December 2025

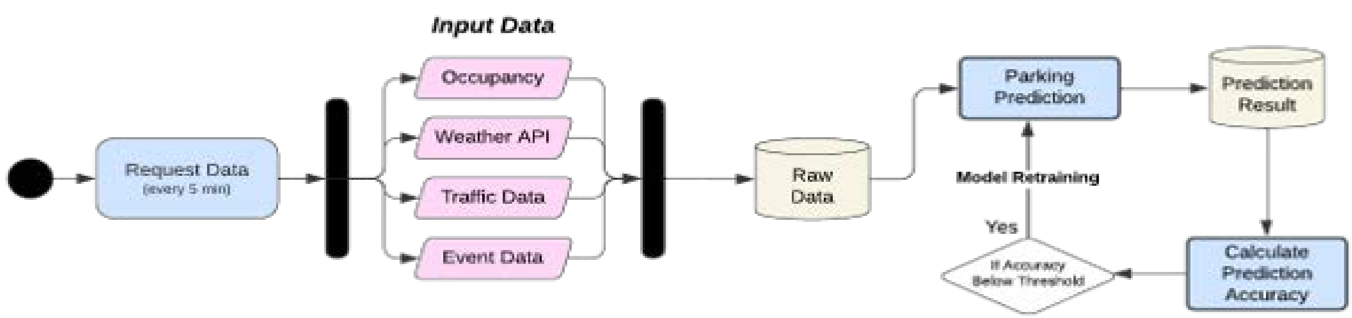

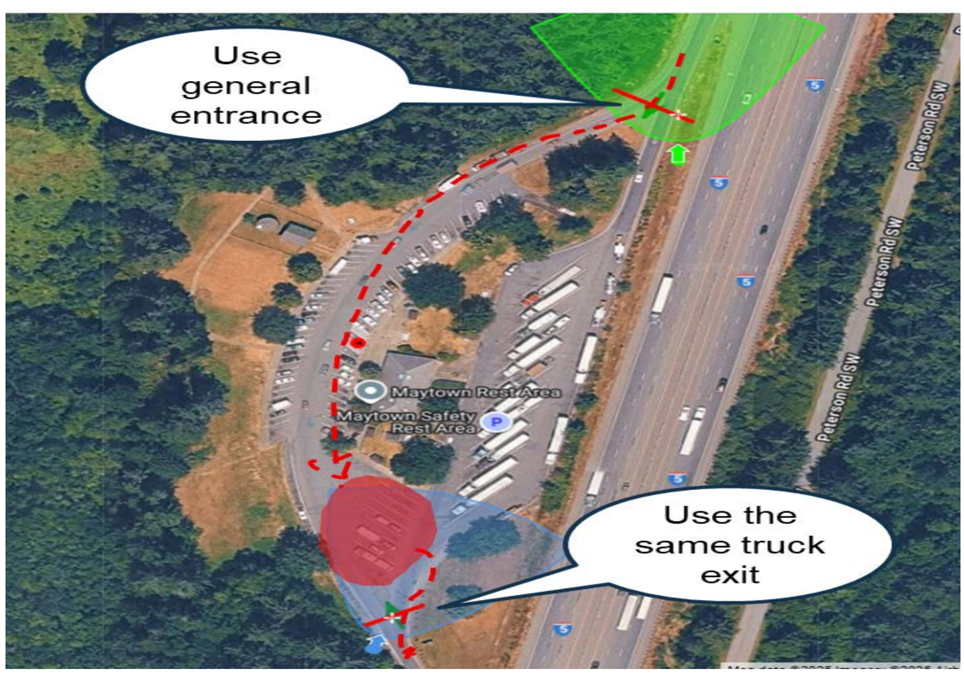

Washington State Department of Transportation
Office of Research & Library Services

WSDOT Research Report

Research Report
Agreement T1461, Task 90
WSDOT FMCSA-Truck Parking Information and Management System (TPIMS)
WA-RD 956.1

Truck Parking Information and Management System (TPIMS)

by

Yinhai Wang
University of Washington

Kehua Chen
University of Washington

Nutvara Jantarathaneewat
University of Washington

Department of Civil and Environmental Engineering
University of Washington, Box 352700
Seattle, Washington 98195

Sponsored by

Washington State Department of Transportation (WSDOT)
310 Maple Park Avenue SE
PO Box 47300
Olympia, WA 98504-7300

Washington State Department of Transportation
Technical Monitor
Matthew Neeley

Prepared for

The State of Washington
Department of Transportation
Julie Meredith, Secretary

December 2025

Technical Report Documentation Page

1. Report No. WA-RD 956.1	**2. Government Accession No.**	**3. Recipient's Catalog No.**
4. Title and Subtitle Truck Parking Information and Management System (TPIMS)		**5. Report Date** December 2025
		6. Performing Organization Code
7. Author(s) Yinhai Wang, Kehua Chen, Nutvara Jantarathaneewat		**8. Performing Organization Report No.**
9. Performing Organization Name and Address Washington State Transportation Center University of Washington, Box 359446 University Tower, 4333 Brooklyn Ave NE Seattle, WA 98195-9446		**10. Work Unit No. (TRAIS)**
		11. Contract or Grant No. Agreement T1461-90
12. Sponsoring Organization Name and Address Washington State Department of Transportation Transportation Building, MS 47372 Olympia, Washington 98504-7372 14 David Strich, Project Coordinator, David.Strich@wsdot.wa.gov		**13. Type of Report and Period Covered** Final research report
		14. Sponsoring Agency Code
15. Supplementary Notes		

16. Abstract

Washington state faces persistent challenges in providing reliable and accessible truck parking information to commercial drivers, resulting in inefficient trip planning, increased fatigue-related safety risks, and operational inefficiencies across major freight corridors. To address these issues, the Washington State Department of Transportation (WSDOT), in partnership with the University of Washington STAR Lab, developed a self-learning and optimizing Truck Parking Information and Management System (TPIMS). The four-year project integrated statewide stakeholder engagement, comprehensive detection-system review, site evaluations across I-5 and I-90 corridors, and the development of a multi-source data management platform that consolidates real-time occupancy, parking events, weather, and traffic data. Multiple sensing technologies—including space-by-space and entry/exit systems—were assessed through field deployment, calibration, and validation to determine their accuracy, cost-effectiveness, and long-term maintainability.

Building on these data streams, the team advanced a machine-learning–based prediction algorithm capable of forecasting near-term parking availability using spatial–temporal patterns and heterogeneous feature embeddings. A robust TPIMS database and server architecture were implemented to support real-time data flow, historical archiving, and multi-source fusion. A standards-aligned application programming interface (API) was developed to provide open, programmatic access to occupancy and prediction data for third-party developers, private truck stops, and other agencies, enabling broad dissemination of truck parking information. User testing and interface enhancements further refined system usability and operational readiness.

The project demonstrated the feasibility and scalability of a statewide TPIMS, delivering a validated detection framework, an operational prediction model, and an interoperable API that together enhance situational awareness and improve safety for commercial drivers. These outcomes provide a strong foundation for future expansion of TPIMS to additional sites and neighboring states, supporting national efforts to improve freight mobility and reduce risks associated with inadequate truck parking.

17. Key Words Truck parking, parking information management system, parking prediction, Washington state		**18. Distribution Statement** No restrictions.	
s19. Security Classification (of this report) Unclassified.	**20. Security Classification (of this page)** Unclassified.	**21. No. of Pages** 109	**22. Price** NA

Form DOT F 1700.7 (8-72) Reproduction of completed page authorized

Disclaimer

The contents of this report reflect the views of the authors, who are responsible for the facts and accuracy of the data presented herein. This document is disseminated through the Washington State Department of Transportation. The contents do not necessarily reflect the views or policies of Washington State Department of Transportation or the Federal Highway Administration. This report does not constitute a standard, specification, or regulation.

TABLE OF CONTENTS

TABLE OF FIGURES

TABLES

1. INTRODUCTION AND BACKGROUND

Truck parking space availability for commercial drivers traveling through Washington state is limited. To help commercial drivers plan their trips and maximize the use of available parking, Washington State Department of Transportation (WSDOT), in partnership with the University of Washington (UW) STAR Lab, will develop and install a self-learning and optimizing Truck Parking Information and Management System (TPIMS) at state-owned truck parking facilities at weigh stations and safety rest areas (SRAs).

The project will be approached as a collaborative effort between WSDOT and UW. WSDOT would like to install truck parking technology at all the weigh stations and rest areas that currently have truck parking and are in the T1 Corridors. A total of 273 truck parking stalls are located along Interstate 5 (I-5) throughout 16 rest areas and weigh stations (not including Fort Lewis and Scatter Creek). Detection sensors will be installed in each parking stall or another monitoring technology will be installed so each parking location can be monitored for availability. The TPIMS will collect data from technology installed at weigh stations and rest areas to monitor the available commercial truck parking locations. The data will then be placed in an algorithm that will predict the probability of a stall being available to a commercial driver at one of the truck parking locations at a specific time and location. The algorithm will continuously analyze the data and update its probability model.

The outcomes of this research have clear implications for improving the efficiency of scheduling truck parking and enhancing the robustness of the whole truck operation system. In addition, this project will improve truck parking management at weigh stations and SRAs to mitigate the safety risk caused by fatigued driving and other factors. In addition, the comprehensive review and evaluation of the existing truck parking detection sensors will give WSDOT a better understanding of the pros, cons, and retrofitting feasibility of the existing truck parking detection products. This will provide guidance on the expansion of TPIMS implementation in more places.

2. RESEARCH OBJECTIVES

1) **STAKEHOLDER ENGAGEMENT:** To engage with vendors for vehicle occupancy detection technologies, internal and external partners for collaboration on the information management and dissemination system, and public stakeholders about the features of the TPIMS
2) **DETECTION SYSTEM SELECTION AND SITE EVALUATION:** To conduct a review of available detection systems and make site visits and evaluations.
3) **INSTALLATION:** To develop an installation plan and install the detection systems at the selected sites, set up a data management system, and perform validation and testing of the systems.
4) **API DEVELOPMENT:** To develop an application programming interface (API) for third-party partners to build applications.
5) **USER TESTING:** The WSDOT team will conduct pilot testing of the systems' application to connected vehicle (CV) operators and the public through surveys, focus groups, etc.

3. LITERATURE REVIEW

3.1 TPIMS DEPLOYMENTS IN OTHER STATES

3.1.1 The WisTransPortal System

Several Midwestern and Great Plains states have already deployed large-scale Truck Parking Information Management Systems (TPIMS) through the Mid-American Association of State Transportation Officials (MAASTO) coalition, providing a mature reference model for regional data sharing, public dissemination, and third-party API access (Figure 3-1). Since 2019, these states—such as Illinois, Indiana, Iowa, Kansas, Kentucky, Michigan, Minnesota, Ohio, and Wisconsin—have operated integrated TPIMS platforms that aggregate real-time and archived parking availability data from more than 150 public and private sites. Their systems include web-based archived data portals, query interfaces for historical analysis, and standardized public JavaScript Object Notation (JSON) feeds that update every 5 minutes for use by app developers, freight operators, and navigation service providers. The coordinated architecture demonstrates the feasibility of multi-state interoperability, unified data standards, and scalable dissemination mechanisms that support both operational needs and research use cases. These deployments provide valuable lessons for Washington state as it modernizes its own multi-source TPIMS, particularly in areas of API standardization, data archiving practices, and public-facing information services.

Figure 3-1 TPIMS Archive Portal on the WisTransPortal System, providing access to performance reports, archived truck parking data, and public data feeds from eight participating MAASTO states.

3.1.2 California Statewide Truck Parking Study (CalTrans, 2022)

CalTrans' 2022 Truck Parking Availability System (TPAS) study outlines a coordinated, multistate approach along the I-10 corridor involving Arizona, New Mexico, and Texas. The system aims to provide real-time information that supports better parking decisions, reduces unauthorized parking, and minimizes time spent searching for spaces. The study proposes a phased implementation: short-term efforts focus on developing a Concept of Operations (ConOps) for site prioritization, data collection, and information dissemination; mid-term actions emphasize deployment along top-priority freight corridors; and long-term plans target a full statewide TPAS deployment.

3.1.3 Oregon Truck Parking Information Management System (TPISM) (ODOT, 2020 and 20204)

Oregon DOT's 2024 Truck Parking Information Management System (TPIMS) is part of a multistate initiative along the I-5 corridor in partnership with Washington and California. The project deploys sensors and cameras at truck parking locations to collect occupancy data and transmit real-time availability information directly to truck drivers. ODOT plans to disseminate this information through multiple channels, including digital roadside signs, in-cab systems, and smartphone or travel-information apps. Supported by federal Infrastructure for Rebuilding America (INFRA) grant funding, the system aims to improve safety, reduce search time, and create a connected parking-information network across the West Coast.

3.2 TRUCK PARKING DETECTION TECHNOLOGIES

Truck-parking detection systems rely on two primary counting methods—individual space (slot-by-slot) detection and in/out counting—each suited for different facility sizes and operational needs. Slot-by-slot systems directly monitor each parking space using technologies such as magnetic sensors, radar, or video imaging. These provide high accuracy and avoid cumulative errors but can become costly for large truck lots because of the need for many sensors. By contrast, in/out counting systems install sensors at entrances and exits to track vehicles entering and leaving. These solutions are highly scalable and less dependent on pavement markings, although they are prone to error accumulation over time and require periodic ground-truth corrections.

A wide range of sensing technologies has been used in pilot deployments. Magnetic sensors offer low-power, wireless, and durable designs but may require multiple sensors per space. Laser and radar sensors provide precise detection and classification, especially in challenging lighting or weather conditions, although they can be sensitive to physical obstructions and environmental effects. Video image processing, including AI-enhanced and 3D reconstruction methods, enables coverage of large areas with fewer devices and can resolve occlusion issues, but it requires higher bandwidth and processing capacity. GPS-based methods provide insights into regional parking demand but lack the penetration and granularity needed for real-time occupancy estimation.

Communication systems vary from wired to cellular and LoRa-based networks, depending on vendor design. Information dissemination channels include roadside variable message signs (VMS), websites, mobile apps, and in-cab devices, although user surveys indicate limited awareness among drivers, highlighting persistent challenges in effective information delivery.

Pilot deployments in the U.S. and Europe have demonstrated that most modern technologies can achieve over 95 percent detection accuracy when properly calibrated. However, each technology remains sensitive to site-specific factors such as lot geometry, driver behavior, environmental conditions, and installation quality. Studies have consistently shown that regular supervision, calibration, and design adjustments (such as clear markings and structured layouts) are crucial to performance. Overall, although multiple technologies have proved effective, large-scale, unattended deployment remains challenging, and the industry continues to face knowledge gaps regarding long-term reliability, scalability, and cross-site performance.

3.3 PREDICTION ALGORITHMS FOR PARKING AVAILABILITY

3.3.4 Truck Parking Occupancy Prediction: Xgboost-LSTM Model Fusion (Gutmann et al., 2021)

This study proposed a hybrid machine-learning approach for forecasting truck parking occupancy, addressing the growing challenge that truck drivers face in locating safe and legal parking—particularly when planning rest stops in advance. Unlike many existing systems that only provide real-time availability, the proposed model predicts occupancy 30, 60, 90, and 120 minutes ahead, enabling drivers to make informed decisions about downstream parking options. The method relies exclusively on historical occupancy data, making them widely applicable to any detection technology capable of recording past usage patterns. The model architecture fuses Extreme Gradient Boosting (XGBoost)—a strong performer in traffic forecasting—with long short-term memory (LSTM) networks, which are well suited for capturing temporal dependencies in time series. A feed-forward neural network integrates the outputs of both models to improve overall prediction accuracy. Trained on one year of occupancy data from a single rest area, the hybrid model achieved low forecasting errors, with root mean square error (RMSE) values of 2.1, 2.9, 3.5, and 4.1 trucks for the four prediction horizons, respectively. This demonstrated that a lightweight, single-source, meta-model framework can reliably predict near-term truck parking demand and can be easily integrated into existing monitoring systems.

3.3.5 Prediction Truck Parking Occupancy Using Machine Learning (Slavova et al., 2022)

Several studies have explored features that influence parking occupancy prediction, often incorporating factors such as weather, traffic flow, holidays, and special events. However, evidence from truck-parking–specific research has suggested that these variables may play a snaller role for freight operations. For example, Slavova et al. found that unlike car parking—for which rainfall, temperature, and holiday effects are strong predictors—truck parking behavior is driven primarily by temporal patterns and historical occupancy, with minimal influence from weather or special events. Using 1.5 years of real-world data from a truck parking facility in the Netherlands, the authors demonstrated that decision tree models trained solely on time-of-day, weekday, and past occupancy achieved superior performance for short-term truck parking prediction, delivering fast training, real-time inference, and high accuracy (a one-hour RMSE of 0.0029). Their findings highlighted two important takeaways:

- Truck parking demand is largely independent of weather conditions and holidays, and
- Simple, single-source models—such as decision trees using only temporal and occupancy patterns—can outperform more complex multi-feature approaches in real-time operational settings.

3.3.6 Truck Parking Pattern Aggregation and Availability Prediction by Deep Learning (Yang et al., 2022)

In the pilot project, an advanced deep-learning model known as the Truck Parking Occupancy Prediction (TPOP) neural network was developed to support short-term parking availability forecasting. This work built on three major components originally introduced by Yang et al. (2022). First, the research team implemented a TPIMS at two pilot sites—Scatter Creek Safety Rest Area and the Nisqually Weigh Station. This system enabled space-by-space parking status data collection, real-time data processing, and generation of availability information that served as the foundation for prediction model development. Second, the team conducted a comprehensive investigation of long-term slot-level truck parking patterns across the United States. To analyze recurring temporal behaviors, they introduced a sequence-based pattern similarity approach, the Advanced Sequence Alignment Method (ASAM), which aggregates large-scale parking sequences and uncovers underlying behavioral patterns that inform occupancy changes over time. Third, leveraging these insights, the project developed the TPOP deep-learning prediction model, an attributes-aware sequence-to-sequence architecture that integrates temporal patterns and site-specific attributes to forecast short-term occupancy. The model demonstrated strong accuracy, achieving mean absolute percentage error (MAPE) values of 4.19 percent, 4.84 percent, 5.07 percent, and 5.82 percent at 2-, 4-, 8-, and 16-minute prediction horizons, respectively (Yang et al., 2022). These results indicated that TPOP can reliably capture minute-level fluctuations in parking demand and outperform traditional statistical prediction methods. Figure 3-2 provides a comparison of prediction performance across the time horizons.

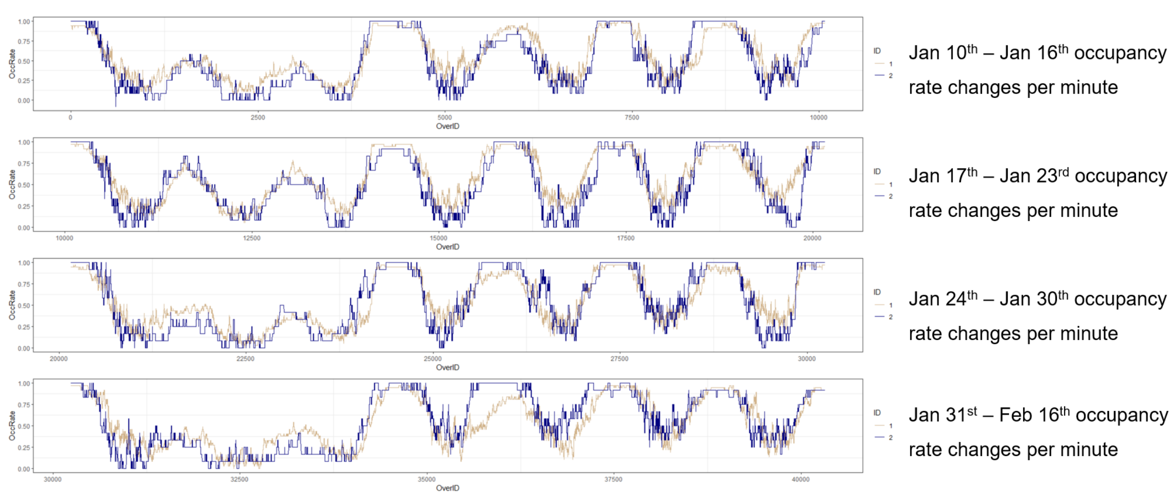

ID 1 and 2 represent two different truck parking lots here.

Figure 3-2 Comparison of TPOP prediction result at 2 min, 4 min, 8 min, and 16 min ahead (Yang et al. 2022)

Building on the original prediction framework, several potential improvements and modifications were identified for future development. First, the TPIMS, originally implemented at two pilot sites, could be expanded to multiple sites and enhanced by integrating data from both slot-by-slot sensors and in/out detectors, allowing the system to operate effectively under diverse sensing configurations. Second, the analysis of long-term slot-level parking patterns could be strengthened by incorporating additional contextual factors, such as surrounding land use, distance to adjacent parking facilities, driver preferences, and impacts from major events such as COVID-19. These enhancements would also support modeling of network-level parking patterns, reflecting how parking demand shifts across nearby locations. Third, the TPOP deep-learning prediction model could be extended to perform network-wide availability prediction, and future work should investigate whether different detection technologies—such as magnetic sensors, radar, or cameras—require customized prediction models to account for differences in data characteristics and accuracy. Together, these improvements would create a more robust, scalable, and context-aware prediction system for statewide truck parking management.

3.3.7 Spatio-Temporal Parking Occupancy Forecasting Integrating Parking Sensing Records and Street-Level Images (Gong et al., 2023)

This study presented a spatio-temporal parking occupancy forecasting framework that integrates large-scale parking sensing records with street-level imagery to improve prediction accuracy. Unlike previous research that treated parking lots independently, this work explicitly modeled correlations between parking facilities by incorporating three types of similarity measures: (1) spatial correlations derived from travel distance and Bayesian probability–based relationships between neighboring car parks; (2) visual similarities obtained through a Convolutional Neural Network (CNN) applied to street-view images to capture the built environment and structural characteristics of each lot; and (3) activity-type similarities computed using cosine similarity. These components were combined with Graph Convolutional Networks (GCN) and Gated Recurrent Units (GRU) to learn both the spatial dependencies and temporal dynamics of parking demand. Using over ten million sensing records from parking lots in Ningbo and Beijing, the proposed Temporal-GCN–based Correlated Parking Prediction Model (CPPM) significantly outperformed baseline methods and demonstrated strong transferability across cities with different sizes and development levels. Beyond prediction accuracy, the study also revealed meaningful behavioral insights, including distinct parking preferences among commuters in the two cities, providing valuable information for urban parking management and planning.

3.4 WASHINGTON STATE TRUCK PARKING PILOT PROJECT REVIEW

The Washington State Truck Parking Pilot Project evaluated the feasibility of deploying real-time truck parking monitoring using slot-by-slot magnetic sensors at two locations: the Scatter Creek Safety Rest Area and Nisqually Weigh Station. The pilot system architecture (Figure 3-3 and Figure 3-4) demonstrated how field sensors communicated occupancy data to a central server, with data pulled at 5-minute intervals for processing, storage, and visualization.

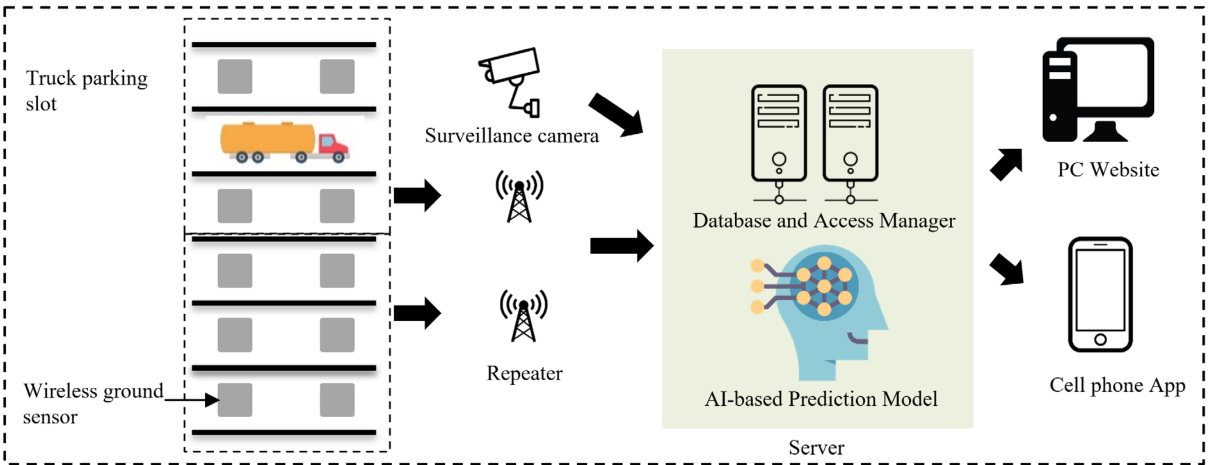

Figure 3-3 Pilot TPIMS architecture (Yang et al. 2022)

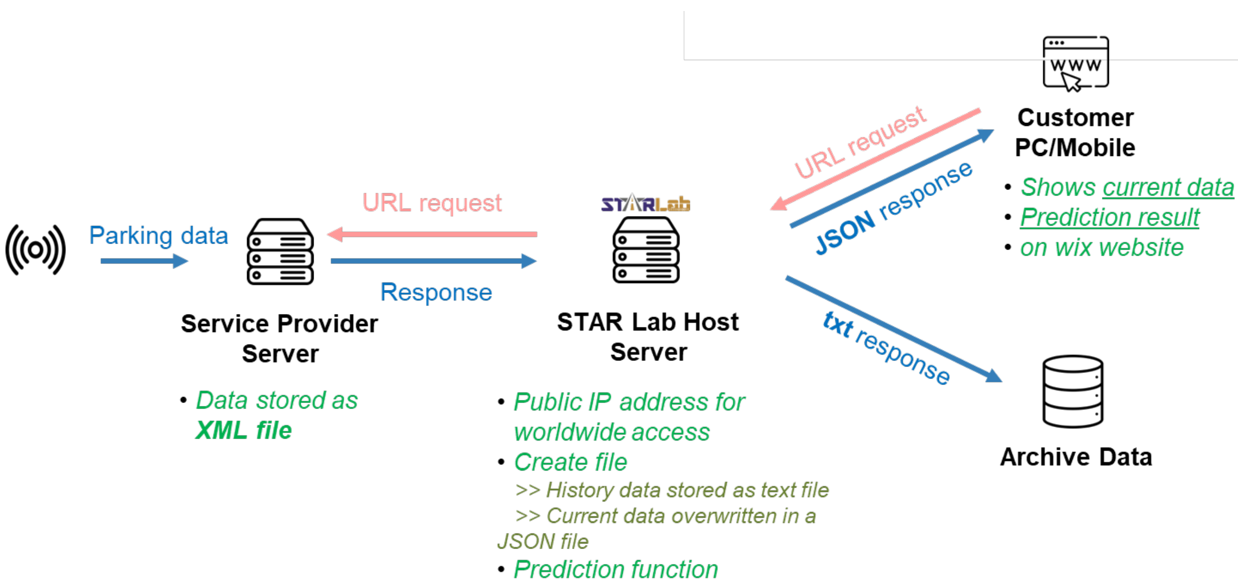

Figure 3-4 Pilot system architecture

For prediction, the pilot project implemented a non-homogeneous Poisson (NHP) process as part of the deployment pipeline. This statistical model estimates how parking availability changes based on historical temporal patterns. At each prediction cycle, the system first retrieves the current number of available spaces from the live feed and then computes a time-varying rate parameter derived from historical data from the same weekday and hour in past weeks. By comparing availability across consecutive hours in previous weeks, the model identifies typical arrival/departure trends and sums these trends across the selected prediction horizon (e.g., 1-hour, 5-hours, or 12-hours ahead). The expected change is scaled by total lot capacity and added to the current availability to estimate future open spaces. Predicted availability is then published as JSON files for use in the online display system.

Results from the pilot were integrated into the WSDOT Truck Parking Information Website (Figure 3-5). Each site's page displayed site information, real-time slot availability, predicted future availability, and a map-based visualization of truck parking spaces, offering drivers a clearer view of parking conditions.

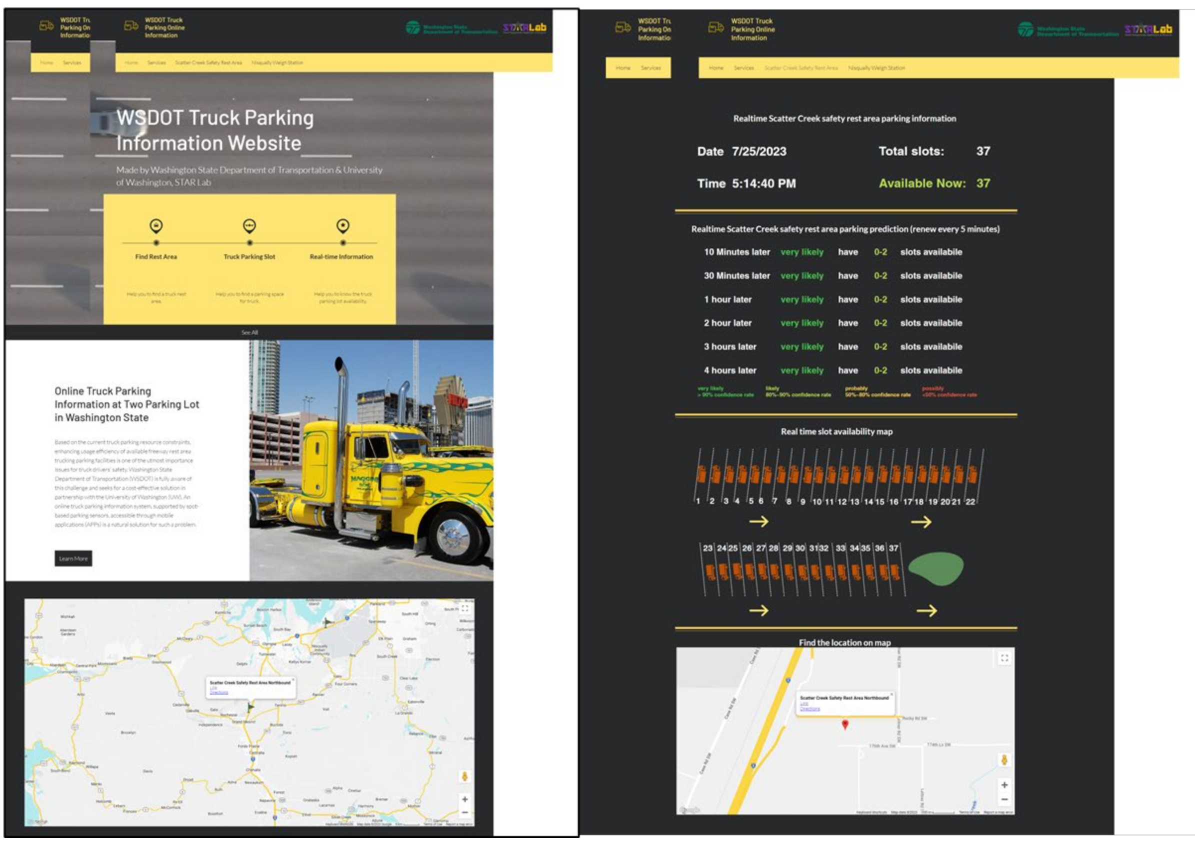

Figure 3-5 WSDOT Truck Parking Information Website

While the pilot project demonstrated the feasibility of real-time truck parking monitoring and short-term prediction, several practical and technical challenges were identified during deployment. These limitations highlight important considerations for future system expansion, long-term operations, and statewide scalability. Key limitations included the following:

- **Sensor reliability and maintenance burden:** Slot-by-slot magnetic sensors provided high accuracy but required substantial maintenance. Sensors frequently failed because of environmental exposure, installation challenges, and power issues, making long-term operation costly.
- **Scalability and installation constraints:** The pilot approach required one sensor per parking stall, leading to high installation labor and material costs. Scaling this approach statewide would not be cost-effective.
- **Data output format limitations:** The pilot system displayed results only on a website interface. For broader integration—such as navigation apps, commercial fleet systems, or third-party services—the system would require standardized, flexible APIs and machine-readable formats (e.g., JSON/GTFS-like feeds).
- **Limited generalization of prediction model:** The NHP-based prediction relied heavily on historical patterns and might not adapt well to disruptions, atypical traffic patterns, or seasonal changes without continuous recalibration.

4. SYSTEM REQUIREMENTS AND DESIGN FRAMEWORK

4.1 STAKEHOLDER NEEDS ASSESSMENT

4.1.1 Vendor Engagement

Vendor outreach began in June 2022 to survey the full range of truck-parking detection technologies, including radar-based, video-based, and in-ground sensor systems. The companies contacted included TSPS, Wavetronix, RTMS, TCS International, PNI, CIVIC Smart, SensIT/Nwave, Indect, and Sensys. Across all engagements, the team followed a consistent framework to collect both sensor-level and system-level information such as detection accuracy, sensing principle, communication architecture, power requirements, installation needs, maintenance expectations, data volume, communication costs, availability of APIs and web tools, and monitoring capabilities. Vendors were also asked to share details of any demonstration sites to help evaluate real-world performance and scalability

Demo Site

TSPS and Wavetronix provided example configurations of their demonstration sites at Indian John Hill EB, and State DOT Weigh Station SCALE DuPont, respectively, to illustrate how their systems operate in a real truck-parking environment. The demo layout included radar units placed at both the entrance and exit of the lot to capture vehicle movements, supported by a pan-tilt-zoom (PTZ) camera positioned centrally to provide visual verification and supplemental monitoring.

TCS also shared detailed information on its demonstration site at Maytown, including potential camera mounting locations and expected coverage patterns. Each vendor also supplied estimates for equipment, installation, and monthly operations and maintenance costs, along with a schematic showing the sensor coverage areas and field of view.

Additional details are provided in Appendix 4-1. Several practical considerations were identified for the demonstration sites study, including the need to determine whether existing cameras at weigh stations can serve as validation cameras, and the requirement to obtain raw sensor data formats from vendors to support API development and integration into the statewide TPIMS system.

4.1.2 Internal WSDOT Engagement with IT

During internal coordination with WSDOT's IT team, the primary focus was on identifying the long-term hosting environment for the truck-parking system and planning the migration of system components to WSDOT-managed servers. IT staff outlined the requirements for moving the current data ingestion, processing, and API services from temporary or contractor-managed environments into WSDOT's production infrastructure. Discussions covered server specifications, security standards, network access, database configuration, and integration with existing WSDOT authentication and monitoring tools. The IT team also provided guidance on the migration timeline, necessary documentation, and compliance with statewide security and data retention policies. This engagement ensured that the system architecture and deployment plan would align with WSDOT's operational frameworks and support a smooth transition into the agency's long-term hosting environment.

4.1.3 External Stakeholder Engagement

The project team actively engaged with external stakeholders to align the development and deployment of the truck parking prediction system with operational needs, partner expectations, and user requirements. Key engagements involved WSDOT program staff, Transpo Group as the INFRA grant consultant team, and industry partners such as DriveWyze.

Beginning in July 2024, multiple coordination meetings were held with WSDOT and Transpo Group to discuss the truck parking prediction algorithm in the context of planned INFRA-funded deployments. These discussions focused on algorithm reliability, performance under over-capacity parking conditions, data inputs, and overall readiness for operational use. Follow-up meetings in November 2024 allowed stakeholders to ask detailed technical questions regarding the prediction methodology, system assumptions, and limitations. Across these sessions, all questions were addressed and formally documented, providing transparency and building confidence in the algorithm's design and planned enhancements.

In January 2025, WSDOT invited the University of Washington team to participate in a collaboration with DriveWyze, a private-sector partner focused on delivering information directly to truck drivers. The discussion centered on developing a DriveWyze application that would leverage the APIs developed under this project to disseminate real-time and predictive truck parking information. This engagement established an ongoing bi-weekly coordination process and reinforced the project's role as a foundational data provider for third-party traveler information platforms.

Additional stakeholder engagement in May 2025 focused on how truck parking data should be presented to drivers. Discussions emphasized the importance of a user-friendly, visually engaging interface, with a strong preference for mobile-oriented access to support trip planning rather than desktop-based tools. Feedback from this session informed design considerations for future traveler-facing applications and reinforced the need for clear, intuitive presentation of both real-time and predictive information.

Overall, these external stakeholder engagements ensured that the truck parking prediction system was technically sound, operationally relevant, and aligned with partner deployment strategies and end-user expectations.

4.2 DETECTION SYSTEM EVALUATION AND SITE CHARACTERISTICS

4.2.1 Review of Candidate Detection Systems

For the TPIMS deployment, three candidate detection solutions were selected for evaluation—TSPS (in/out), Wavetronix (in/out), and TCS (video-based slot detection). Each system was reviewed using two complementary perspectives: sensor-level information (e.g., sensing principle, accuracy, communication method, power requirements, and installation needs) and system-level information (e.g., data interfaces, maintenance needs, integration requirements, and overall operational considerations). Details are shown in Table 4-1 and Table 4-2.

Table 4-1 Sensor information summarized

Sensor Items	TPSP (in/out)	Wavetronix (in/out)	TCS (VDO slot)
Parking Sensor Type/Name	Omnisight MegaRadar	Wavetronix Smartsensor HD	Mistall solar-powered time-lapse camera
Detection Accuracy	98%	95%+	High accuracy with AI-based video processing (slot-level)
Detetion Principle (radar/video/geomagnetic/ultrasnoic)	Radar/Video Edge Detection	Radar	Video imaging + AI/ML for vehicle detection & dwell time tracking
In-lot Communication Component and Protocol	CAT6 - TCP/IP	Ethernet (10/100 Mbps) over single twisted pair	4G LTE cellular modem (no on-site network required)
Lot-to-Server Communication Component and Protocol	HTTPS	Expanse Arc → Internet → HTTPS/API	Cellular → Cloud server via HTTPS / REST API (JSON)
Communication Distance	Wired - with extender is 1km	Standard Ethernet distance (up to ~100 m per segment; extendable with cabinet devices)	Anywhere with 4G cellular coverage
Power Supply of Sensor/Communication Components	PoE Switch or PoE injector	120 or Solar	Solar panel + battery (off-grid)
Installation Requirement	No controllers required	Mounted roadside/side-fire; requires power + Ethernet to cabinet; non-intrusive installation	Mount camera on existing pole/lamp post/building with view of spaces; attach solar panel + modem
Maintenance Frequency/Costs/Requirements	$50 / month / radar unit	None	Minimal; solar/battery system; remote monitoring; recurring software/internet fee required
Internet Connectivity Requirement/Costs	Internet is required	Ethernet	Uses 4G cellular; included in recurring cost (~$14,976/year per site for software + data + services)
Input Voltage Requirement	PoE Switch is 48V DC	24 volt	Solar-powered (battery-supported), no external voltage required
Power Access Point Distance Requirement (~ ft)	None	Under 2000 ft	None (off-grid solar system)

Sensor Items	TPSP (in/out)	Wavetronix (in/out)	TCS (VDO slot)
What Is Needed for the Installation for the Sensing Components	Simple single wire connection - RAM mounts included with device	Pole, and mount	Camera unit, solar panel, battery pack, 4G modem, mounting hardware; junction box if multiple cameras
Other Comments	Installation is flexible. The radar units only need a PoE CAT6 connection with internet access. If there is existing network you can install just using PoE injectors. Otherwise you can use a PoE switch.	Wavetronix is the most reliable sensor in all conditions and requires minimal to no maintenance. Wavetronix fully believes that once you purchase the device you own all the data, so there are no ongoing payments for software or data collection/usage.	Asset-light, no existing infrastructure needed; captures dwell time, counts, analytics; up to 3 cameras per solar mount

Table 4-2 System information summarized

System information	TPSP (in/out)	Wavetronix (in/out)	TCS (VDO slot)
Clear Picture for Entire Sensing System for a Truck Parking Lot	The radars are capable of capturing images / video of entry/exit events. We recommend a PTZ camera so that we can ground truth the lot.	Two XP20 radar sensors at entrance & exit connected to Expanse Arc cabinet unit; data streamed via Internet to API	Solar-powered cameras capture slot-level images every 3 minutes; AI identifies occupancy and dwell time; data sent via cellular to cloud dashboard and APIs
System Type (entrance/exit, space-by-space or hybird)	Entrance/Exit recommended. We can also do spot specific.	Entrance/exit (in/out counting)	Slot-by-slot occupancy monitoring
System Components	MegaRadar units are operated without a controller. They just need power and internet	XP 20 Arcsystem	Mistall cameras, solar panel + battery, 4G cellular modem, cloud backend database, TCS/Mistall analytics dashboard, REST APIs
Power Supply Requirement/Voltage of Different Components (except sensor)	poe 802.3at	No	Solar panel + internal battery (no AC power required except for optional junction boxes)

System information	TPSP (in/out)	Wavetronix (in/out)	TCS (VDO slot)
Data Type and Volume (per day)	Counts are real time	Based on intervals and sever Capacity	Time-lapse images every 3 minutes; processed occupancy/dwell-time events; moderate cellular data usage
Communication Cost (per month)	Whatever internet costs	None	Included in recurring yearly fee (~$14,976/year per site)
Communication Bandwidth Requirement	Edge detection requires very little internet bandwidth. Radars can be configured to transmit video 24x7 which would require more bandwidth (not enabled by default)	Low—edge radar detection generates small data packets	Moderate (periodic time-lapse images + metadata upload)
Data Services API Available or Not	Yes, data api is available	Yes	Yes—REST API (JSON), documentation + sample code (Python, C#)
Web/App Available or Not	Web app for doing counting / ground truthing and another web app for managing the hardware	No	Yes—mobile apps (iOS/Android), mobile web, IVR phone line, SMS, digital signage integration
Maintenance Information Available or not	No maintenance is required for the radar units or PTZ camera	Yes	Provided via vendor; cloud dashboard shows system health
Sensor Working Status Monitoring System Available or Not	Yes sensors are monitored to make sure they are working properly	Yes	Yes—remote system health via cloud platform
Special Needs for Different Components to Install/Use/Maintain	No special needs	No	Requires clear camera view; depends on solar exposure; multiple cameras share one solar + modem mount; cellular coverage required

4.2.2 Design Constraints for Washington Sites

Our objective was to determine which detection system would be most suitable for each site by assessing the physical feasibility of installation. This assessment was conducted primarily through off-site investigations using available secondary information to understand existing infrastructure conditions. To evaluate installation feasibility, we examined two main sources of information:

- Google Street View: This was used to identify the presence of light poles at site entrances and exits, which are critical for mounting in/out counting sensors. We also checked whether a server room or equipment cabinet already existed on-site to support system deployment.
- Plan and Drawing Review: Site plans were inspected to understand existing power and communication infrastructure, including AC electrical layouts and cabinet configurations, which affect installation complexity and system compatibility.

Additional detailed findings for each site are provided in Appendix 4-1.

From the site assessment results, several notable issues were identified that could affect the feasibility of sensor deployment:

1. **Shared entrances/exits with passenger vehicles**, making it difficult to isolate truck movements for accurate in/out counting (Figure 4-1).

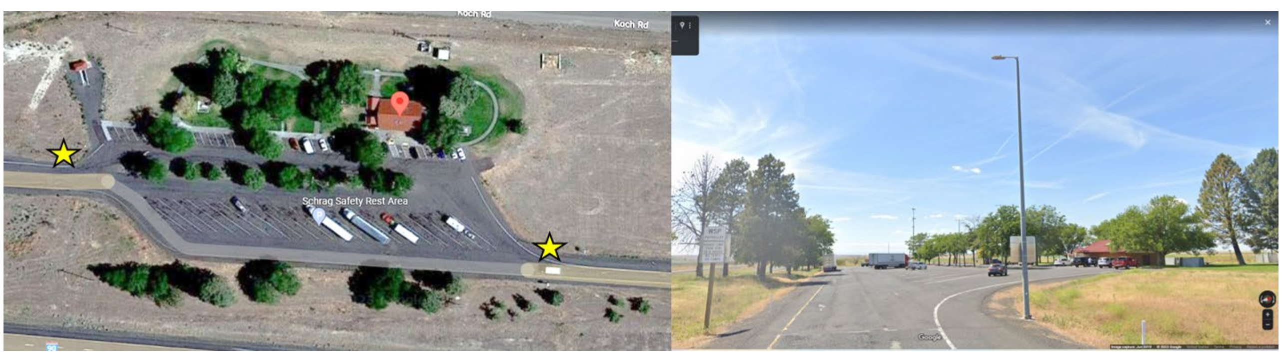

Figure 4-1 Shared entrances/exits with passenger vehicles at Schrag WB.

2. **Limited ground-level imagery,** with only aerial photos available at some locations (e.g., Toutle River NB, Ridgefield PoE).
3. **Separate truck and RV parking areas,** which complicate sensor placement and stall allocation (e.g., Smokey Point NB and Smokey Point SB (Figure 4-2)).

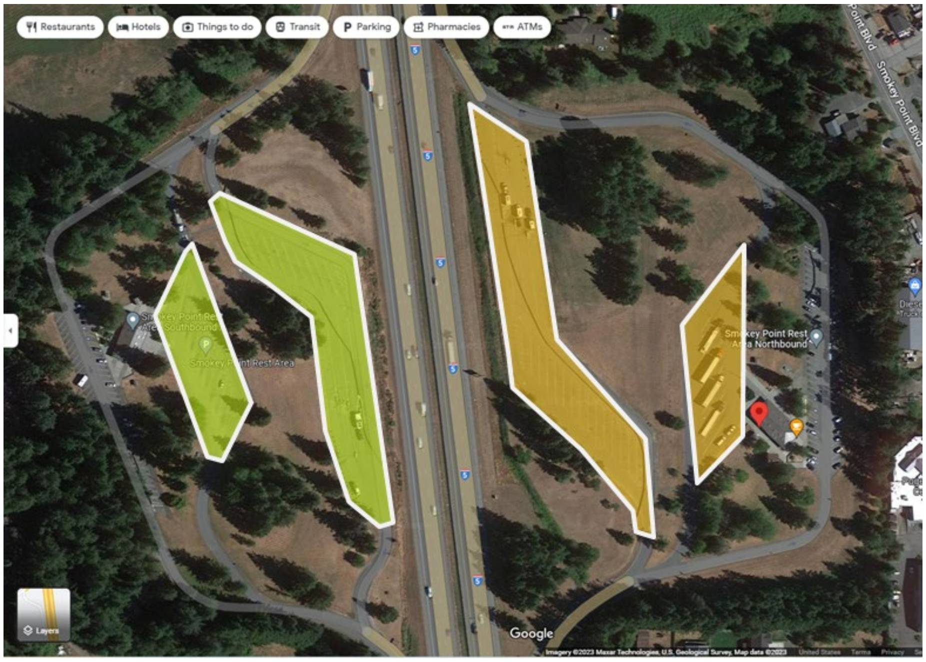

Figure 4-2 Separate truck and RV parking areas at Smokey Point NB and Smokey Point SB

4. **Uncertain cabinet status,** where existing equipment appears inactive or cannot be confirmed from Street View, requiring on-site inspection (e.g., Indian John Hill WB, SeaTac SB (Figure 4-3)).

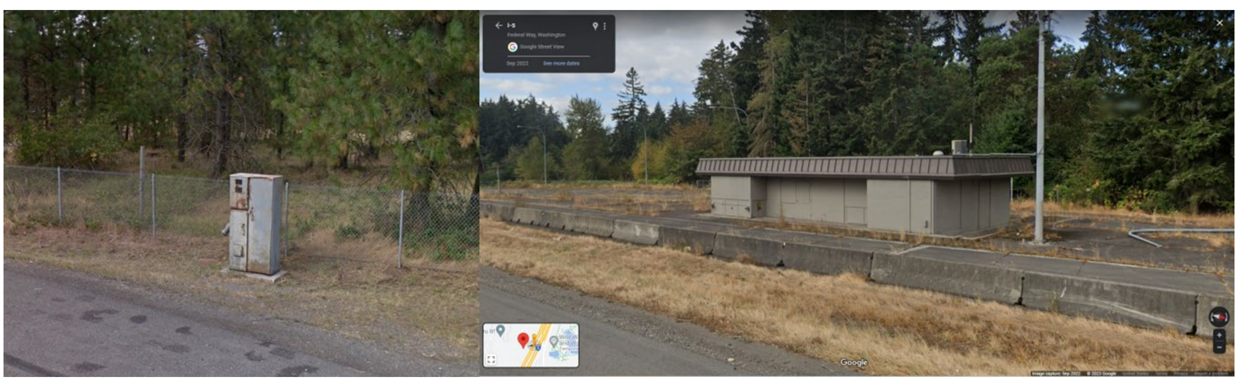

Figure 4-3 Uncertain cabinet status at Indian John Hill WB (left),and SeaTac SB (right).

4.2.3 Site Evaluation Methodology

We utilized a two-step selection framework. This framework considered both system-level performance requirements and site-specific operational constraints.

Step 1: Assign a Technology Based on the Level of Precision Needed

Each Washington truck-parking site differs in parking demand, operational risk, and the potential consequences of inaccurate reporting. To identify the most suitable detection technology, we evaluated how much precision would be required at each location. Sites with high utilization, chronic shortages, or significant safety implications (e.g., long rural gaps between rest areas) require high-precision sensing because even a small counting error can lead to large discrepancies over time. These locations typically favor stall-level (slot-by-slot) detection, which minimizes cumulative error and provides real-time ground truth at the stall level. In contrast, sites with moderate or low demand—and where an occasional error is less consequential—can be effectively served by in/out counting systems, which offer scalable coverage at significantly lower cost. Key metrics included the following:

1. Cost of Miss (C(x,y)). This metric estimated the operational impact if a truck driver bypassed a parking site because the system incorrectly reported the lot as full or nearly full. For each site, we computed how many additional miles a driver would have to travel to reach the nearest available parking alternatives. The "cost" was calculated by identifying the **three closest neighboring sites** and measuring the distance needed to reach each of them. Sites with long inter-site distances—especially in rural corridors—were assigned a **high cost of miss**, reflecting the safety and fatigue risks associated with extended travel without rest. The sample results are presented in Table 4-3, and the complete list of sites and calculated distances is provided in Appendix 4-2.

2. Parking Demand (D(t)). We used historical parking demand data from Washington Truck Parking Assessment from 2022 (Washington Truck Parking Assessment) to quantify how frequently each lot reached or approached capacity (Figure 4-4). High-demand sites experience frequent occupancy peaks and are sensitive to even small errors. Low-demand sites have more operational flexibility and can tolerate less precise detection. For this analysis, each site's demand level was classified into three categories: very high, high, and medium.

Table 4-3 Sample cost of miss calculation from distance to closet point

Area	Point 1 Location	Point 1 Dist (mi)	Point 2 Location	Point 2 Dist (mi)	Point 3 Location	Point 3 Dist (mi)	Average Distance
Gee Creek NB	Ridgefield PoE	4.06	Gee Creek SB	4.11	Toutle River NB	43.27	17.15
Gee Creek SB	Gee Creek NB	5.59	Ridgefield PoE	5.92	Toutle River NB	45.13	18.88
Ridgefield PoE	Gee Creek SB	5.32	Gee Creek NB	10.83	Toutle River NB	39.21	18.45
Toutle River NB	Toutle River SB	5.60	Maytown SB	42.63	Gee creek SB	47.34	31.86
Toutle River SB	Toutle River NB	4.17	Gee Creek SB	41.79	Ridgefield PoE	42.02	29.33
Maytown SB	Toutle River SB	38.76	Toutle River NB	42.88	SeaTac NB	57.34	46.33
SeaTac NB	SeaTac SB	4.15	Maytown SB	51.45	Smokey Point NB	66.63	40.74
SeaTac SB	SeaTac NB	6.73	Maytown SB	47.31	Smokey Point NB	73.36	42.47

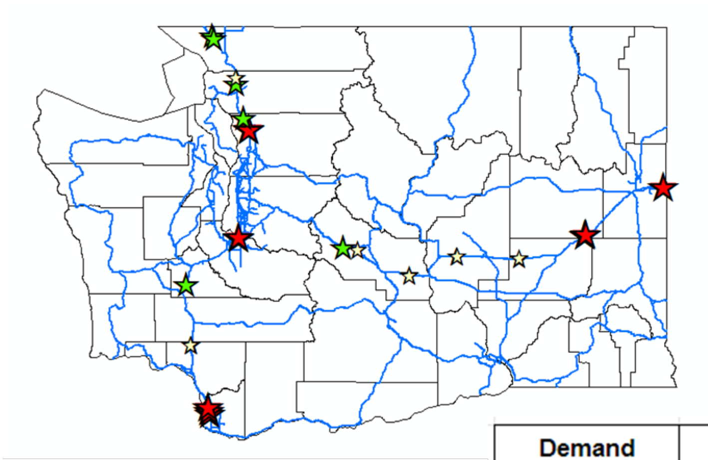

Demand	Count
Mid	12
High	6
Very High	10

Figure 4-4 Parking demand

After computing the Cost of Miss and Parking Demand scores, each site was assigned to one of three recommendation categories (Figure 4-5):

*a. High Demand + High Cost of Miss***:** These locations are heavily used and have limited backup parking options nearby.

- Detection accuracy requirement: very high
- Error tolerance: very low
- Parking information necessity: high (for safety and operational planning)
- Recommendation: Slot-by-slot detection
  (Provides the highest fidelity and prevents cumulative counting errors).

b. High Demand + Low Cost of Miss: These sites are busy but surrounded by other nearby alternatives.

- Detection accuracy requirement: moderate
- Error tolerance: higher (because drivers have fallback options)
- Parking information necessity: high (to reduce cruising and improve lot turnover)
- Recommendation: In/out counting detection
  (Provides good scalability while meeting operational needs).

c. Low Demand + High Cost of Miss: These sites are not frequently congested but are highly isolated, meaning drivers face long distances to the next available rest area.

- Detection accuracy requirement: high (to prevent unnecessary bypassing)
- Error tolerance: very low
- Parking information necessity: high (because of limited regional supply)
- Recommendation: Slot-by-slot detection
  (Ensures accuracy even when occupancy is low but operational risk is high).

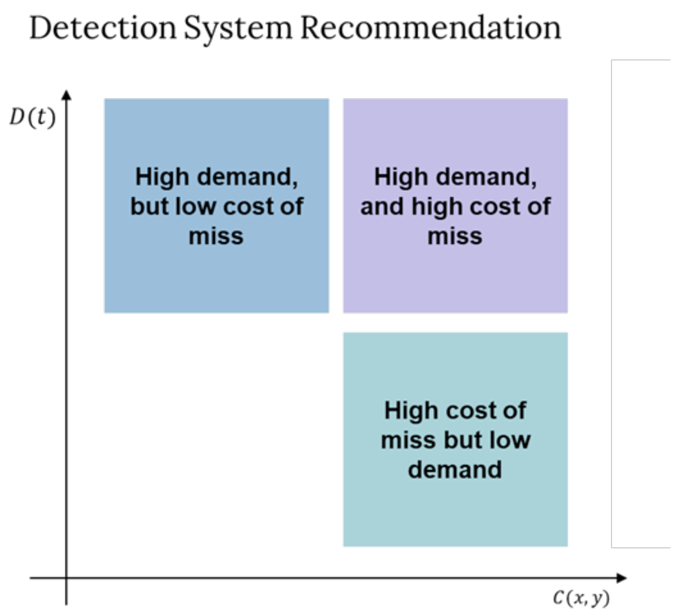

Figure 4-5 Framework for determining appropriate detection technology based on demand and cost-of-miss.

In addition, we further examined two supplementary metrics to refine the assessment and better understand installation feasibility:

3. Area of the Parking Lot The physical size of the lot directly affects installation complexity and cost. Smaller lots typically require fewer sensors, shorter cable runs, and have lower risk of cumulative error—resulting in lower installation cost. Larger lots, by contrast, often require longer conduit and cabling, additional mounting infrastructure, and more labor, which can substantially increase installation cost and complexity for both in/out and stall-level systems. For this analysis, each site's size was classified into three categories: small, medium, and large, based on number of stalls **(**Figure 4-6**)**.

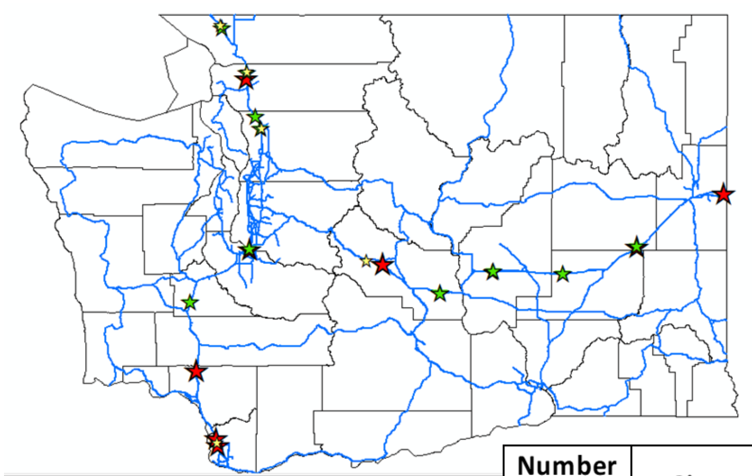

Number of Stall	Size	Count
<= 11	S	5
12-19	M	13
>= 20	L	10

Figure 4-6 Area of parking lot

4. Area Type (Rural vs. Urban) Each site was classified as either rural, urban, or urbanized, based on FHWA Highway Urban Area definitions (Figure 4-7). This classification was important because it affects the availability and reliability of supporting infrastructure needed for sensor installation.

- Urbanized areas: populations ≥ 50,000
- Urban areas: populations between 5,000 and 49,999
- Rural areas: populations < 5,000.

Rural sites often pose greater installation challenges, including limited electrical service, fewer internet service providers, and longer distances to connect utilities. In contrast, sites located in urban or urbanized areas generally have more accessible power and communication infrastructure, reducing installation barriers and overall deployment complexity.

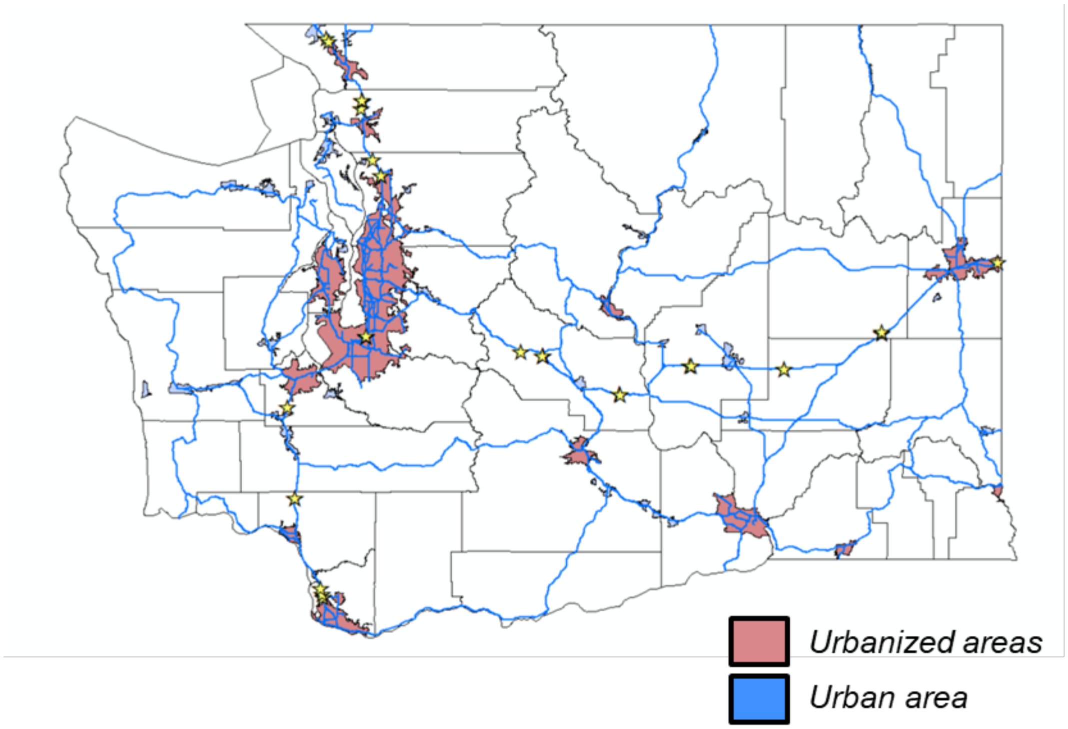

Figure 4-7 Area type data from WSDOT Geospatial Open Data Portal (WSDOT Geospatial Open Data Portal, accessed on Oct 5, 2022).

<u>Step 2: Conduct a Site-by-Site Feasibility Assessment</u>

Once the precision level had been defined, each site was evaluated individually to determine the feasibility of installing and maintaining the recommended technology. This included examining the following:

- **Ease of installation:** pavement drilling limitations, pole locations, power access, line-of-sight, weather exposure
- **Ease of maintenance:** sensor accessibility, battery replacement cycles, susceptibility to damage
- **Total lifecycle cost:** hardware, installation, communications, and recurring maintenance needs
- **Benefit-cost analysis** evaluating whether the operational benefits of improved accuracy justify the installation and maintenance costs for each site.
- **Conduit availability** assessing whether existing electrical or communication conduits are present—or whether new trenching and conduit installation would be required—which may significantly impact cost and feasibility.

This step incorporated aerial imagery, ground photography, and on-site conditions to confirm that the proposed technology could be effectively deployed and would remain reliable over the long term. After applying the two-step framework, each technology was assessed for its real-world feasibility across WSDOT sites, considering factors such as weather exposure, site geometry, maintenance access, and communication reliability (Table 4-4). This would ensure that the selected systems were not only accurate in theory but also practical, scalable, and sustainable for long-term statewide deployment.

Table 4-4 Comparison across detection alternatives

Criteria	Slot-by-Slot (Magnetic/Rada-Based)	Shot-by-Slot (Video-Based)	In/Out Counting (Radar + Camera Calibration)
Precision / Accuracy	Very high; no cumulative error; strong performance even under atypical truck behavior	Moderate; affected by lighting, weather, occlusion	Moderate; improved with camera calibration; prone to cumulative error
Installation Requirements	Installed at each stall (in-pavement or surface-mount); higher effort; lot reconfiguration sensitive	Requires elevated mounting with clear views; minimal ground work	Installed at entrance and exit only; minimal field work
Maintenance Needs	Distributed maintenance across many sensors; battery replacement cycles	Requires camera cleaning, repositioning, and view maintenance	Maintenance concentrated at access-point sensors; periodic ground truthing
Environmental Sensitivity	Magnetic sensors affected by cross-interference; radar stable in all weather	Sensitive to rain, fog, snow, night-time lighting	Radar performs well in all weather; cameras used for verification
Scalability	Scales poorly for large facilities because of cost per stall	Moderate scalability (1 camera per many stalls)	Scales extremely well; cost independent of stall count
Power & Communications	Often battery/LoRa; may require multiple gateways	Solar + cellular feasible; moderate bandwidth needs	Radar typically requires AC power and Ethernet/cellular; low bandwidth
Best Use Case	**High-demand, high-importance locations needing maximum precision**	**Medium-demand locations with moderate precision needs and feasible visibility**	**Low-demand or highly constrained sites where cost and scalability dominate**

4.2.4 Key Findings from Site Assessments

First Round of Recommendations

The first round relied solely on secondary information gathered from publicly available sources—such as Google Street View, aerial imagery, site plans, and the Washington Truck Parking Assessment—to evaluate each site's operational needs and classify them based on demand, cost of miss, area type, and lot size.

Our site assessment produced several key findings that guided the selection of detection technologies. First, we developed a ranking framework that used Cost of Miss and Parking Demand to cluster all sites into four tiers representing different levels of required detection precision. Each site received a score for Cost of Miss (11–28 miles = 1, 28–44 miles = 2, and >44 miles = 3) and for Parking Demand (medium = 1, high = 2, very high = 3). Figure 4-8 shows analysis results for Cost of Miss calculation with tier cluster.

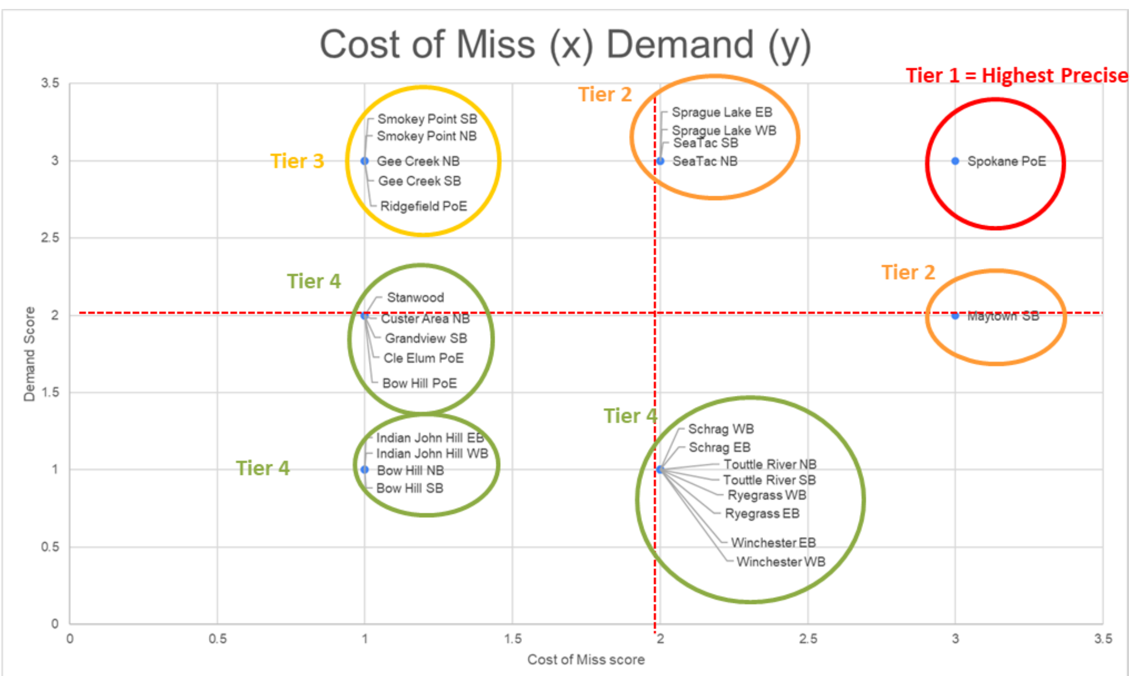

Figure 4-8 Cost of miss and demand analysis results

To support technology selection, Table 4-5 provides a comparison of the three candidate detection approaches—slot-by-slot (magnetic/radar-based), slot-by-slot (video-based), and in/out counting with camera calibration. The table outlines key attributes such as precision requirements, conduit needs, maintenance considerations, installation constraints, applicability under different demand levels, and sensitivity to cost-of-miss impacts.

After assigning sites to tiers, we refined the prioritization by ranking sites within each tier based on additional factors, including area type, parking lot size, demand, and cost-of-miss score. Smaller lots require greater accuracy because of limited error tolerance, while urbanized sites benefit from higher-precision detection to support safety and traffic management needs. This tiered and ranked evaluation framework provided a systematic method for determining which sites would require high-precision stall-level detection and which could be effectively served by lower-cost in/out counting systems. Table 4-6 shows the ranking results.

Table 4-5 Comparison of key characteristics for candidate detection technologies for first round recommendations

Key Characteristics	**Slot-by-Slot (Magnetic/Rada-Based)**	**Slot-by-Sot (VDO)**	**In/Out Counting (with Camera Calibration)**
Level of Precision	• High • Low error accumulation	• Moderate • Error accumulation	• Moderate • Higher accuracy with camera calibration
Conduit Availability	Required	Required	Not required
Maintenance Needs	Required at each detector	Required at each detector	Required at radar unit
Installation	Out of the pavement (lot configuration changes)	A location that can capture most of the entire parking area	At the entrance/exit of the parking lot
Demand	High parking demand	High to low parking demand	Low parking demand (high with camera calibration)
Cost of Miss	High cost of miss	High to low cost of miss	Low cost of miss (high with camera calibration)

Table 4-6 Ranking results

Region	Interstate	Mile Post	Direction	Name	# of Stalls	SIZE	Demand	Average Distance to Cost of Miss (mile)	Area Type	Tier	Site Inspection Light Pole @ Entrance	Site Inspection Light Pole @ Exit	Site Inspection Server Room /Big Cabinet	Site Inspection Conduit	Plan Inspection Result	1st round Recommendation System
ER	090	299	WB	Spokane PoE	21	L	Very High	77.67	Urbanized	1	0	Y	Y	NA	NA	Slot-by-slot (Magnetic/Rada-based)
NWR	005	188	SB	SeaTac SB	14	M	Very High	42.47	Urbanized	2	Y	N	Y	NA	NA	Slot-by-slot (Magnetic/Rada-based)
NWR	005	140	NB	SeaTac NB	22	L	Very High	40.74	Urbanized	2	Y	Y	Y	NA	NA	Slot-by-slot (Magnetic/Rada-based)
ER	090	243	WB	Sprague Lake WB	13	M	Very High	37.3	NA	2	Y*	Y*	N	NA	NA	Slot-by-slot (Magnetic/Rada-based)
ER	090	242	EB	Sprague Lake EB	20	L	Very High	36.55	NA	2	Y*	Y*	N	NA	NA	Slot-by-slot (Magnetic/Rada-based)
OR	005	93	SB	Maytown SB	13	M	High	46.33	NA	2	Y	Y*	N	NA	NA	Slot-by-slot (Magnetic/Rada-based)
NWR	005	207	NB	Smokey Point NB	11	S	Very High	13.64	Urbanized	3	Y	Y	N	NA	NA	Slot-by-slot (VDO) OR in/out counting with camera calibration
NWR	005	207	SB	Smokey Point SB	12	M	Very High	15.13	Urbanized	3	Y	Y	N	NA	NA	Slot-by-slot (VDO) OR in/out counting with camera calibration
SW	005	11	NB	Gee Creek NB	22	L	Very High	17.15	Urbanized	3	Y	Y	N	NA	NA	Slot-by-slot (VDO) OR in/out counting with camera calibration
SW	005	13	SB	Gee Creek SB	11	S	Very High	18.88	NA	3	Y*	Y*	N	NA	NA	Slot-by-slot (VDO) OR in/out counting with camera calibration

Region	Interstate	Mile Post	Direction	Name	# of Stalls	SIZE	Demand	Average Distance to Cost of Miss (mile)	Area Type	Tier	Site Inspection Light Pole @ Entrance	Site Inspection Light Pole @ Exit	Site Inspection Server Room /Big Cabinet	Site Inspection Conduit	Plan Inspection Result	1st round Recommendation System
SW	005	15	NB	Ridgefield PoE	32	L	Very High	18.45	NA	3	Not Sure	Not Sure	Y	NA	Has no illumination plan	Slot-by-slot (VDO) OR in/out counting with camera calibration
SCR	090	80	WB	Cle Elum PoE	7	S	High	27.92	NA	4	Y	N	N	NA	NA	Slot-by-slot (VDO) OR in/out counting with camera calibration
NWR	005	269	SB	Grandview SB	11	S	High	23.46	NA	4	Y*	N	Not Sure	NA	NA	Slot-by-slot (VDO) OR in/out counting with camera calibration
NWR	005	268	NB	Custer Area NB	18	M	High	26.96	NA	4	Y*	N	N	NA	NA	Slot-by-slot (VDO) OR in/out counting with camera calibration
NWR	005	214	NB	Stanwood	13	M	High	14.7	NA	4	Y	Y	Y	NA	Has Electrical Service Cabinet	Slot-by-slot (VDO) OR in/out counting with camera calibration
NWR	005	235	SB	Bow Hill PoE	20	L	High	15.02	NA	4	Y	N	Y	NA	Contradict with Google Maps	Slot-by-slot (VDO) OR in/out counting with camera calibration
ER	090	199	EB	Schrag EB	19	M	Mid	36.49	NA	4	Y*	Y*	Not Sure	NA	NA	in/out counting with camera calibration
NCR	090	162	WB	Winchester WB	12	M	Mid	34.04	NA	4	Y	Y	N	NA	NA	in/out counting with camera calibration
SCR	090	126	WB	Ryegrass WB	12	M	Mid	33.7	NA	4	Y*	Y*	N	NA	NA	in/out counting with camera calibration
SCR	090	126	EB	Ryegrass EB	12	M	Mid	32.96	NA	4	Y	Y	N	NA	Not sure if the same direction?	in/out counting with camera calibration

Region	Interstate	Mile Post	Direction	Name	# of Stalls	SIZE	Demand	Average Distance to Cost of Miss (mile)	Area Type	Tier	Site Inspection Light Pole @ Entrance	Site Inspection Light Pole @ Exit	Site Inspection Server Room /Big Cabinet	Site Inspection Conduit	Plan Inspection Result	1st round Recommendation System
															(EB/WB) Has only junction box and under-ground transformer	
NCR	090	161	EB	Winchester EB	12	M	Mid	28.77	NA	4	Y	N	N	NA	NA	in/out counting with camera calibration
ER	090	199	WB	Schrag WB	19	M	Mid	28.4	NA	4	Y*	Y*	Y	NA	NA	in/out counting with camera calibration
SW	005	54	NB	Toutle River NB	26	L	Mid	31.86	NA	4	Y*	N	N	NA	Has only Junction Box	in/out counting with camera calibration
SW	005	55	SB	Toutle River SB	25	L	Mid	29.33	NA	4	Y*	Y*	N	NA	NA	in/out counting with camera calibration
NWR	005	235	SB	Bow Hill SB	11	S	Mid	11.47	NA	4	Y	Y*	N	Y	NA	in/out counting
NWR	005	235	NB	Bow Hill NB	12	M	Mid	14.71	NA	4	N	Y*	N	Y	NA	in/out counting
SCR	090	89	EB	Indian John Hill EB	30	L	Mid	21.57	NA	4	Y*	Y*	N	NA	NA	in/out counting
SCR	090	89	WB	Indian John Hill WB	20	L	Mid	20.73	NA	4	Y*	Y*	Y	NA	NA	in/out counting

Second Round of Recommendations

The second round of recommendations built on the initial tiering framework but incorporated additional site inspections and a deeper evaluation of costs, installation feasibility, and long-term maintenance needs (Figure 4-9).

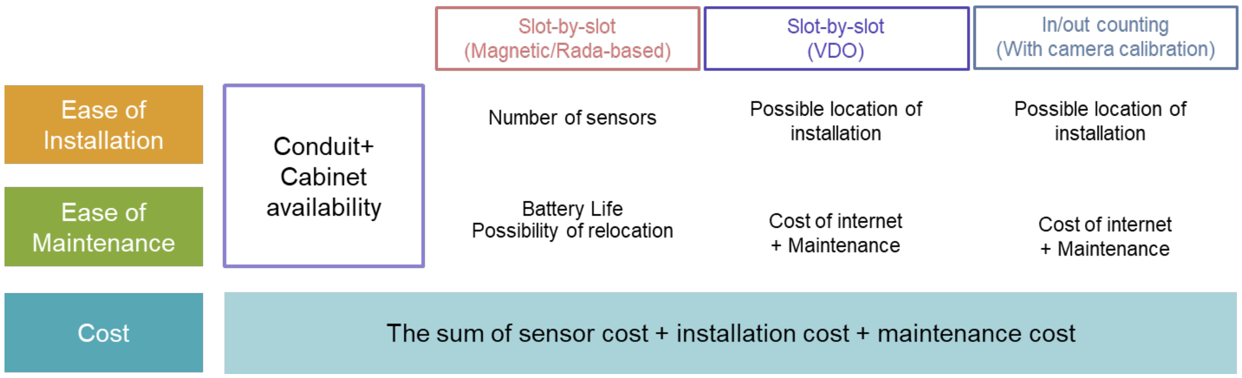

Figure 4-9 Criteria used for second-round technology evaluation

We evaluated installation cost and maintenance cost as two key dimensions for comparing the three candidate detection technologies. Figure 4-10 summarizes all cost components considered in this assessment, including conduit and cabinet installation, sensor installation labor, number of required sensors, pole installation needs, battery life, battery replacement cycles, potential relocation costs, and recurring service or communication fees. By aggregating these items, we were able to estimate the total lifecycle cost for each technology and determine which option would offer the best balance of practicality, maintainability, and long-term affordability for each site.

			Slot-by-slot (Magnetic/Rada-based)	Slot-by-slot (VDO)	In/out counting (With camera calibration)
Ease of Installation	1	Conduit/Cabinet Installation (USD)	TBD	TBD	TBD
	2	Sensor Installation cost (USD)	TBD	TBD	TBD
	3	Sum installation cost (USD)	(1) + (2)		
	4	Number of Sensor (unit)	✓	1-2 unit	2 unit
	5	Price per Sensor (USD)	TBD	TBD	TBD
	6	Sum Sensor Cost (USD)	(4) x (5)		
	7	Pole Installation (USD)		✓	✓
Ease of Maintenance	8	Battery Life (yr)	TBD		
	9	Price per Battery (USD)	TBD		
	10	Number of Sensor (unit)	✓		
	11	Sum Battery Cost (USD)	Service Life/(8) x (9) x (10)		
	12	Possibility of Relocation (yr)	TBD	TBD	
	13	Reinstallation Cost	TBD	TBD	
	14	Sum of Reinstallation Cost(USD)	Service Life/(12) x (13)		
	15	Internet Cost (USD)	39.99/month + etc.		
	16	Maintenance Cost (USD)	TBD	TBD	TBD
	17	Sum of Service Cost (USD)	(15) + (16)		

Figure 4-10. Cost components used to assess installation and maintenance requirements

We also proposed four sites for the pilot study, each chosen to represent different operational and physical characteristics:

1. **Spokane PoE** – A weigh station with very high parking demand and the highest cost of miss in the network (Figure 4-11).

# of Stalls	SIZE	Demand	Cost of Miss (mile)	Area Type	Tier
21	L	Very High	77.67	Urbanized	1

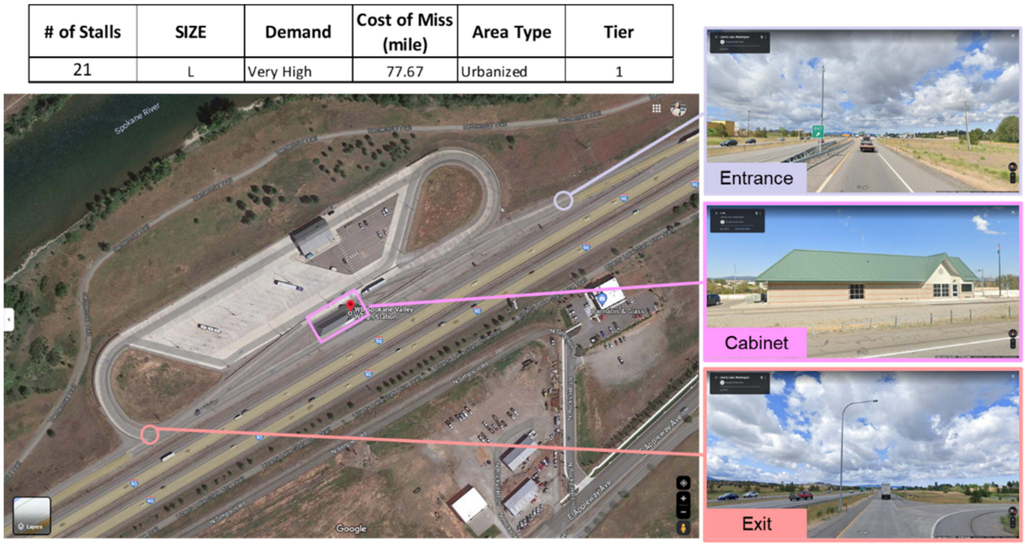

Figure 4-11 Proposed site for the pilot study: Spokane PoE.

2. **Maytown SB** – A high-demand site where trucks and passenger vehicles share the same exit, creating challenges for in/out detection (Figure 4-12).

# of Stalls	SIZE	Demand	Cost of Miss (mile)	Area Type	Tier
13	M	High	46.33		2

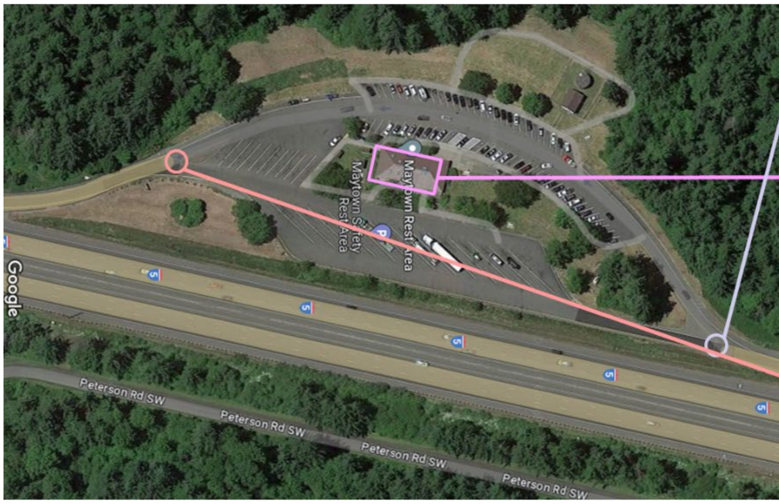

Figure 4-12 Proposed site for the pilot study: Maytown SB).

3. **Toutle River NB** – A site with a separate truck parking area, but where trucks and passenger vehicles also share the same exit, requiring careful sensor placement (Figure 4-13).

# of Stalls	SIZE	Demand	Cost of Miss (mile)	Area Type	Tier
26	L	Mid	31.86		4

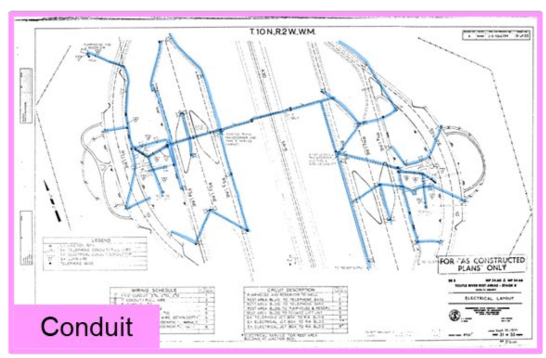

Figure 4-13 Proposed site for the pilot study: Toutle River NB.

4. **Bow Hill PoE** – A weigh station with high parking demand, offering an opportunity to evaluate system performance in a busy, controlled environment (Figure 4-14).

# of Stalls	SIZE	Demand	Cost of Miss (mile)	Area Type	Tier
20	L	High	15.02		4

Figure 4-14 Proposed site for the pilot study: Bow Hill PoE.

These sites collectively captured a diverse set of deployment conditions to fully assess the feasibility and performance of candidate detection technologies.

5. SYSTEM ARCHITECTURE AND IMPLEMENTATION

5.1 OVERALL SYSTEM ARCHITECTURE

The Washington State TPIMS architecture is designed to support real-time parking data ingestion, centralized data processing, prediction generation, and multi-platform dissemination. The system integrates data from multiple detection vendors and ensures that standardized, MAASTO-compliant parking information is made available to third-party developers, internal operators, and truck drivers through web and mobile applications.

Figure 5-1 illustrates the high-level communication pathway between the sensing systems, WSDOT servers, and external interfaces. Parking data generated by the field sensors flow from the service-provider cloud toward WSDOT through periodic URL requests issued by the WSDOT server. Upon receiving the vendor-specific JSON or binary response, WSDOT converts the payload into the MAASTO-defined data structure and exposes it through the public API. External applications—including the public website (npm/React), the prototype interface on Wix/Velo, and the internal Streamlit dashboard—retrieve real-time occupancy and prediction results from this API. The system also supports two-way communication with UW's prediction environment: WSDOT transmits the latest parking data to UW, and UW returns updated multi-horizon predictions, which are then embedded into the TPIMS API response.

As this project emphasized the development and testing of a strong prediction algorithm, all prediction functions initially run on UW servers. This setup allows rapid iteration, controlled experimentation, and detailed monitoring of model behavior. UW requires secure access to WSDOT servers to update prediction logic, manage versioning, and monitor potential anomalies or data inconsistencies.

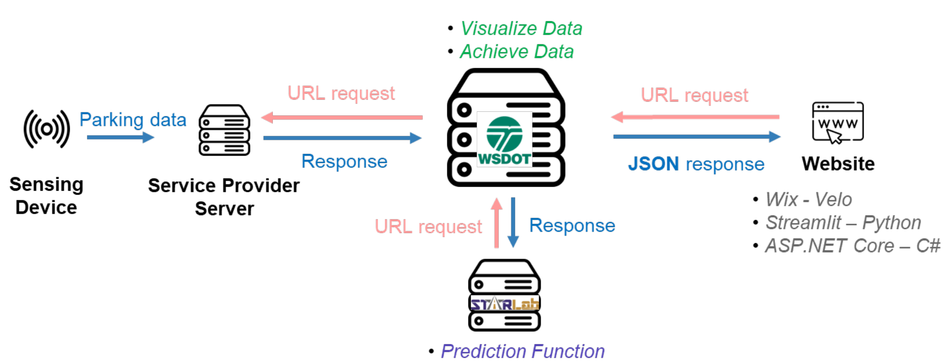

Figure 5-1 High-level system data exchange architecture between field sensors, WSDOT servers, UW prediction services, and public-facing applications.

Figure 5-2 provides a more detailed depiction of the operational environment during the interim (pilot) phase and the final (WSDOT-hosted) deployment. In the interim architecture, the TPIMS server—along with the prediction module—resides within the UW environment, while sensor data traverse both WSDOT and UW networks. WSDOT video feeds and Omnisight in/out sensor data flow through the WSDOT edge switch and firewall before reaching the UW-hosted system, where data ingestion, prediction, and API generation occur.

In the final architecture, TPIMS services are fully transitioned to a WSDOT-hosted environment. The field sensors transmit data through WSDOT's edge switch into a centralized TPIMS server housed within WSDOT's firewall. The prediction engine—after stabilization and validation—will be deployed on this TPIMS server, enabling WSDOT to operate the full pipeline internally. The Omnisight and Osprey service-provider clouds continue to supply standardized data feeds, which are aggregated, processed, and exposed through the TPIMS API. This architecture streamlines operations, simplifies maintenance, and reduces the number of cross-agency network dependencies while preserving UW's access for periodic updates and performance monitoring.

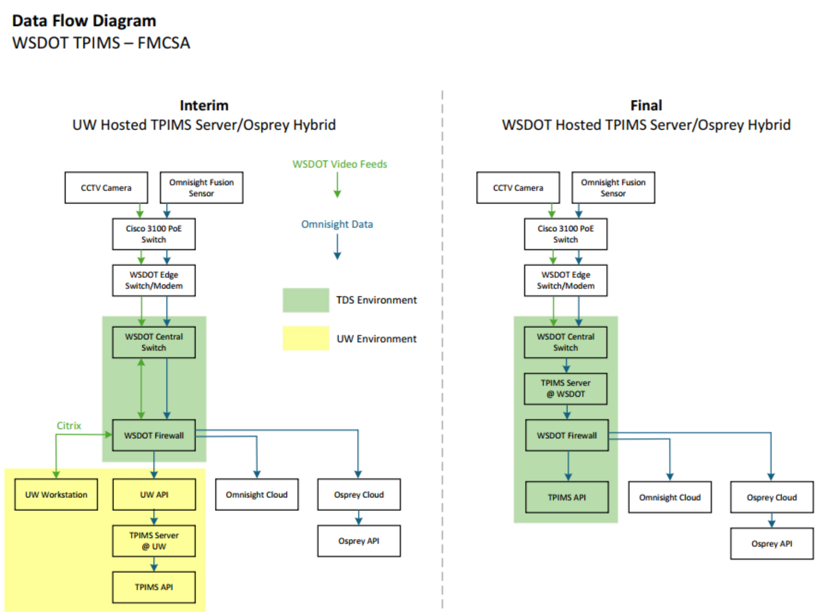

Figure 5-2 Interim and final TPIMS deployment architectures showing data pathways, network environments, and system responsibilities for UW-hosted and WSDOT-hosted configurations.

5.1.1 Selected Detection System Operation and Data Interface

Omnisight In/Out Sensors

The Omnisight in/out sensors generate an event-based data feed by pushing a JSON record each time a vehicle crosses the designated detection zone at the entry or exit of the parking lot. When a vehicle breaks the sensor's beam, the device immediately publishes a structured JSON message to the client endpoint, enabling real-time capture of arrivals and departures. The vendor provides multiple event endpoints—including detection zone events, screen line events, and target-detection events—which contain attributes such as timestamp, direction of movement, object ID, and classification metadata when available. These event messages form the foundation of the occupancy estimation logic, allowing the system to compute current availability by incrementing or decrementing counts based on detected entry and exit events.

The Omnisight cloud dashboard (Figure 5-3) provides a comprehensive interface for monitoring sensor activity, reviewing detection events, and managing data integrations. The main overview displays recent target volumes—aggregated by device and classified by vehicle type—which helps operators quickly identify traffic patterns and potential anomalies at each site. A dedicated "Parking View" presents the sensor coverage areas over a satellite map, allowing users to visually confirm the location of detection zones and ensure proper alignment of inbound and outbound screen lines. Detailed event logs list each detected object along with size estimates, classification, detection device, and timestamp, giving users the ability to audit system performance at a granular level. The dashboard also includes an "Integrations" tab where WSDOT or site operators can configure client endpoints and activate event triggers such as detection-zone entry/exit, screen-line crossings, speeding, wrong-way movements, and target-detected events. This flexible configuration interface enables seamless forwarding of real-time events to external systems—such as the TPIMS API—supporting validation, analytics, and live occupancy computation.

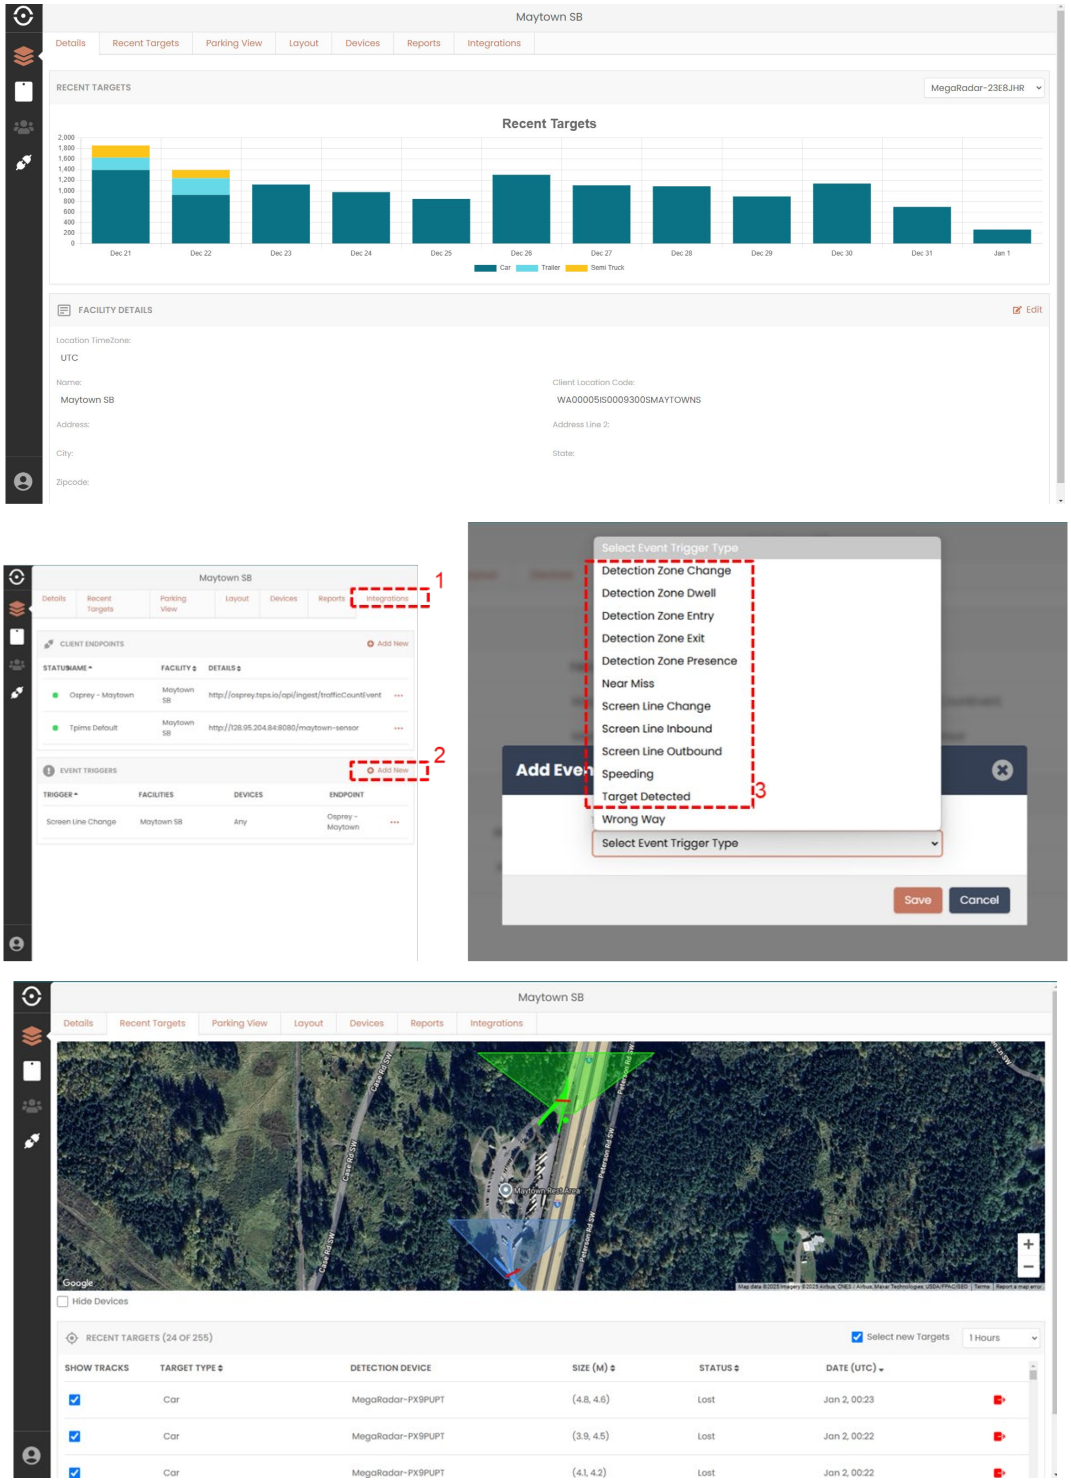

Figure 5-3 Omnisight dashboard

In the early phase of the project, raw sensor data from Omnisight and Sensys were intended to flow directly to the UW server for prediction model development. However, the operational constraints limited direct access to raw data across the WSDOT firewall. As the system matured, the project transitioned to using the processed and standardized parking data supplied by the TSIS TPIMS API. This approach leverages the vendor's established data pipeline—where Omnisight and Sensys detections are already fused, quality-checked, and aggregated into a unified format—reducing complexity on the UW side and ensuring consistency with the data consumed by WSDOT's live applications, including third-party apps, and the TSIS web platform. In the recommended architecture (Figure 5-4), the TSPS API becomes the authoritative data source for both real-time parking status and predictive modeling inputs, enabling UW to develop and serve prediction outputs that seamlessly integrate with WSDOT's operational TPIMS ecosystem.

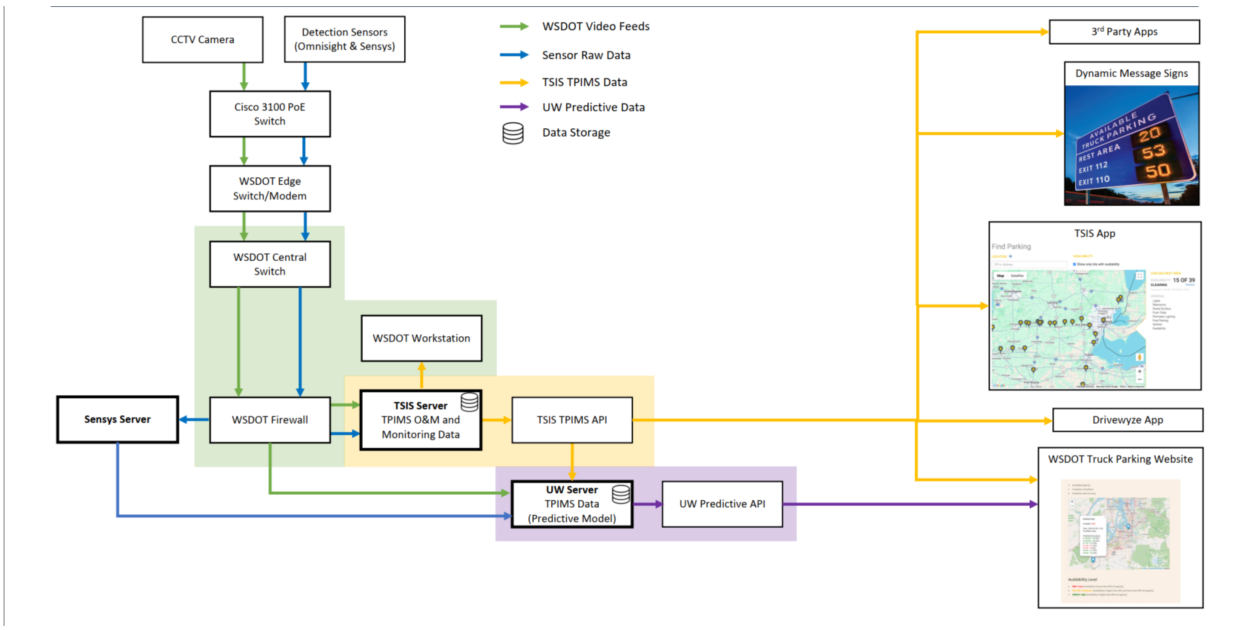

Figure 5-4 TSIS recommended architecture

Streetline Camera Detection Sensor

The Streetline camera-based detection system generates real-time parking availability information by processing video feeds through an embedded computer vision model deployed on the device. Rather than streaming raw video, the sensor performs on-device analytics to identify occupied and unoccupied spaces and transmits only the processed results to the cloud. The resulting JSON feed includes standardized fields such as facility identifiers, total spaces, number of available spaces, number of defined parking areas, and timestamped occupancy updates. An example output from this API is shown on the left, where each facility reports the current availability and metadata needed for downstream applications.

This processed feed is integrated directly into end-user platforms such as the ParkerTruck mobile application (Figure 5-5), which displays available spaces, restroom or ADA-specific stalls, historic occupancy charts, and live camera snapshots for driver verification. The app aggregates Streetline data across multiple states and provides a location selector, enabling truck drivers to navigate to facilities with confirmed availability. The use of processed JSON rather than raw video reduces bandwidth, minimizes privacy concerns, and ensures fast refresh rates suitable for real-time parking guidance and long-distance trip planning.

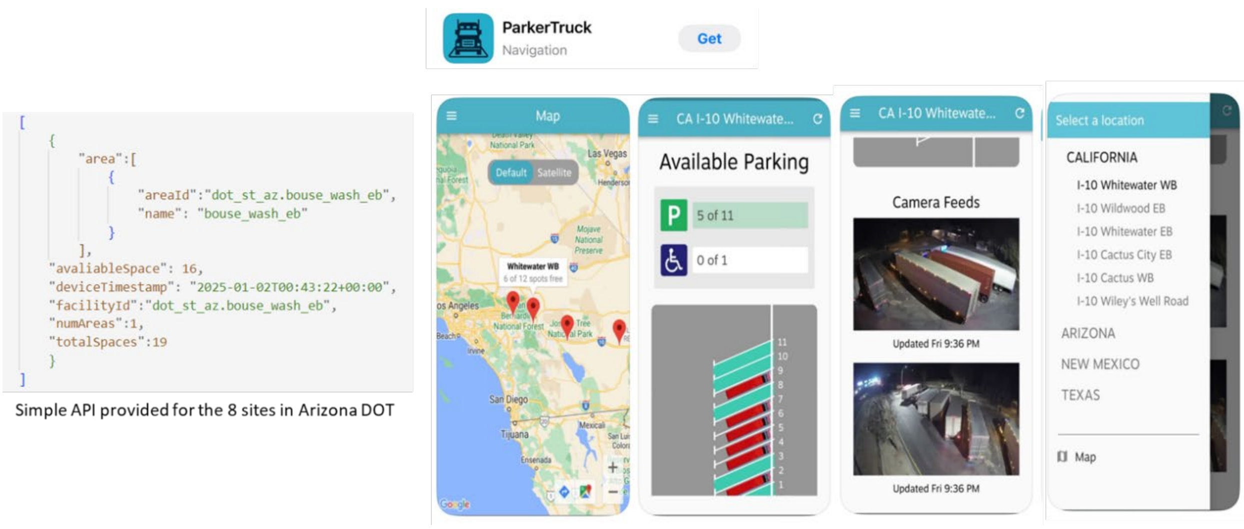

Figure 5-5 Streetline camera-based detection system and example mobile application integration.

During the development period, the ParkerTruck application supported only a limited number of sites, primarily in states already participating in Streetline's deployment. However, there has been ongoing collaboration to integrate the application with the UW-hosted TPIMS site and to incorporate the prediction component. This continued effort aims to enable ParkerTruck to display both real-time availability and multi-horizon occupancy forecasts for Washington sites once the data pipeline and system interfaces are fully aligned.

5.1.2 Server Migration

During the development and testing phases of the TPIMS system, the University of Washington (UW) server served as the primary environment for API hosting, data processing, and prediction model execution. This setup enabled rapid iteration, close monitoring of model behavior, and flexibility in refining the data pipeline while the detection systems were still being validated. From the outset, however, the system was designed with the expectation that all operational components—including the API, web application, and predictive engine—would ultimately migrate to WSDOT's production environment. As the data pipeline stabilized, the project team initiated a phased transition to move data ingestion, storage, API delivery, and prediction processing from UW to WSDOT servers.

The migration required coordination on several technical elements, including securing inbound data access from TSIS Osprey and Sensys, configuring a tunnel for outbound API and web traffic, and enabling HTTPS encryption for all data in transit. Under the planned architecture, WSDOT will continue to pull processed parking data using the existing vendor tokens, while UW's prediction algorithm will be deployed directly within WSDOT's environment to support fully integrated, real-time forecasting (Figure 5-6). The transition also includes shifting the public-facing truck parking webpage from the UW environment to WSDOT's hosting platform.

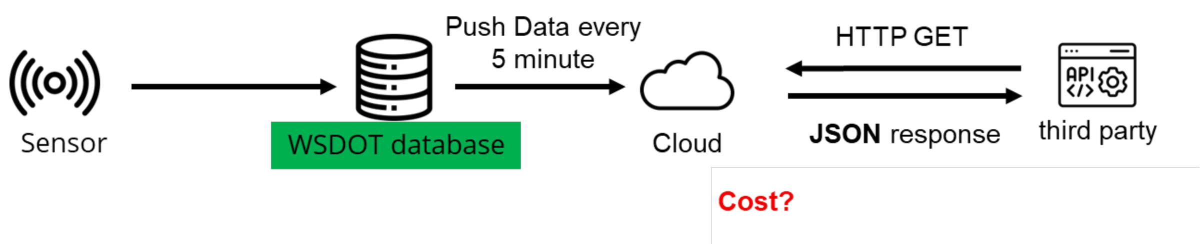

Figure 5-6 Hosting the TPIMS on WSDOT server and using a push-based architecture for data feed

5.2 SENSOR INSTALLATION AND FIELD DEPLOYMENT

5.2.1 Installation Plan and Procedures

Sensor installation and field deployment followed a structured, validation-first approach to ensure reliable system performance. In addition to deploying the primary detection sensors at each site, supplemental camera infrastructure was incorporated as a required component of the installation plan. While detection sensors would provide the operational data stream, camera imagery would serve as an independent ground-truth reference for validating detection accuracy and diagnosing errors. Figure 5-7 shows a site-specific timeline for TPIMS installation and validation.

	Site	Camera (Dec 2025)			System	Jan	Feb-Apr	May				June				July				Aug				Sept				Oct				Sep			
		CCTV Feed Avail (# cams)	*FTP Site Avil (Y/N)*	*Snapshot Avail (# cams)*				Wk 1	Wk 2	Wk 3	Wk 4	Wk 1	Wk 2	Wk 3	Wk 4	Wk 1	Wk 2	Wk 3	Wk 4	Wk 1	Wk 2	Wk 3	Wk 4	Wk 1	Wk 2	Wk 3	Wk 4	Wk 1	Wk 2	Wk 3	Wk 4	Wk 1	Wk 2	Wk 3	Wk 4
1	Ft. Lewis WS NB	Streetline Camera			Detection																														
					Prediction																														
2	Maytown RA	2	Yes	2	Detection																														
					Prediction																														
3	Custer RA NB	1	No	1	Detection																														
					Prediction																														
4	Custer RA SB	1	Yes	1	Detection																														
					Prediction																														
5	Smokey Point RA NB	2	Yes	2	Detection																														
					Prediction																														
6	Smokey Point RA SB	2	Yes	3	Detection																														
					Prediction																														
7	Gee Creek RA NB	2	No	2	Detection																														
					Prediction																														
8	Gee Creek RA SB	3	Yes	2	Detection																														
					Prediction																														
9	SeaTac WS	4	Yes	4	Detection																														
					Prediction																														
10	Toutle River RA NB	2	No	2	Detection																														
					Prediction																														
11	Ridgefield WS	4	No	4	Detection																														
					Prediction																														

Figure 5-7 Site-specific timeline for TPIMS detection and prediction installation and validation based on camera availability

Following installation and system bring-up, each site underwent a phased stabilization process. A one-month validation period was allocated to calibrate detection algorithms, verify consistency across operating conditions, and achieve the target detection accuracy. Once detection performance had been validated, an additional one-month period was allocated for prediction model tuning to ensure that forecast accuracy was achieved using reliable detection inputs. This staged timeline was intended to minimizes error propagation and improve long-term system reliability.

System validation relied on two complementary camera data sources: live closed-circuit television (CCTV) video streams accessed through WSDOT infrastructure for real-time verification, and periodic CCTV image snapshots transferred via WSDOT-managed File Transfer Protocol (FTP) sites for retrospective analysis. Together, these sources provided both timely validation and the ability to investigate historical conditions, supporting robust and defensible deployment.

5.2.2 Site-Level Installation Summaries

Site-level installation planning was conducted for each deployment location to identify optimal placements for sensors, cabinets, communication links, and supplemental cameras. Site plans documented existing infrastructure, including light poles, power availability, cabinets, and wireless connectivity, to support efficient and cost-effective deployment.

For in/out detection sensors, priority was given to poles located at facility entrances and exits. These locations supported accurate vehicle counting while leveraging existing electrical infrastructure, thereby reducing the need for new pole installations and minimizing installation costs. Sensor placement at these points also simplified maintenance and improved long-term system reliability.

In addition to detection sensors, potential camera installation locations were identified at each site. These plans specified candidate poles, and they anticipated camera fields of view and coverage areas to ensure adequate visibility of parking activity for validation purposes. Camera placement was selected to maximize coverage of parking bays, entrances, and internal circulation while minimizing visual obstruction and installation complexity.

Detailed site plans, including pole locations, equipment configurations, wireless links, and camera coverage diagrams, are provided in the Appendix 5-2 site-level installation summaries for reference and implementation support. The Gee Creek NB plan is shown in Figure 5-8.

Figure 5-8 Gee Creek NB plan for sensor and camera installation

5.2.3 Challenges and Mitigation Strategies

Several site-specific challenges were identified during sensor installation and field deployment. These challenges primarily relate to vehicle classification at shared access points, site geometry constraints, and camera coverage limitations. Mitigation strategies were evaluated through collaboration with technology vendors and project partners.

<u>*Vehicle Classification at Shared Entrance/Exit Points*</u>

At some sites, detection sensors were positioned at access points shared by both general-purpose (GP) vehicles and trucks. This configuration reduced classification efficiency, as the sensor had to distinguish between vehicle types operating within the same spatial and temporal window. Misclassification at these locations could lead to inaccurate in/out counts and downstream prediction errors. Figure 5-9 shows an example at Gee Creek NB.

Figure 5-9 Example of vehicle classification challenges at Gee Creek NB site with separate entrances but a shared exit, where overlapping truck, RV, and general-purpose vehicle movements reduce in/out detection accuracy.

Site Geometry and Design Constraints

At Maytown locations, site layout introduced inherent limitations to accurate detection. Trucks and RVs use separate entrances but share a common exit (Figure 5-10). This configuration resulted in repeated exit detections as vehicles circulated within the facility, leading to systematic undercounting and elevated false negatives. Several mitigation strategies were evaluated. Sensor-side approaches included restricting counts to heavy vehicles only, adjusting virtual screen lines, and recalibrating entry and exit detection logic. While filtering for semi-trucks could improve classification precision, it could also undercount non-semi vehicles that occupied truck-designated spaces. Alternative screen line configurations were also considered; however, limitations in sensor range and viewing geometry restricted reliable detection of vehicles in distant stalls.

Relocating sensors to alternative pole locations was identified as a potential mitigation where site geometry permitted. In particular, relocating sensors to positions that would isolate truck-only circulation paths might improve classification performance. In parallel, recalibrating detection zones to accommodate overlapping truck and RV usage was explored as a pragmatic compromise at constrained sites.

Physical site modifications were also considered. Measures such as restriping or installing low physical barriers could better separate GP, RV, and truck movements, thereby improving sensor performance by reducing ambiguous vehicle paths. These interventions, while outside the sensor system itself, might offer long-term benefits for both operations and data quality.

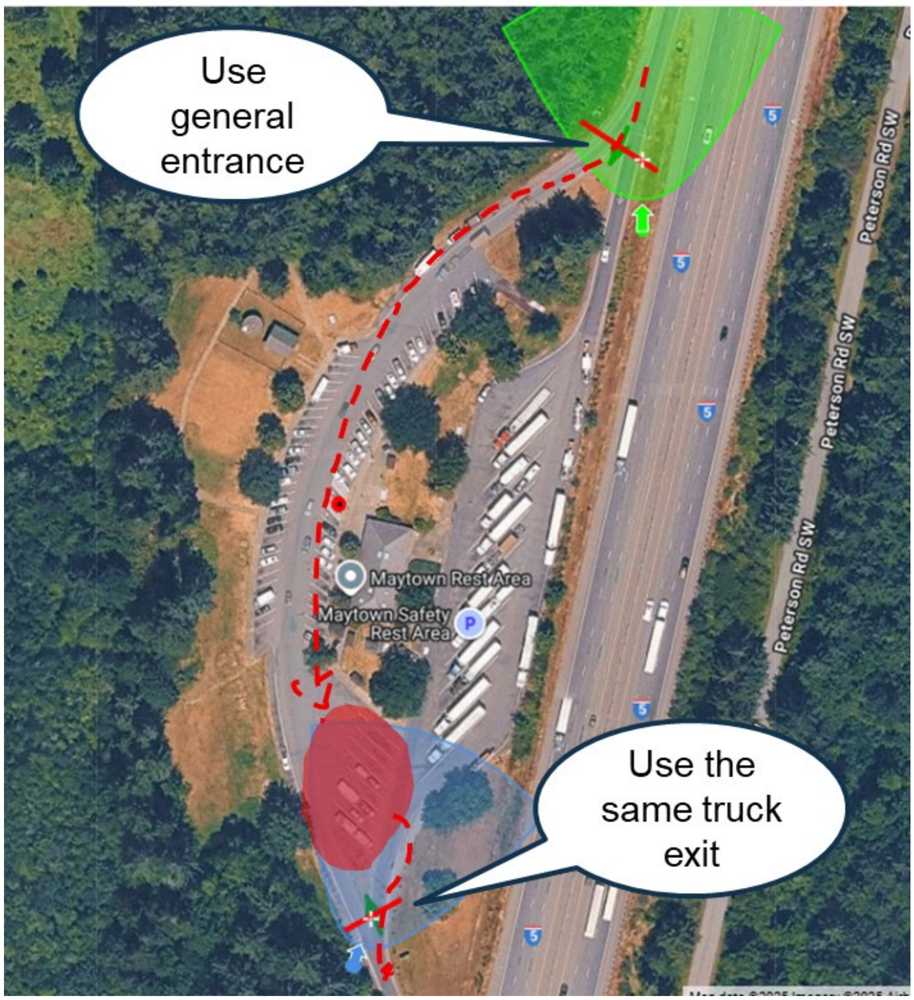

Figure 5-10 Site layout showing shared exit geometry and vehicle circulation patterns affecting in/out sensor performance at the Maytown Safety Rest Area.

Camera Coverage and Environmental Conditions

Camera-based validation also presented several challenges. In some locations, camera coverage was limited because of obstructions such as vegetation, insufficient overlap between camera views, or inadequate nighttime illumination (Figure 5-11). Severe weather conditions could further degrade image quality and reduce validation effectiveness. Mitigation strategies included identifying alternative mounting locations, prioritizing overlapping fields of view for critical areas, and adjusting validation procedures to account for environmental variability.

Figure 5-11 Severe environmental conditions at Gee Creek NB (left) and camera view obstruction caused by vegetation at Toutle River (right).

Camera coverage was systematically evaluated at each site to assess adequacy for validation purposes. This evaluation identified locations where existing camera placement would provide sufficient coverage without modification, locations requiring mounting or orientation adjustments, and locations where additional cameras were necessary to achieve full coverage of critical areas. Results of the camera coverage assessment, including recommended adjustments and supplemental installations, are documented in Appendix 5-2 Camera Coverage.

5.3 DATA MANAGEMENT SYSTEM DEVELOPMENT

5.3.1 Data Schema

The TPIMS database schema is organized around four core tables—StaticPublic, DynamicPublic, DynamicArchive, and Prediction—that together support multi-source truck parking monitoring and forecasting. StaticPublic stores site attributes keyed by siteId, including location (as a text array), ownership, capacity, and high-level site metadata such as relevant highway, reference post, and direction of travel. This table provides the reference layer that other tables join to when reporting or visualizing results at the site level. DynamicPublic captures the most recent operational state for each site with an auto-increment id key, including timestamps, reported available spaces, trend, open/trust status, and short- to mid-horizon prediction outputs (pre_10m–pre_8h), optionally enriched with external context such as weather and WSDOT traffic conditions. Historical snapshots of these operational states are persisted in DynamicArchive, which mirrors the core occupancy fields but adds verification and calibration support variables (e.g., lastVerificationCheck, verificationCheckAmplitude, lowThreshold, and trueAvailable) to facilitate anomaly detection and long-term performance tracking of sensors. Finally, the Prediction table stores detailed model evaluation records for each prediction run, again keyed by an auto-increment id and linked logically by siteId and timeStamp. For each horizon, it records the model output (pre_*), the past available (get_*), and a pre-computed accuracy metric (acc_*), along with contextual fields (weather, traffic_wsdot, trend, open). Figure 5-12 shows the Entity–Relationship (ER) schema for the Multi-Source TPIMS database

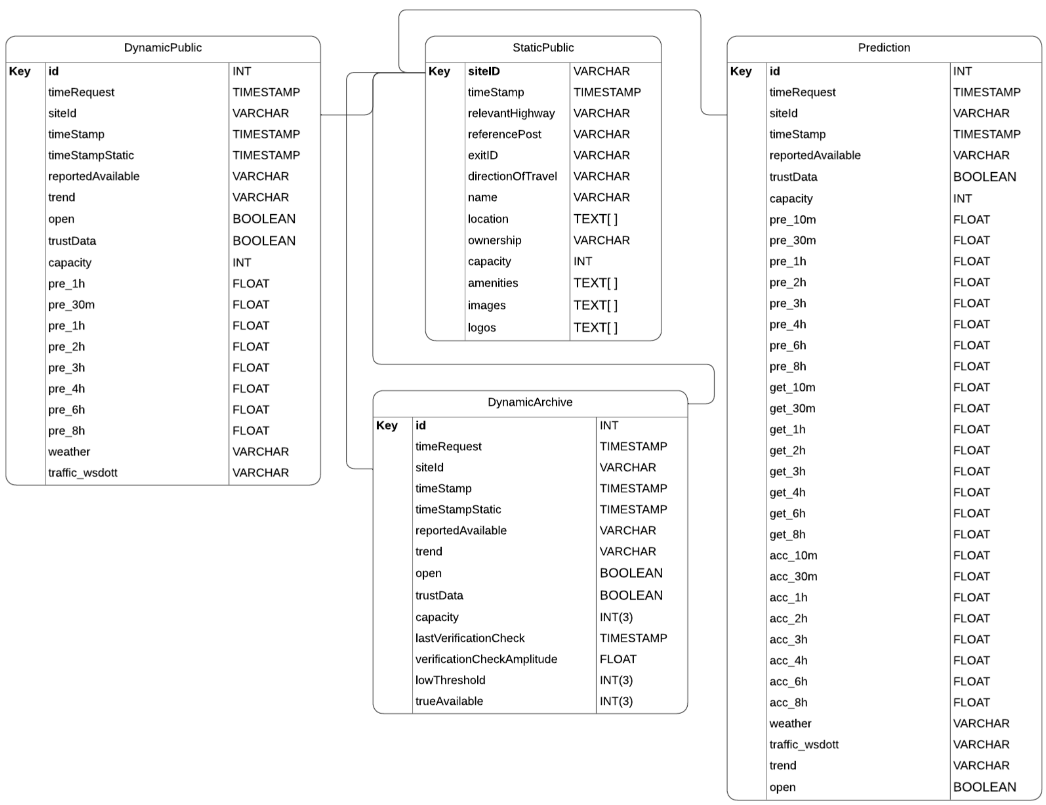

Figure 5-12 Entity–Relationship (ER) schema for the Multi-Source TPIMS database.

5.3.2 Data Pipeline

The TPIMS data pipeline was designed to automatically retrieve and process multi-source data at a fixed 5-minute interval. During each cycle, the system collects updated occupancy, weather, traffic, and event data and stores them in the centralized Raw Data repository. This repository is built on an advanced relational database system that supports both traditional SQL queries and non-relational JSON operations, providing flexibility for handling diverse data structures. The database is free, open-source, and widely used as a primary backend for modern web, mobile, and analytics applications, making it well-suited to support the scalability and reliability requirements of TPIMS. The processed data streams are then used as inputs to generate multi-horizon parking predictions. Prediction outputs are written to the Prediction table together with the corresponding ground-truth occupancy values. Using these stored records, the system continuously computes prediction accuracy for each forecast horizon. To ensure long-term model reliability, the pipeline incorporates a monitoring protocol: if the accuracy for any horizon falls below a predefined threshold, then the system flags performance degradation and suggests a model retraining process using the most recent historical data. This continuous feedback-and-retraining mechanism allows the prediction model to remain adaptive to seasonal patterns, site-specific behavioral changes, and evolving operational conditions (Figure 5-13).

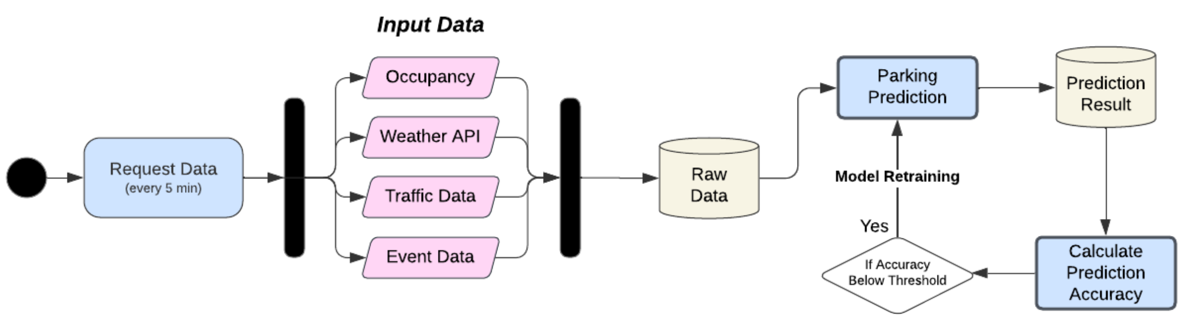

Figure 5-13 End-to-end TPIMS data management and prediction workflow for in/out data.

5.3.3 Quality Control

<u>*Validation*</u>

To ensure the reliability of truck parking detection systems, the project team implemented a structured validation protocol. Validation is done by two sides: vendors and the UW team.

Sensor

From the in/out vendor side, the video footage from on-site cameras is streamed and archived in the vendor's cloud storage. After a 24-hour delay, calibration staff select appropriate time windows for manual verification. During these windows, the client or vendor are able to review the recorded video to manually count vehicle entries and exits and compare them against the sensor-reported values.

Automatic Validation Framework

In addition to the vendor-provided validation workflow, the project team investigated for automatic validation process. Raw video footage could be used directly as a comprehensive source of ground truth for system validation. Video provides several high-value attributes—including vehicle in/out counts, vehicle type classification, slot-level occupancy, parking duration, and full vehicle trajectories—that supplement and strengthen the validation of both radar-based and slot-based detection systems.

As an additional validation alternative, the team also explored the potential use of edge computing to reduce reliance on cloud-based workflows and high-bandwidth communication infrastructure. By processing video or sensor data directly at the device or roadside unit, edge computing can significantly minimize the need for high-speed or broadband data transfer, thereby lowering supporting infrastructure and operational costs. Edge-based systems can perform on-site truck detection, video analytics, and calibration tasks, transmitting only processed summaries or validation metrics rather than full video files. This approach offers a more efficient and cost-effective pathway for long-term system calibration and performance monitoring, especially in rural or bandwidth-constrained environments.

Manual Validation

As part of the validation process, the project team conducted both initial and ongoing accuracy checks to ensure reliable system performance. During the initial setup phase, counting accuracy was verified before each site was formally accepted into operation, supported by daily verification to identify and correct early-stage configuration errors. Over time, the team observed that errors tended to accumulate, particularly for in/out–based systems where small miscounts could compound and result in larger discrepancies. To monitor such drift, the team performed periodic validation using a fixed camera within the sensor's field of view, allowing direct comparison between manual observations and sensor-reported counts. These ongoing checks were typically performed on a weekly or bi-weekly basis, ensuring that sensor performance remained stable and that any degradation or misalignment could be detected and corrected promptly.

Source of Data: SFTP WSDOT

For routine validation, the project team relied primarily on CCTV footage, which provided continuous visual confirmation of vehicle movements and parking activity. WSDOT supported this process by providing a Secure Shell File Transfer Protocol (SFTP) site for the vendor, enabling secure transfer of recorded video (Figure 5-14). The vendor's system captured one-minute video segments from each camera and uploaded them to the SFTP location, where files were stored temporarily and automatically deleted every 24 hours to manage storage capacity. Across sites, the system maintained approximately eight cameras, generating roughly 75 GB of video data over a six-month cycle. This CCTV-based workflow served as the main ground-truth source for validating detection accuracy, monitoring system drift, and confirming correct in/out counts and occupancy behavior.

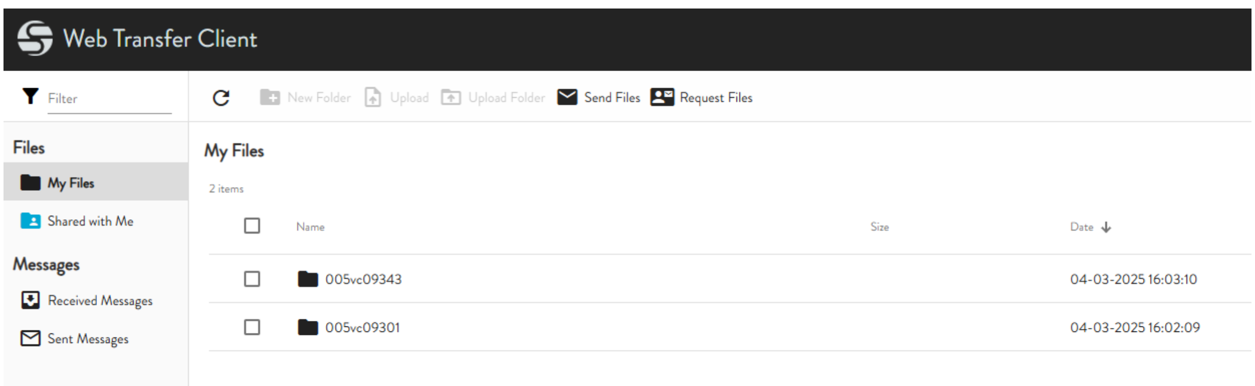

Figure 5-14 Access to WSDOT SFTP (Mar 03, 2025)

Source of Data: WSDOT CCTV

In addition to the SFTP-based video transfer workflow, WSDOT further supported the validation process by installing live CCTV access directly at the UW STAR Lab (Figure 5-15). As part of this setup, WSDOT replaced the network switch at the field cabinet and completed the necessary network configuration, enabling UW to securely view real-time camera feeds from the study site. This live CCTV interface allowed researchers to monitor system behavior, verify sensor performance in real time, and quickly diagnose issues such as misalignment, obstructions, or environmental effects. The availability of live monitoring significantly enhanced the team's ability to perform timely validation and troubleshooting without requiring on-site visits.

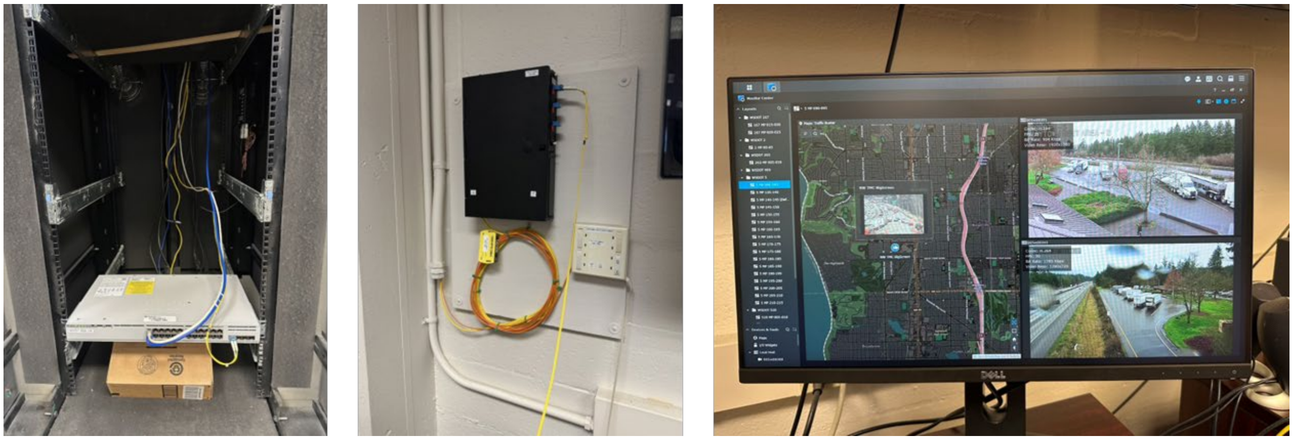

Figure 5-15 CCTV access directly at the UW STAR Lab (Mar 20, 2025).

Anomaly Detection

The current system architecture does not provide automated malfunction notifications, so anomaly detection is performed manually. Issues are identified by monitoring incoming data for abnormalities such as sudden drops in reported availability, unexpected patterns, or complete stoppage of data flow (indicating a potential outage). In future iterations, device-level health metrics could be incorporated to support early detection of sensor degradation. These metrics may include power status, device orientation, RAM usage, radar temperature, ambient temperature, and on-device processing status. Leveraging these operational attributes would enable more proactive maintenance and improve long-term system reliability.

In addition to detecting sensor-related anomalies, the system also experienced server-level anomalies, specifically recurring server freezes that affected data collection and processing. Over several weeks, the UW server went down multiple times, primarily due to the combination of an aging machine and a heavy request load from continuous 5-minute data ingestion and processing tasks. The team implemented several fixes and configuration adjustments and began closely monitoring system performance to isolate triggers. Notably, previous freezes tended to occur during weekends, prompting targeted observation during those periods. These efforts ensured system stability and minimized disruptions to the end-to-end data pipeline.

To maintain uninterrupted service during periods of server instability, the team implemented a temporary dual-endpoint solution. The original UW server continued to handle core responsibilities—including data storage and all API endpoints such as static public, dynamic public, dynamic archive, and prediction—along with supporting the UW website interface. In parallel, a temporary virtual machine (VM) hosted in the cloud was deployed to ensure that Drivewyze could continue receiving real-time parking information without interruption. This VM hosted a simplified version of the API, providing static and dynamic public data only. By splitting responsibilities across these two endpoints, the system remained operational for all stakeholders while long-term server fixes and performance adjustments were under way.

5.4 PREDICTION ALGORITHM DEVELOPMENT AND DEPLOYMENT

5.4.1 Long Short-Term Memory (LSTM) Prediction Architecture

The truck-parking prediction model is built using a deep sequence learning approach, specifically a long short-term memory (LSTM) neural network. LSTM is a type of recurrent neural network (RNN) designed to learn temporal dependencies by retaining information from previous time steps through its internal memory cell and gating mechanisms (Figure 5-16). This makes LSTM well-suited for time-series problems where past observations influence future outcomes.

In our framework, the model is trained on sequences of labeled historical data—each sequence representing past parking availability (and additional attributes such as traffic conditions where applicable). Once trained, the model can generate predictions for new, unseen sequences, enabling short-term and multi-hour forecasts of parking availability. This is conceptually similar to other sequence-classification tasks, such as determining whether a movie review is positive or negative based on the sequence of words; however, in our application, the input sequence is a continuous stream of occupancy and contextual features.

The deployment pipeline utilizes the trained LSTM weights to run real-time inference. Observed parking data from the last required time window are fed into the model, which outputs predicted availability for multiple future horizons. These predictions are then provided to the TPIMS system through our FastAPI-based service for operational use and dissemination to truck drivers.

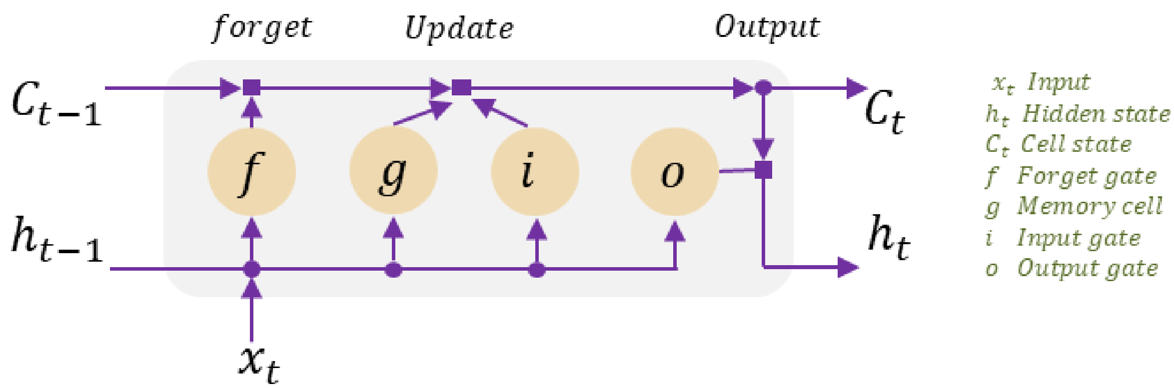

Figure 5-16 Structure of an LSTM unit, illustrating the input, forget, and output gates that regulate information flow through the cell state.

<u>*Traffic Data Integration*</u>

STAR Lab Traffic Data

In addition to historical parking data, we incorporate traffic data as an additional model attribute. One of our alternatives is to incorporate corridor-level traffic data from WSDOT loop detector data that are archived and processed by the STAR Lab data pipeline, covering more than 900 detector cabinets across Washington state. As shown in Figure 5-17, the traffic data available from loop detector cabinets do not fully correspond to all truck parking locations statewide. To improve prediction quality and support broader system scalability, future efforts will investigate additional data sources to expand coverage across all parking facilities.

Figure 5-17 Traffic data at the STAR Lab

WSDOT Traveler Information API

As an additional traffic data source, the project explored the WSDOT Traveler Information API, which provides a unified gateway to multiple statewide traffic and roadway information feeds. Among the available interfaces (e.g., traffic flow, travel times, weather, highway alerts, cameras), the Traffic Flow API was identified as the most relevant for truck-parking prediction. This interface provides congestion level indicators for highway segments using a categorical scale (0–5), where 1 = WideOpen, 2 = Moderate, 3 = Heavy, 4 = Stop-and-Go, and 5 = No Data. For each record, the API also supplies useful metadata including station location (latitude, longitude, milepost, road name), direction of travel, region, and timestamp. In our prototype analysis, the fields utilized were FlowReadingValue, latitude/longitude, milepost, and timestamp, allowing us to spatially link traffic conditions with nearby truck parking facilities.

However, the coverage of the Traveler Information API is not geographically comprehensive. Many of the available flow stations are concentrated in the Puget Sound urban corridor, with sparse or no coverage along several major freight routes and rural segments where truck parking lots are located. As illustrated in the Figure 5-18 parking overlay map, a number of WSDOT flow stations do not fall within effective proximity to the targeted truck parking sites. Consequently, while the API could serve as a valuable supplementary data source, its incomplete statewide distribution would limit its ability to support consistent traffic–parking integration for all sites. Future work will consider additional data sources or hybrid integration approaches to improve coverage across rural and eastern Washington corridors.

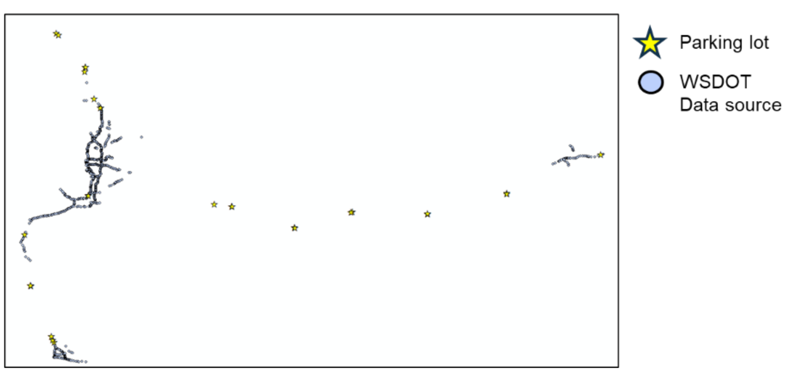

Figure 5-18 WSDOT Traveler Information API location and truck parking site

WSDOT Traffic Count (TCDS)

In addition, WSDOT also maintains a Traffic Count (TCDS) website that provides a more comprehensive statewide dataset, including 15-minute traffic counts and hourly averaged volumes (Figure 5-19). Although this source does not offer real-time data, its broad geographic coverage and consistent historical records make it highly suitable for model development. These aggregated traffic measurements could be incorporated as auxiliary features during the training or testing phases of the prediction algorithm, enabling the model to learn long-term traffic–parking relationships even at sites where real-time flow feeds are unavailable.

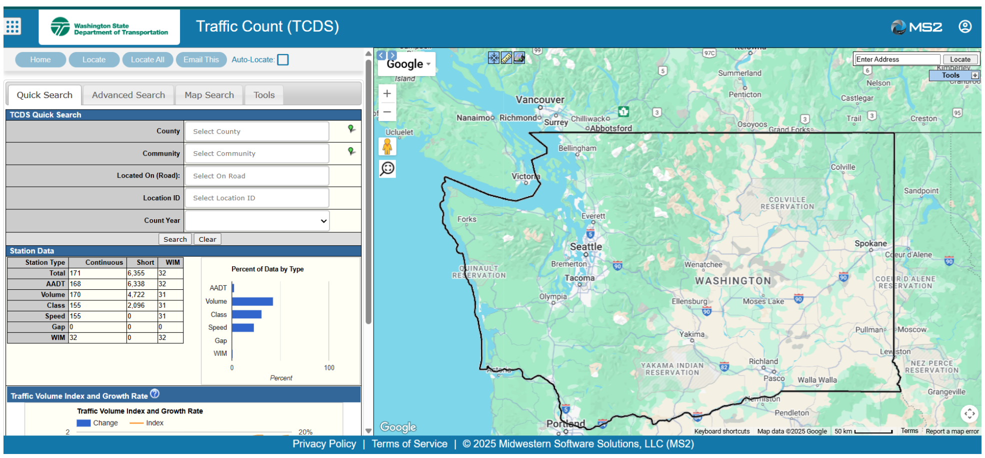

Figure 5-19 WSDOT Traffic Count (TCDS).

5.5 SYSTEM EVALUATION AND PERFORMANCE ASSESSMENT

5.5.1 Detection System Evaluation

<u>Challenges</u>

Duplication

One of the first issues observed during system evaluation was the occurrence of duplicated detections. This typically happened when consecutive screenshots were taken only a few seconds apart, causing the system to register the same truck more than once—especially when another vehicle was positioned closely behind the first truck. Such proximity made it difficult for the detection algorithm to distinguish whether the sensor was capturing a new vehicle or repeatedly detecting the same one. To address this, we implemented validation against synchronized video footage to confirm true vehicle counts and identify false duplicates. In addition, further analysis was conducted on closely timed detections, particularly cases where the same device reported detections within less than 5 seconds, allowing the team to refine filtering rules and reduce duplication errors.

Incorrect Classification

Another issue identified during recent evaluations was incorrect classification of detected objects. This problem had not been observed in previous rounds of testing and only emerged during the most recent evaluation period. In this case, the system occasionally misidentified non-truck objects—or trucks in unusual positions—as occupied or unoccupied parking events, leading to inconsistencies in the reported availability. Because this represented a newly emerging issue, additional verification was performed using video evidence to confirm the true object type. Ongoing investigation will focus on identifying the root cause, assessing whether recent environmental or system changes contributed to the misclassification, and updating detection rules accordingly.

Negative Availability Reporting

A newly identified issue involved the API occasionally returning negative values for trueAvailable in the dynamic archive. This typically stemmed from counter overflow or accumulated miscounts that pushed the internal availability estimate below zero. Such values were not compatible with the MAASTO standard, which requires non-negative, clearly interpretable availability fields. Although these occurrences were infrequent, they could affect downstream systems relying on consistent data formats. This issue has been flagged for correction, and additional safeguards—such as lower-bound clipping and enhanced validation checks—will be implemented to prevent negative availability values from entering the public API feed.

Design Leak (General Entrance → Truck Exit Misrouting)

During validation, we identified a site-specific design leak in which some vehicles entered the rest area through the general entrance but later exited through the truck-only exit (Figure 5-20). This movement pattern bypassed the intended detection flow and caused the system to register an exit without a corresponding entrance event. As a result, the availability counter decreased incorrectly, leading to cumulative miscounts over time. The issue was confirmed using the validation camera, which captured smaller vehicles triggering the truck exit screenline. Potential mitigation strategies include adjusting the exit screenline placement, refining classification rules to differentiate trucks from non-trucks at the exit, and applying logic that reconciles unmatched exit events during post-processing. This challenge highlighted the importance of aligning physical site geometry with the detection logic to prevent structural sources of counting error.

Figure 5-20 Design leak at Maytown.

To mitigate this issue, several potential adjustments were considered. Sensor-side strategies included implementing vehicle-type filtering to count only semis, although this would risk underrepresenting actual occupancy when non-semis legitimately used truck stalls. Other options evaluated included refining or adding screenlines to distinguish general-purpose (GP) movements, relocating the sensor farther north to isolate the true semi exit, and recalibrating both entry and exit screenlines to better separate traffic streams. Physical site modifications were also discussed, such as restriping or installing low-height barriers to more clearly delineate GP and semi zones, reducing the likelihood that RVs or passenger vehicles would cross the truck exit screenline. Addressing this design leak will require a combination of detection logic adjustments and potential site-level improvements to align actual vehicle paths with the intended sensing geometry.

Validation Approaches

The evaluation of the detection system followed a step-by-step approach designed to confirm sensor functionality before assessing overall accuracy. Three types of validation were conducted.

1. **Sensor Functionality Validation (Slot-by-Slot Devices)**
   For slot-by-slot sensors, we first verified that each device was able to detect *any* vehicle correctly. API-reported availability was compared directly with the corresponding video feed from the device. This step ensured that the sensor hardware and data transmission were functioning as expected before deeper accuracy analysis.
2. **Vehicle Classification Validation (In/Out Sensors Only)**
   Some locations shared the same entrance and exit for general parking and designated truck-parking areas. For these sites, classification accuracy was critical to distinguish vehicles that actually used the truck parking lot. We grouped vehicle types into two categories based on operational use:
   - **Truck group:** Bobtail, Bus, Camper, Semi-truck, Utility vehicle
   - **Car group:** Passenger car, Passenger pickup/light truck
     The sensor's ability to correctly classify vehicles into these categories was validated to ensure accurate occupancy counting for truck-only stalls.
3. **Counting Validation (Both In/Out Sensors and Slot-by-Slot Sensors)**
   Once the sensor functionality and classification accuracy had been confirmed, we evaluated counting performance. During the initial deployment period, ground-truth verification was performed daily. To monitor accuracy over time, we compared CCTV snapshots with API-reported availability at multiple intervals following each ground-truth check: 30 minutes, 1 hour, 2 hours, 4 hours, 6 hours, 8 hours, 12 hours, and 24 hours. This multi-interval evaluation helped assess how quickly accuracy degraded, how reset periods performed, and whether sensors exhibited systematic drift or error patterns.

<u>*Results*</u>

Across all three validation components—sensor functionality, vehicle classification, and counting performance—the evaluation revealed that the detection systems were generally capable but exhibited site-specific issues that affected reliability. Sensor functionality checks using entrance/exit video confirmed that slot-by-slot devices could detect vehicle movements but occasionally produced misclassifications or missed detections, particularly when the site was at or near capacity (e.g., SeaTac WS, Smokey Point NB/SB).

Vehicle classification validation for in/out sensors showed varying error rates across sites, ranging from approximately 1 percent at SeaTac WS to 6 percent at Maytown, with notable boundary cases such as passenger pickup trucks being inconsistently classified at Smokey Point SB. Counting validation results demonstrated substantial variation in accuracy across locations and time periods. High-performing sites such as Gee Creek SB (up to 98.77 percent) and Ridgefield PoE NB (97.14 percent) showed stable behavior, while others such as Maytown and SeaTac WS exhibited error accumulation, drift, or mismatches caused by site layout constraints, classification leakage, or inconsistent nighttime detection.

Streetline systems at Ft. Lewis consistently achieved 100 percent accuracy during multiple test periods, although snapshot synchronization issues occasionally caused visual inconsistencies. Overall, the results indicated that while the detection technologies were functional, ongoing calibration, improved classification models, better syncing, and periodic validation remained necessary to ensure the accuracy and robustness required for statewide deployment.

5.5.2 Prediction Algorithm Evaluation

Evaluation Metrics

To rigorously assess the performance of the truck-parking time-series prediction models, we applied three widely used forecasting metrics: mean absolute percentage error (MAPE), mean squared error (MSE), and root mean squared error (RMSE). Each provided complementary insights into the accuracy, stability, and operational reliability of the model. MAPE captured the relative percentage error, MSE reflected overall stability by heavily penalizing large deviations, and RMSE provided practical interpretability by expressing prediction error in the same unit as the actual number of available spaces.

Because prediction accuracy directly affects the usefulness of the statewide truck parking system, we adopted a performance requirement that each model should achieve at least 80 percent accuracy across all prediction horizons. This would ensure that both short-term and longer-term predictions were reliable enough to support routing decisions, traveler information dissemination, and operational planning.

1. Mean absolute percentage error (MAPE) MAPE quantifies the average magnitude of prediction error as a percentage of the actual value:

$$\text{MAPE} = \frac{1}{n}\sum_{t=1}^{n} \left| \frac{A_t - P_t}{A_t} \right| \times 100 \qquad \text{Eq. 5.1}$$

where:

- A_t= actual available spaces at time t
- P_t= predicted available spaces
- n= number of time points

MAPE is intuitive and scale-independent, making it suitable for comparing performance across sites with different capacities. It directly expresses prediction accuracy as a percentage, which supports the operational goal of providing at least 80 percent accuracy for all prediction horizons. MAPE can become unstable when actual availability values are very small (close to zero). For this reason, we applied safeguards such as minimum denominators or adjusted MAPE windows.

2. Mean squared error (MSE) MSE measures the average squared difference between actual and predicted values:

$$\text{MSE} = \frac{1}{n}\sum_{t=1}^{n}(A_t - P_t)^2 \qquad \text{Eq. 5.2}$$

Because errors are squared, larger mistakes are penalized more heavily. This is valuable in truck-parking prediction, where large deviations (e.g., predicting many more available spots than exist) could mislead drivers and reduce system trust.

3. Root mean squared error (RMSE) RMSE is the square root of MSE:

$$\text{RMSE} = \sqrt{\text{MSE}} \qquad \text{Eq. 5.3}$$

RMSE provides an interpretable metric in the same unit as the prediction target (number of spaces). This would help system operators understand the practical magnitude of forecasting errors—for example, an RMSE of 3 would mean that predictions were off by roughly three spaces on average.

<u>Prediction Result at Demo Site</u>

During the initial phase of model development, sensor installations at the pilot locations were not yet completed; therefore, historical data from a demonstration site were used for training and evaluation. Five months of parking availability data (January 1–June 1, 2022) from a site with a maximum capacity of 60 spaces were analyzed (Figure 5-21). The original dataset contained approximately 100,000 records collected at irregular intervals. To enable sequence-based model training, the data were resampled and interpolated to a uniform 5-minute interval, resulting in a final dataset of roughly 300,000 rows. Examination of the processed data indicated strong temporal regularities, including pronounced daily cycles and clear differences between weekday and weekend parking patterns.

Model performance was evaluated across prediction horizons ranging from 10 minutes to 8 hours. Short-horizon predictions demonstrated the highest accuracy, with an average error of 2.29 percent at the 10-minute horizon. As the prediction horizon increased, model accuracy declined gradually, reaching approximately 12 percent error for horizons beyond two hours (see Figure 5-22). Both the average error trend and the time-series comparisons between predicted and observed availability illustrated that the LSTM-based model captured the overall temporal dynamics of truck parking occupancy, despite increasing uncertainty at longer lead times.

Overall, these results indicated that the model provides reliable short-term forecasts—most relevant for operational decision-making and traveler information—while maintaining acceptable performance for longer-range predictions during this early stage of system deployment.

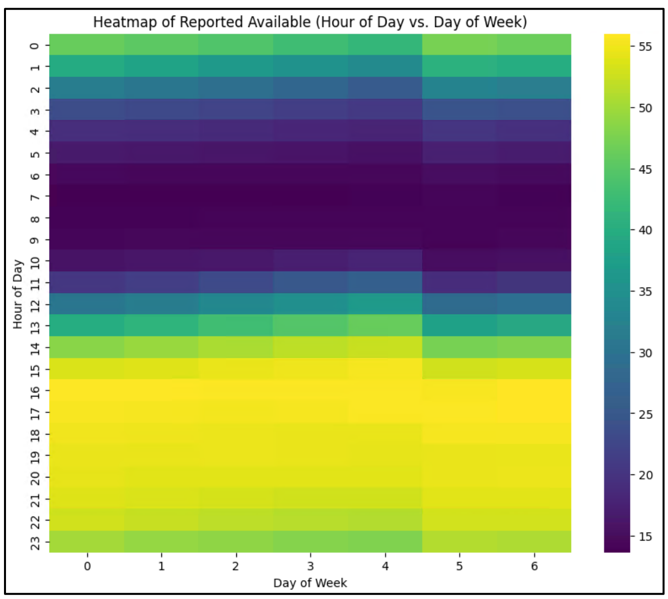

Figure 5-21 Data distribution at the demo site

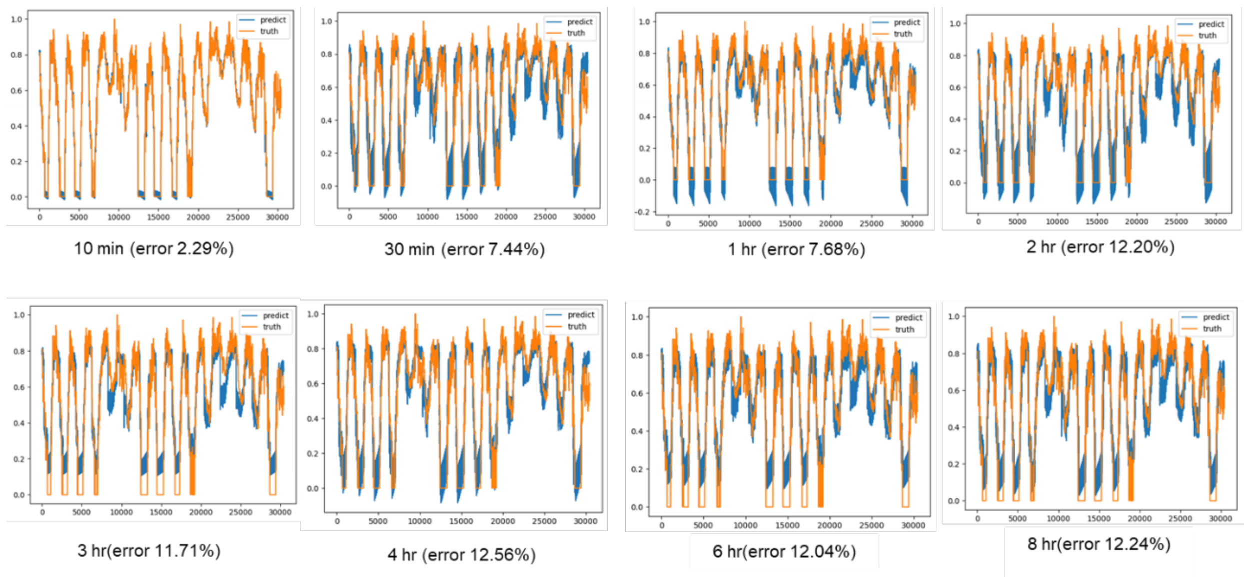

Figure 5-22 Prediction accuracy at the demo site

Prediction Result at Ft. Lewis

After the detection system had been fully deployed and verified, Ft. Lewis became the first site with a complete and reliable historical dataset, enabling a more comprehensive evaluation of the prediction model. Six months of continuous parking availability data (Jan 01, 2025 - July 31, 2025) were used to train and test the model at this location. Like the demo site, the Ft. Lewis data exhibited clear temporal patterns, with parking demand consistently increasing during evening and nighttime hours and decreasing during daytime periods. This daily cycle, along with noticeable weekday–weekend differences, is evident in the heatmap of average reported availability (Figure 5-23).

Model performance was evaluated across multiple prediction horizons ranging from 10 minutes to 8 hours. Results showed that the model performs well for short-term forecasts, with low MAE, MAPE, and RMSE values at the 10-minute and 30-minute horizons. As the horizon increased, prediction error gradually rose—consistent with typical time-series forecasting behavior—yet the model continued to capture the overall occupancy trajectory even at longer horizons (Figure 5-24). Time-series plots comparing actual and predicted availability across all horizons demonstrated that the model reliably followed peak congestion periods at night and lower daytime demand, reinforcing the consistency of temporal patterns observed at both the demo site and Ft. Lewis.

Overall, these findings confirmed that Ft. Lewis provided a stable and representative dataset, and that the LSTM-based prediction model maintained strong short-term performance while offering reasonable accuracy for longer-term forecasts under real-world field conditions.

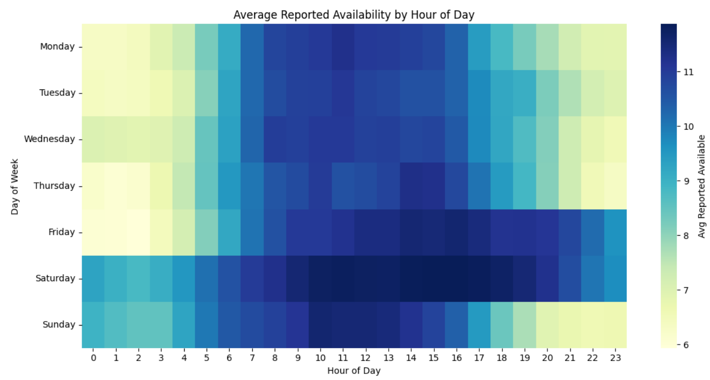

Figure 5-23 Average reported availability at Ft. Lewis (Jan 01, 2025 - July 31, 2025)

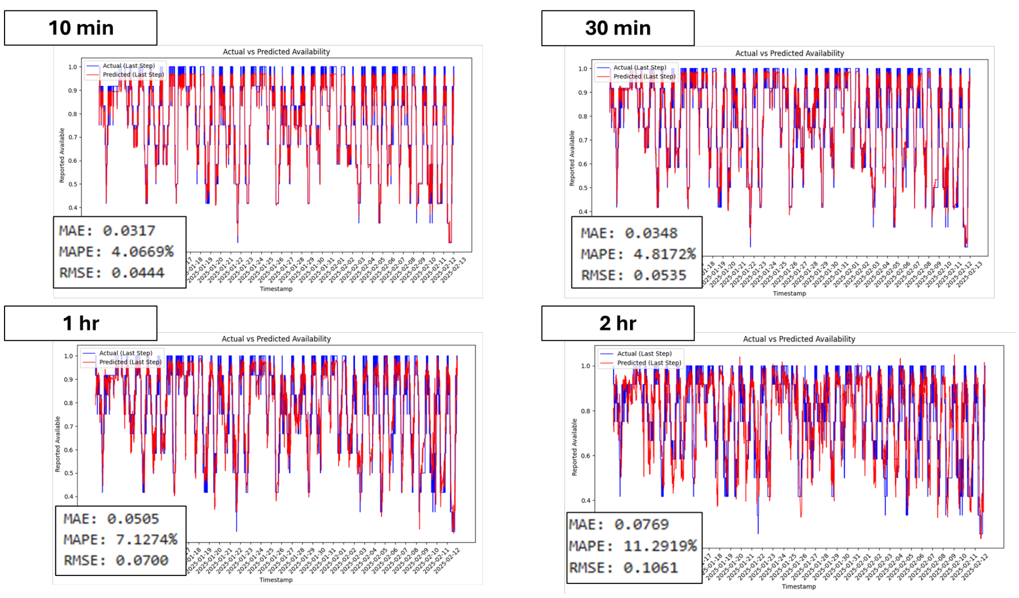

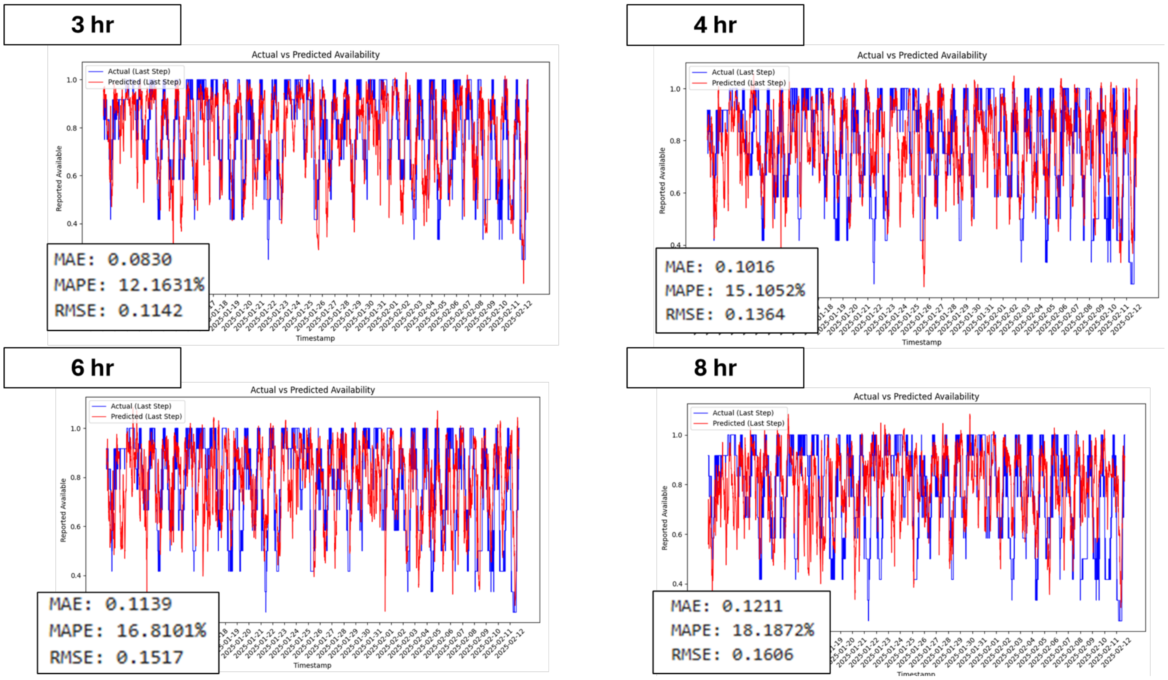

Figure 5-24 Prediction accuracy at Ft. Lewis (last month).

6. API DESIGN AND INTERFACE

6.1 API

6.1.1 MAASTO Data Exchange Specification

The Truck Parking Information Management System (TPIMS) API follows the MAASTO Data Exchange Specification, ensuring interoperability with other states and third-party partners. All data services adopt standard HTTP RESTful conventions, with responses provided in JavaScript Object Notation (JSON). Authentication is stateless and does not rely on cookies or session management; instead, each API request requires an authentication key issued to trusted partners such as the Mid-America Freight Coalition (MAFC) or state DOTs. The API provides three categories of datasets:

1. **Dynamic Public Feed (Dataset #1):** Real-time availability and operational status for all TPIMS truck parking sites. This feed can be accessed openly or through an authenticated endpoint that includes partner-specific metadata.
2. **Static Public Feed (Dataset #2)**: Site characteristics and location descriptors—updated as needed—supporting mapping applications, third-party navigation tools, and integration with commercial routing systems.
3. **Dynamic Archive-Only Feed (Dataset #3)**: Accessible only to trusted partners, this dataset provides the most recent timestamped occupancy entries for internal system monitoring, performance measurement, and validation.

Each entry in the Dynamic Feed includes standardized fields such as

- reportedAvailable – number of available truck parking spaces (or "Low" when thresholds are triggered)
- trend – parking flow state (CLEARING, STEADY, FILLING, or NULL)
- open – whether the site is open or closed
- trustData – data quality indicator
- capacity – total number of spaces at the site.

Parking flow states are derived from short-interval availability cycles (e.g., every 5 minutes). Changes in occupancy—expressed as absolute and percentage deltas—determine whether the site is filling, steady, or clearing (Figure 6-1). These classifications support real-time decision-making and enhance the prediction model's awareness of demand patterns.

Time	Availability	Delta	% Delta	%Flow	Flow State
12:00	20				
12:05	18	-2	-4.0%		
12:10	10	-8	-16.0%		
12:15	9	-1	-2.0%		
12:20	8	-1	-2.0%		
12:25	8	0	0.0%		
12:30	9	1	2.0%	-22.0%	Filling
12:35	7	-2	-4.0%	-22.0%	Filling
12:40	3	-4	-8.0%	-14.0%	Filling
12:45	-1	-4	-8.0%	-20.0%	Filling
12:50	-1	0	0.0%	-18.0%	Filling
12:55	0	1	2.0%	-16.0%	Filling
13:00	0	0	0.0%	-18.0%	Filling
13:05	1	1	2.0%	-12.0%	Filling
13:10	1	0	0.0%	-4.0%	Steady
13:15	1	0	0.0%	4.0%	Steady
13:20	1	0	0.0%	4.0%	Steady
13:25	1	0	0.0%	2.0%	Steady
13:30	1	0	0.0%	2.0%	Steady
13:35	2	1	2.0%	2.0%	Steady
13:40	4	2	4.0%	6.0%	Clearing
13:45	6	2	4.0%	10.0%	Clearing
13:50	7	1	2.0%	12.0%	Clearing
13:55	6	-1	-2.0%	10.0%	Clearing
14:00	8	2	4.0%	14.0%	Clearing
14:05	7	-1	-2.0%	10.0%	Clearing
14:10	6	-1	-2.0%	4.0%	Steady
14:15	7	1	2.0%	2.0%	Steady
14:20	7	0	0.0%	0.0%	Steady
14:25	7	0	0.0%	2.0%	Steady
14:30	9	2	4.0%	2.0%	Steady
14:35	12	3	6.0%	10.0%	Clearing
14:40	16	4	8.0%	20.0%	Clearing
14:45	18	2	4.0%	22.0%	Clearing
14:50	22	4	8.0%	30.0%	Clearing

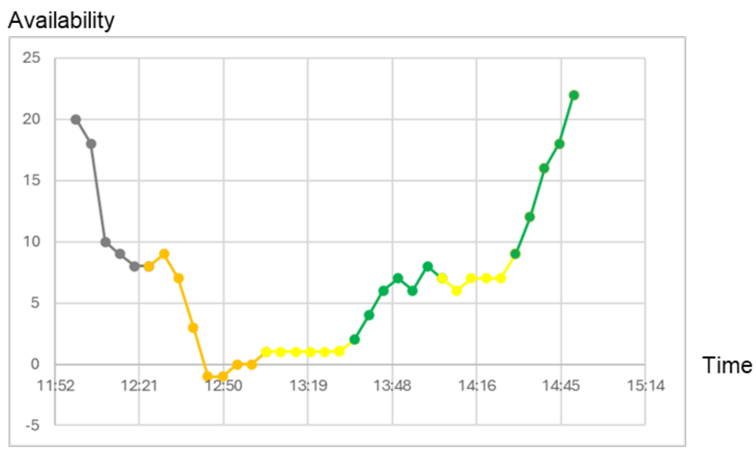

Figure 6-1 Example of availability-based flow-state thresholds, showing how changes in occupancy over 5-minute intervals classify parking conditions as filling, steady, or clearing.

In alignment with MASSTO security practices, the API supports a multi-key access control structure, allowing different users or organizations to receive API keys with defined roles, permissions, and expiration dates. Multiple API keys can be issued and independently managed to support auditing, usage monitoring, and revocation if necessary. Usage limits may also be assigned to ensure equitable access and prevent excessive or unintended system load. Under this structure, WSDOT receives full-access keys enabling retrieval of all API datasets, including prediction and archival feeds, while general public users receive keys restricted to the standard public endpoints only. The UW team is currently working with WSDOT to finalize and implement this key-management framework so that it aligns with operational security requirements and long-term system governance.

6.1.2 API Development

The TPIMS API was developed to provide standardized, programmatic access to real-time truck parking information, prediction outputs, and site metadata for all monitored TPIMS locations. The API follows the MASSTO specification (Minimum Acceptable Standards for Secure Truck Parking Information Exchange) for data structure, update frequency, timestamp format, and security practices. At the same time, WSDOT extends the MASSTO standard by including short- and mid-term parking predictions—becoming the first state to deploy predictions within its TPIMS data service.

The service is implemented using the FastAPI framework, enabling lightweight, high-performance RESTful endpoints, automatic JSON serialization, asynchronous processing, and built-in OpenAPI documentation. FastAPI also allows seamless integration with the deep-learning prediction model, which processes real-time sensor data and produces forecasts every 5 minutes.

In addition to Python-based implementation considerations, system requirements from the WSDOT IT group emphasized compatibility with the .NET framework, which is widely used across state systems for large-scale, enterprise-level applications. The .NET ecosystem provides strong guarantees on performance, security, and scalability, and supports the development of enterprise-grade services in C#. To meet this requirement, the team evaluated interoperability between the FastAPI service and .NET-based systems (Figure 6-2). Testing confirmed that .NET applications can successfully call the FastAPI endpoints over standard HTTP REST protocols without compatibility issues. This validation ensured that the proposed API architecture is fully accessible to existing Microsoft-based environments while still leveraging the strengths of FastAPI for machine-learning integration, resulting in a solution that satisfies both operational IT requirements and research-driven functionality.

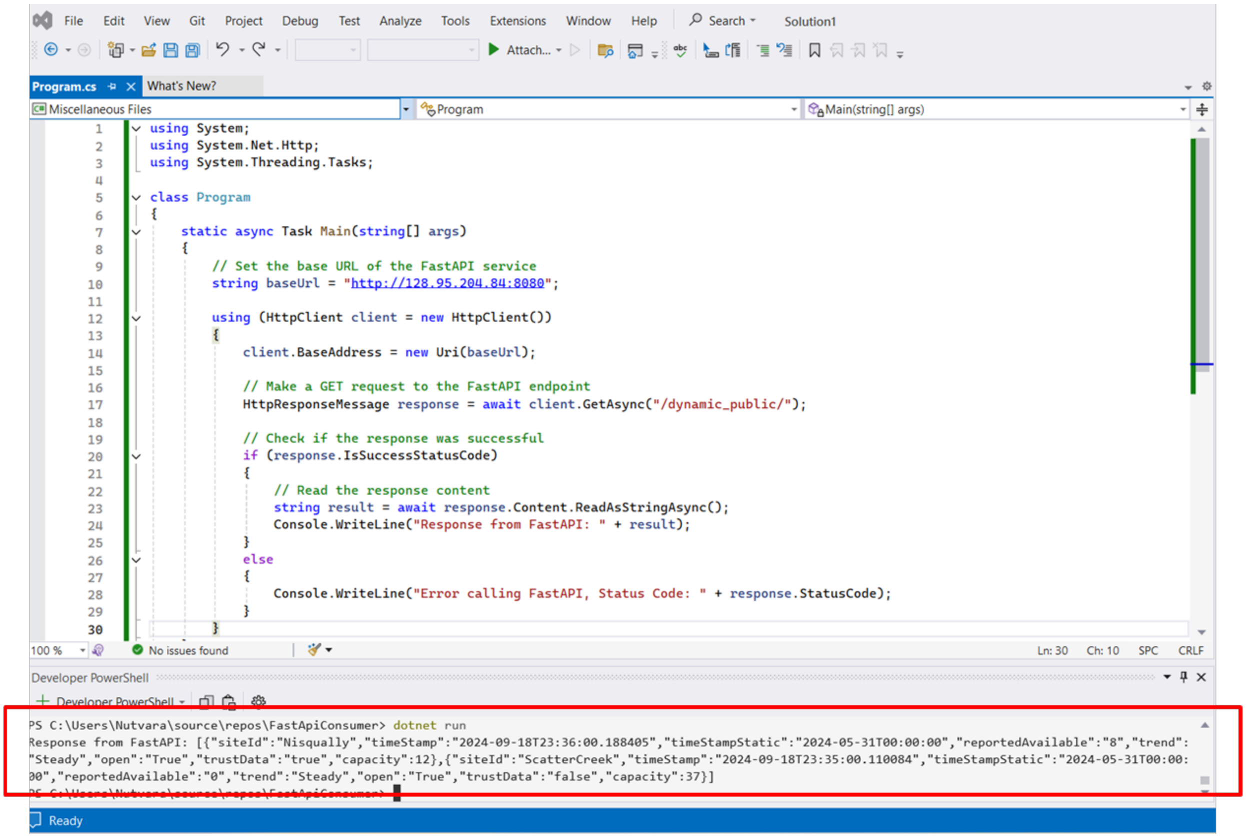

Figure 6-2 .NET framework compatibility testing.

Several TPIMS sites frequently experience overflow parking, where trucks park outside of designated stalls. so the project team established a consistent reporting rule to ensure that API outputs remain meaningful and operationally useful (Figure 6-3). In situations where the number of available marked stalls falls below the predefined low-availability threshold (typically <10 percent of total capacity), the system reports “LOW” rather than a numeric value to more accurately reflect real conditions. As illustrated in the four cases, when two marked stalls remain unoccupied (Case I and III), the system reports the numeric value (2/10). However, when only one stall or none remains (Case II and IV), overflow behavior typically prevents drivers from accessing the last spaces, making the practical availability near zero; therefore, the API reports “LOW/10.” This rule allows the TPIMS feed to better represent true parking usability during congested conditions while remaining consistent with operational expectations and MASSTO-compatible data formatting.

Figure 6-3 Examples of how availability is reported under overflow conditions using the LOW threshold rule.

<u>Static Site Information (static_public)</u>

This endpoint delivers the static metadata necessary for mapping, system integration, and user-facing display applications. It contains information such as location identifiers, site characteristics, ownership, amenities, and capacity—elements that change infrequently and therefore serve as a stable reference layer for all TPIMS operations. These static attributes enable third-party developers to support routing, visualization, and traveler-information services by supplying consistent site-level details. Key fields include the **referencePost**, which provides the milepost location; **directionOfTravel** indicating whether the site serves northbound or southbound traffic; **capacity**, representing the maximum number of truck parking spaces; and **amenities**, such as restrooms or vending facilities, that help enhance the usability of the site. Example response is below:


```
[
  {
    "siteId": "WA00005IS0001170NFTLEWSWS",
    "timeStamp": "2025-11-20T12:00:00.000Z",
    "relevantHighway": "I-5",
    "referencePost": "MP 72",
    "exitID": "72A",
    "directionOfTravel": "NB",
    "name": "Toutle River Safety Rest Area",
    "location": {
      "latitude": 46.12345,
```

```
      "longitude": -122.12345
    },
    "ownership": "WSDOT",
    "capacity": 40,
    "amenities": ["RESTROOMS", "VENDING"],
    "images": ["https://..."],
    "logos": ["https://..."]
  }
]
```

Real-Time Parking Availability (dynamic_public)

This endpoint is for delivering real-time dynamic parking availability. It provides the most up-to-date status for each TPIMS site, including the number of available spaces, the current trend in parking activity, and operational indicators such as whether the site is open or whether the data are considered reliable. The feed is refreshed every 5 minutes and is fully synchronized with the service-provider data stream. Key fields include **reportedAvailable**, which captures the current number of open truck parking spaces; **trend**, which classifies activity as FILLING, STEADY, or CLEARING; **open**, indicating whether the facility is operational or affected by construction or closure; **trustData**, which reflects data reliability; and **capacity**, representing the site's total truck parking capacity. Together, these elements support navigation applications, traveler messaging, and real-time monitoring dashboards.

```
[
  {
    "siteId": "WA00005IS0001170NFTLEWSWS",
    "timeStamp": "2025-11-21T10:00:00.123Z",
    "timeStampStatic": "2025-11-20T12:00:00.000Z",
    "reportedAvailable": "12",
    "trend": "STEADY",
    "open": true,
    "trustData": true,
    "capacity": 40
  }
]
```

Archival Records with Verification Information (dynamic_archive)

This endpoint provides the most recent archived record for each site, along with key verification parameters used to ensure long-term data quality and system reliability. The archival dataset supports calibration activities, post-analysis of sensor behavior, and ongoing monitoring to detect anomalies or drift in reported availability. Each record includes the lastVerificationCheck, indicating when the site was last manually or automatically validated against ground truth; the verificationCheckAmplitude, showing the magnitude of any adjustment applied during verification; the lowThreshold, defining the cutoff for issuing a “LOW” availability message; and trueAvailable, which records the validated ground-truth number of available spaces. Together, these fields provide essential context for evaluating sensor performance and maintaining the accuracy and trustworthiness of the TPIMS data stream.


```
[
  {
    "siteId": "WA00005IS0001170NFTLEWSWS",
    "timeStamp": "2025-11-21T10:00:00.123Z",
    "timeStampStatic": "2025-11-21T09:55:00.000Z",
    "reportedAvailable": "12",
    "trend": "STEADY",
    "open": true,
    "trustData": true,
    "capacity": 40,
    "lastVerificationCheck": "2025-11-21T08:30:00.000Z",
    "verificationCheckAmplitude": 3.0,
    "lowThreshold": 5,
    "trueAvailable": 11
  }
]
```


Prediction

This endpoint represents WSDOT’s major extension beyond the current MAASTO TPIMS Data Exchange Specification, which only defines real-time status fields such as availability, trend, and open/closed indicators and does not include any prediction elements. Because WSDOT is the first state to operationalize truck-parking prediction within a TPIMS environment, the existing standard must be thoughtfully extended while maintaining full backward compatibility.

During the pilot project, prediction results were displayed only through a custom web interface that provided textual likelihood messages (e.g., “very likely to have 0–2 spaces available in 1 hour”). These outputs were not accessible through the official TPIMS API and therefore could not be used by third-party applications or partner agencies.

In the current implementation, predictions are incorporated directly into the dynamic feed as optional JSON fields, making them easily accessible to external applications. The enhanced endpoint now provides multi-horizon forecasts—including 10 minutes, 30 minutes, 1 hour, 2 hours, 3 hours, 4 hours, 6 hours, and 8 hours—generated by a deep-learning time-series model trained individually for each site. The model ingests the most recent observed availability and produces all forecast horizons in a single inference step, with predicted values floored to integers for operational usability.

To ensure transparency and support continuous quality assurance, the API also reports accuracy metrics for each prediction horizon. Accuracy is calculated using a relative-error percentage, defined as:

```
accuracy = max(0, min((1 - abs(reportedAvailable - predictedAvailable) / reportedAvailable) * 100, 100))
```

This score reflects how close the prediction was to reality:

- 100 percent means the prediction was exact.
- 50 percent means the prediction differs from the actual value by half.
- The value is bounded between 0 and 100 to avoid negative accuracy or inflation.

Special handling is applied when the reported availability is “LOW” or 0, where a very small placeholder value is used to avoid division-by-zero errors while still preserving meaningful comparisons.

By embedding both predictive outputs and accuracy indicators directly into the TPIMS API, WSDOT transforms the system from a purely real-time status feed into a forward-looking decision-support service. This enhancement enables more effective routing, proactive parking advisories, and improved logistics planning while remaining fully compatible with clients that only consume standard MASSTO-defined fields.


```
[
  {
    "siteId": "WA00005IS0001170NFTLEWSWS",
    "timeStamp": "2025-11-21T10:00:00.123Z",
    "timeStampStatic": "2025-11-21T09:55:00.000Z",
    "reportedAvailable": "12",
    "trend": "STEADY",
    "open": true,
    "trustData": true,
    "capacity": 40,

    "predict_10m": 11,
    "predict_30m": 10,
    "predict_1hr": 9,
    "predict_2hr": 7,
    "predict_3hr": 6,
```

```
      "predict_4hr": 5,
      "predict_6hr": 4,
      "predict_8hr": 3,

      "acc_10m": 95,
      "acc_30m": 93,
      "acc_1hr": 90,
      "acc_2hr": 88,
      "acc_3hr": 85,
      "acc_4hr": 82,
      "acc_6hr": 80,
      "acc_8hr": 78
    }
  ]
```


API Exception Cases

Certain truck-parking locations—particularly weigh stations that allow parking only during specific hours—require the API to account for temporary or scheduled closures. In these situations, the system must override the normal occupancy reporting and explicitly indicate that the site is closed, following the MASSTO standard (open: true / false / null). When a closure occurs, the API returns open: false regardless of sensor-detected availability, ensuring downstream systems and traveler information signs reflect the correct operational status (Figure 6-4).

In addition to handling scheduled closures, the API also provides administrative flexibility to manually adjust the trustData status for each site. Certain locations may undergo maintenance, sensor recalibration, or validation activities, during which their occupancy data should not be used for public dissemination. Through the internal management interface, administrators can toggle trustData between true, false, or null, allowing the system to immediately flag whether the reported availability is reliable. Downstream partners—such as Drivewyze—actively use this field to determine whether a site's information should be displayed to drivers. This prevents user confusion and ensures that only verified, high-confidence data are integrated into third-party applications and traveler-information services.

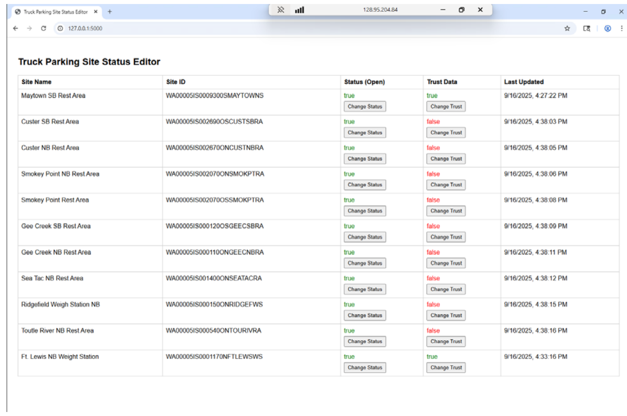

Figure 6-4 API Exception cases interface.

6.2 USER INTERFACE

6.2.1 Review of Previous Truck Parking System Interface

To contextualize the development of the Washington State TPIMS interface, the team also reviewed existing truck-parking information platforms used in other regions and the private sector. Two representative examples are shown below.

Trucks Park Here

This platform is a multi-state initiative designed to help drivers *find parking* by aggregating information from 511 traveler-information systems and participating private truck stops into a unified map (Figure 6-5). The platform also provides a "Developer Information" section, offering APIs for third-party integration—an approach aligned with the TPIMS objective of supporting external app developers.

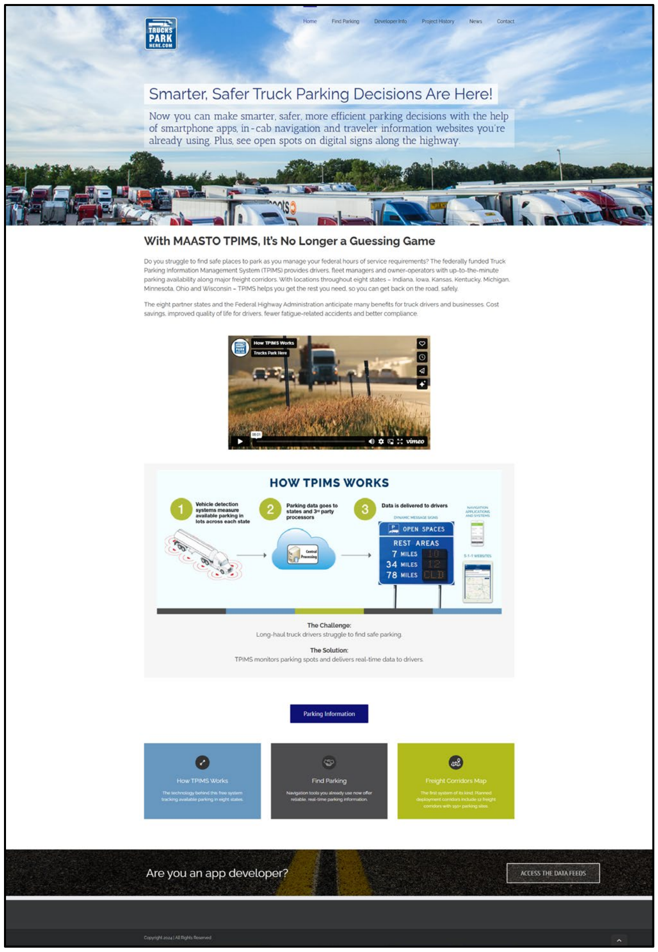

Figure 6-5 Example interface (https://trucksparkhere.com/).

Truck Parking Club

This is a private-sector platform that aggregates live parking availability from partner properties and supports a booking/reservation system for private lots. The interface allows drivers to view real-time space counts, filter nearby parking, and reserve spaces directly through the application (Figure 6-6). This model demonstrates features that could inform future TPIMS enhancements, including mobile-friendly map design, real-time status updates, and integration with privately owned parking supply.

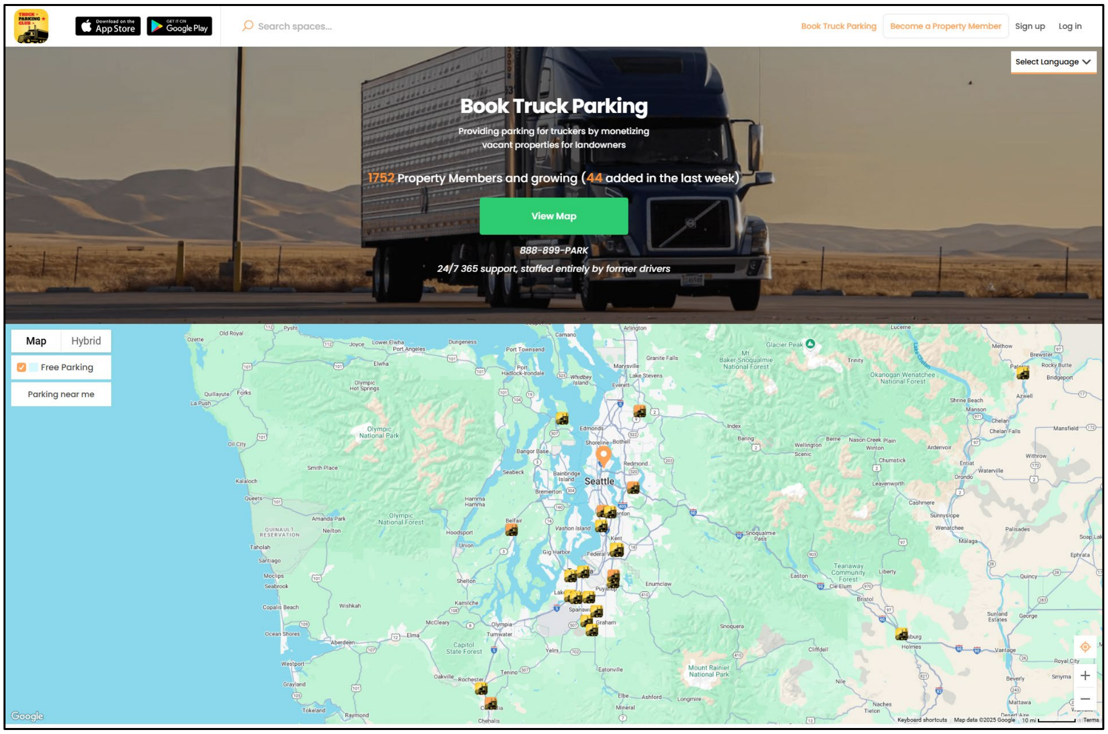

Figure 6-6 Example interface (https://truckparkingclub.com/).

6.2.2 Pilot Prototype Interface (Wix)

The pilot interface was initially developed on a Wix website. This platform provided a clean and visually appealing layout, which allowed for rapid prototyping and stakeholder demonstrations (Figure 6-7). However, Wix requires an ongoing monthly subscription and has limited flexibility for advanced customization. In particular, it restricts the use of custom programming languages and scripts, which limits our ability to integrate dynamic features, external APIs, and custom data visualizations necessary for the full TPIMS deployment. The pilot interface is available at https://nutvaraj.wixsite.com/wsdotparking2

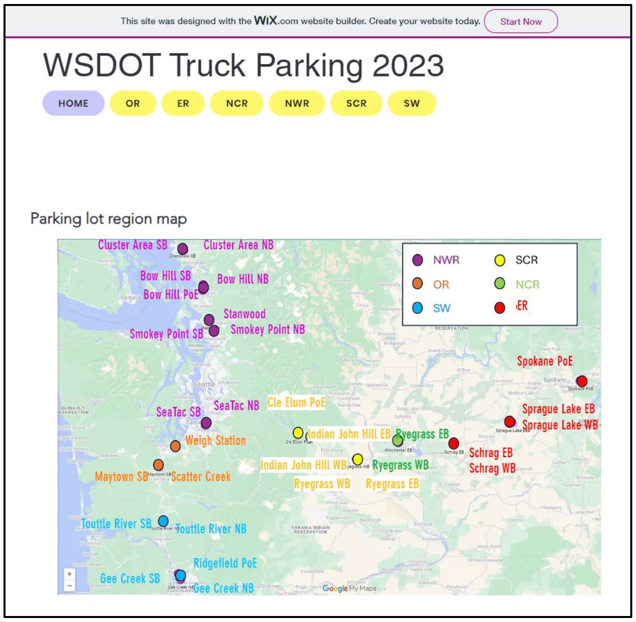

Figure 6-7 Pilot user interface developed on Wix.

6.2.3 Public Interface for Dissemination (npm/React)

During evaluation of the prototype with the Streamlit interface (Figure 6-8), several usability and accessibility issues were identified. Stakeholders noted that the interface should focus on reporting the chance of finding a parking spot rather than simply displaying current availability, aligning the presentation with how drivers make decisions. Mobile usability was also highlighted, as the prototype did not fully meet standards for cellphone-friendly layouts and responsiveness. Additional features—such as providing optional access to CCTV camera snapshots for validation—were suggested, although feasibility and privacy considerations remain open questions. Stakeholders also raised concerns regarding operational costs associated with integrating custom mapping tools. Finally, the existing red/yellow/green color scheme was deemed unsuitable for color-blind users, indicating the need for a more inclusive visualization strategy in future versions.

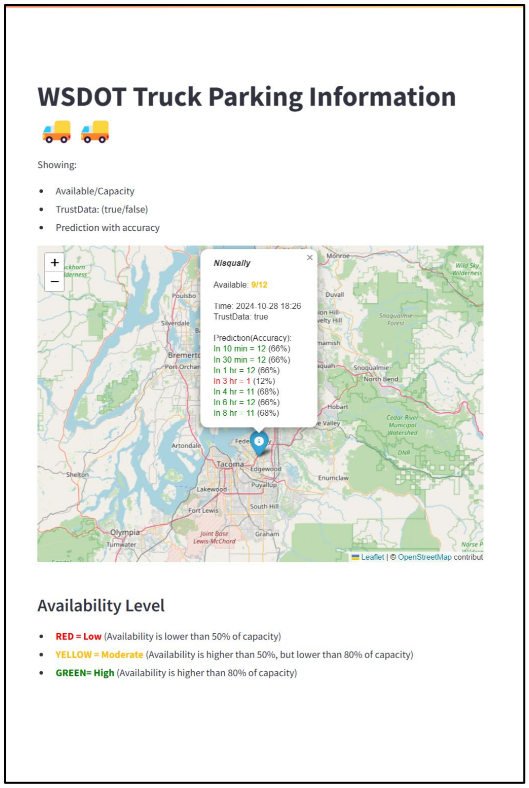

Figure 6-8 Prototype developed with Streamlit.

Therefore, we developed a web-based dashboard that displays truck parking prediction data on npm (Figure 6-9). The dashboard can be accessed at http://128.95.204.84:3000/ and is designed to run on any modern web browser, making it usable across different devices including desktops, tablets, and smartphones. Data on the dashboard are refreshed every 5 minutes, providing near real-time information. The system currently integrates data from eleven sites, allowing users to view both reported availability and predictive analytics for truck parking across multiple locations.

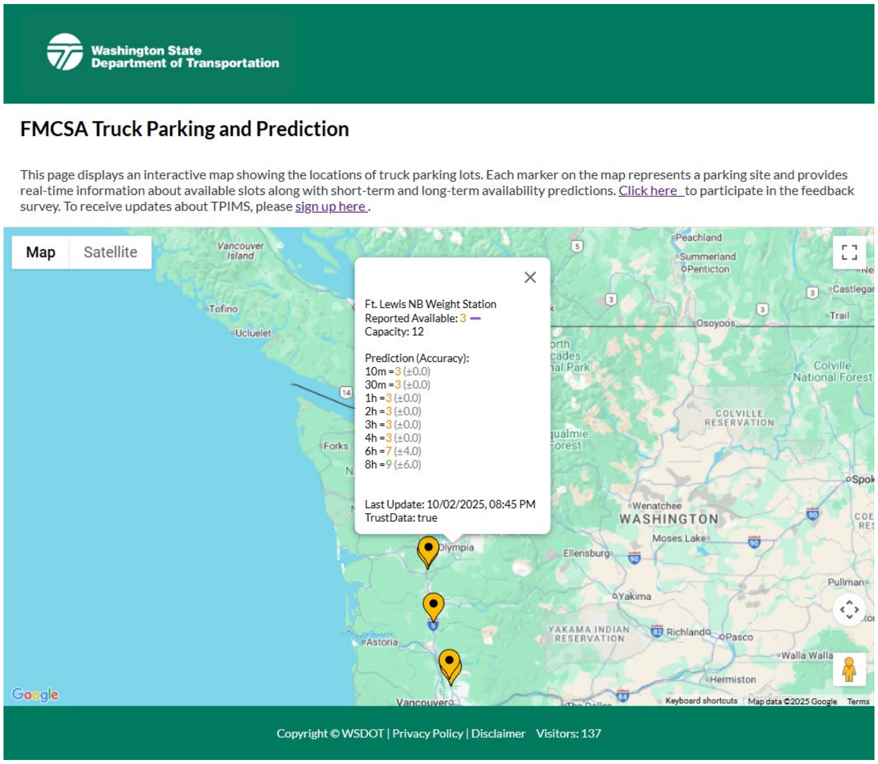

Figure 6-9 UW TPIMS dashboard

In the production public interface, the prediction output is presented differently from the raw accuracy values reported through the API. Instead of displaying percentage accuracies for each prediction horizon, we convert the accuracy metric into an intuitive uncertainty range (e.g., 5 ± 2 spaces). This is calculated by rearranging the accuracy equation and solving for the possible error around each predicted value. The interface automatically updates these values every 5 minutes, giving drivers a clearer representation of prediction confidence without requiring them to interpret accuracy formulas.

$$Accuracy = 1 - \frac{|TruthValue - Predicted|}{TruthValue} \qquad \text{Eq. 6.1}$$

In future versions, the system will incorporate conformal prediction, a statistically grounded framework that generates valid prediction intervals based on historical error patterns. This will provide uncertainty bounds that are theoretically calibrated and more reflective of real-world model performance, allowing users to better judge the reliability of occupancy forecasts at different time horizons.

6.2.4 Internal Operator Dashboard (Streamlit)

To support day-to-day monitoring and system management, the team developed an internal operator dashboard using Streamlit (Figure 6-10). This tool provides WSDOT and UW staff with real-time visibility into the health and performance of the truck-parking system. The dashboard integrates several key functions, including monitoring current availability across all sites, detecting anomalies such as stalled sensors or inconsistent data patterns, and reviewing prediction accuracy at multiple time horizons. The interface also includes a section for calibration-related information; however, the automated calibration workflow is still under development, and the dashboard currently serves primarily as a placeholder for future integration. Designed exclusively for internal use, this dashboard enables rapid diagnosis, troubleshooting, and quality control—capabilities that are not part of the public interface but are essential for maintaining system reliability during both development and deployment.

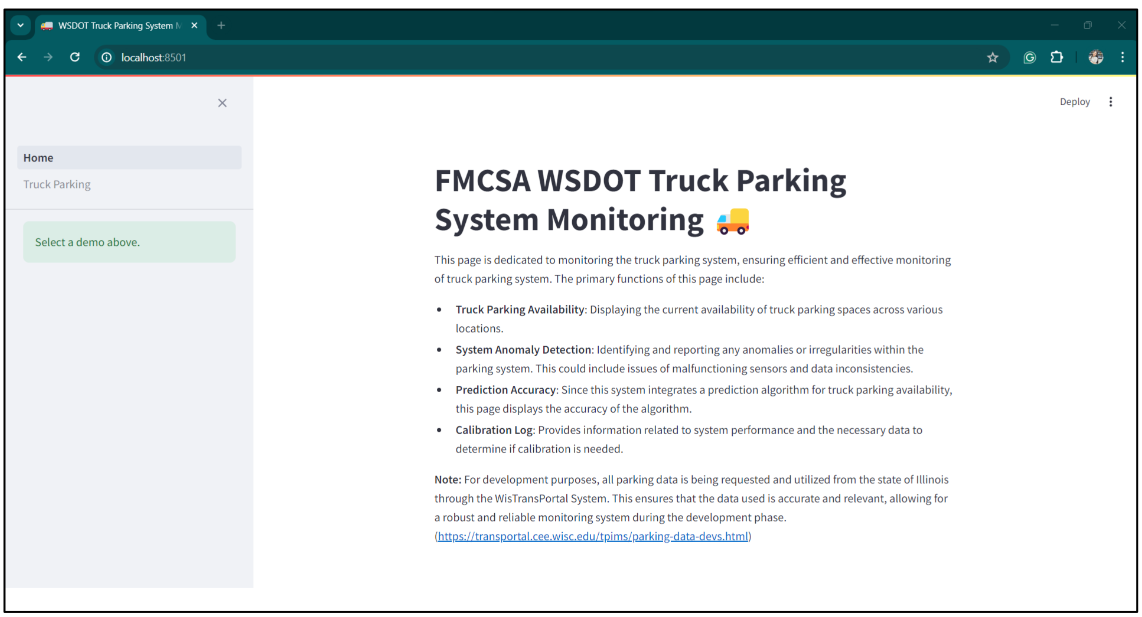

Figure 6-10 Internal operator dashboard.

Trucks Parking Availability Module

_The *Truck Parking Availability* module provides operators with a detailed, query-ready view of occupancy data across all monitored sites (Figure 6-11). Users can examine raw data in tabular form—showing reported availability, capacity, and calculated availability ratios at each timestamp—or interactively filter and select specific sites for closer inspection. The dashboard also supports custom queries, enabling staff to isolate particular time windows or locations for diagnostics. Visual charts automatically update based on the selected sites, allowing operators to observe availability trends over time and quickly identify unusual patterns that may signal sensor malfunctions or data inconsistencies. This combination of tabular and visual analytics provides a flexible and efficient workflow for verifying system behavior during development and ongoing operations.

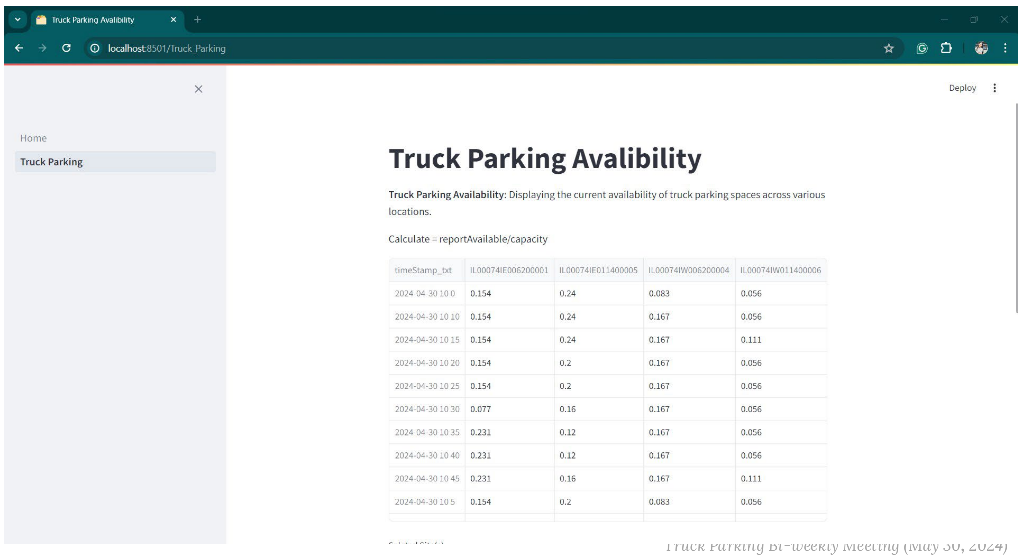

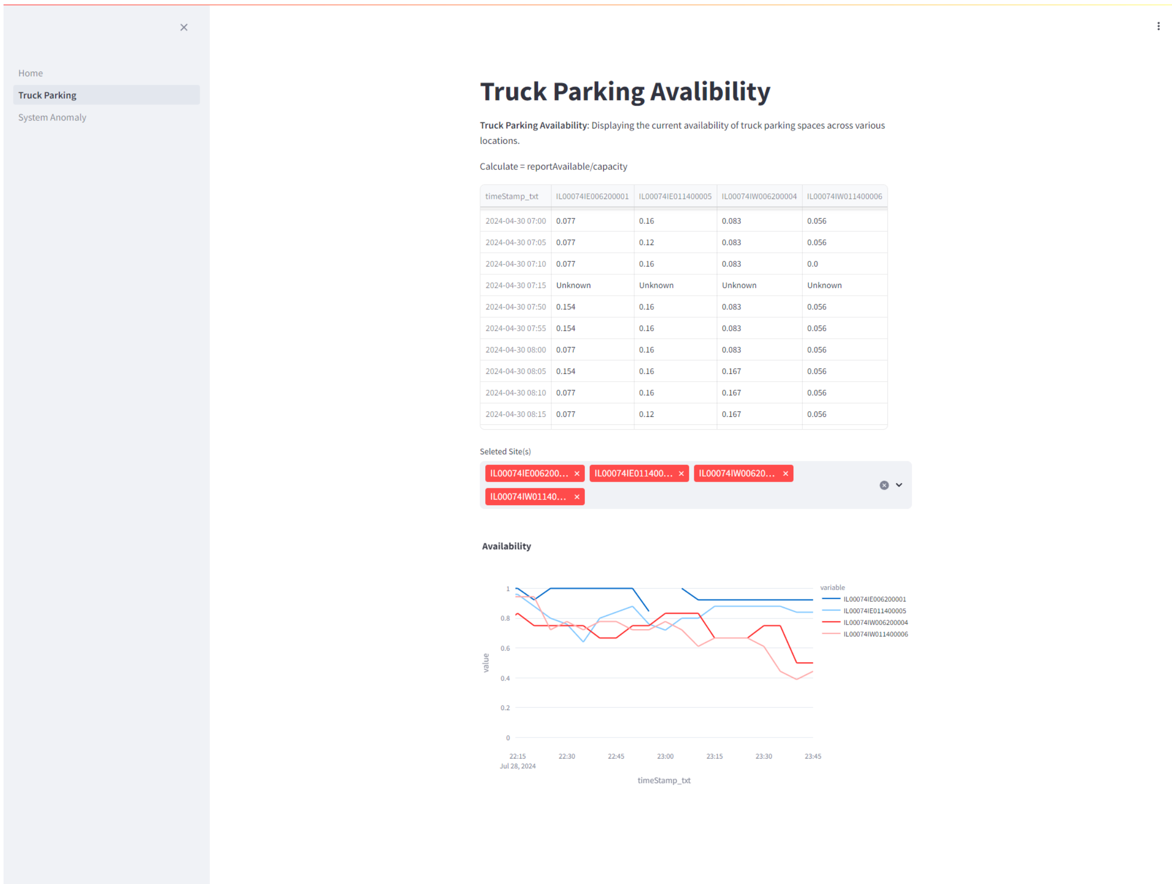

Figure 6-11 Truck Parking Availability module.

System Anomaly Detection Module

The *System Anomaly Detection* module helps operators quickly identify irregularities in data flow from each site (Figure 6-12). In the raw availability tables, values labeled None indicate a system-level failure, meaning our backend was unable to load data from the sensor or process them correctly. In contrast, values labeled Unknown represent a sensor-level failure, where the system is running, but the sensor did not return usable data. To simplify monitoring, the dashboard converts these conditions into three numerical status codes—1 (normal operation), 0 (sensor not responding), and –1 (system down)—and displays them both in tabular format and in a time-series plot. This visualization allows staff to quickly pinpoint when failures occur, differentiate between backend issues and field-sensor problems, and prioritize troubleshooting accordingly.

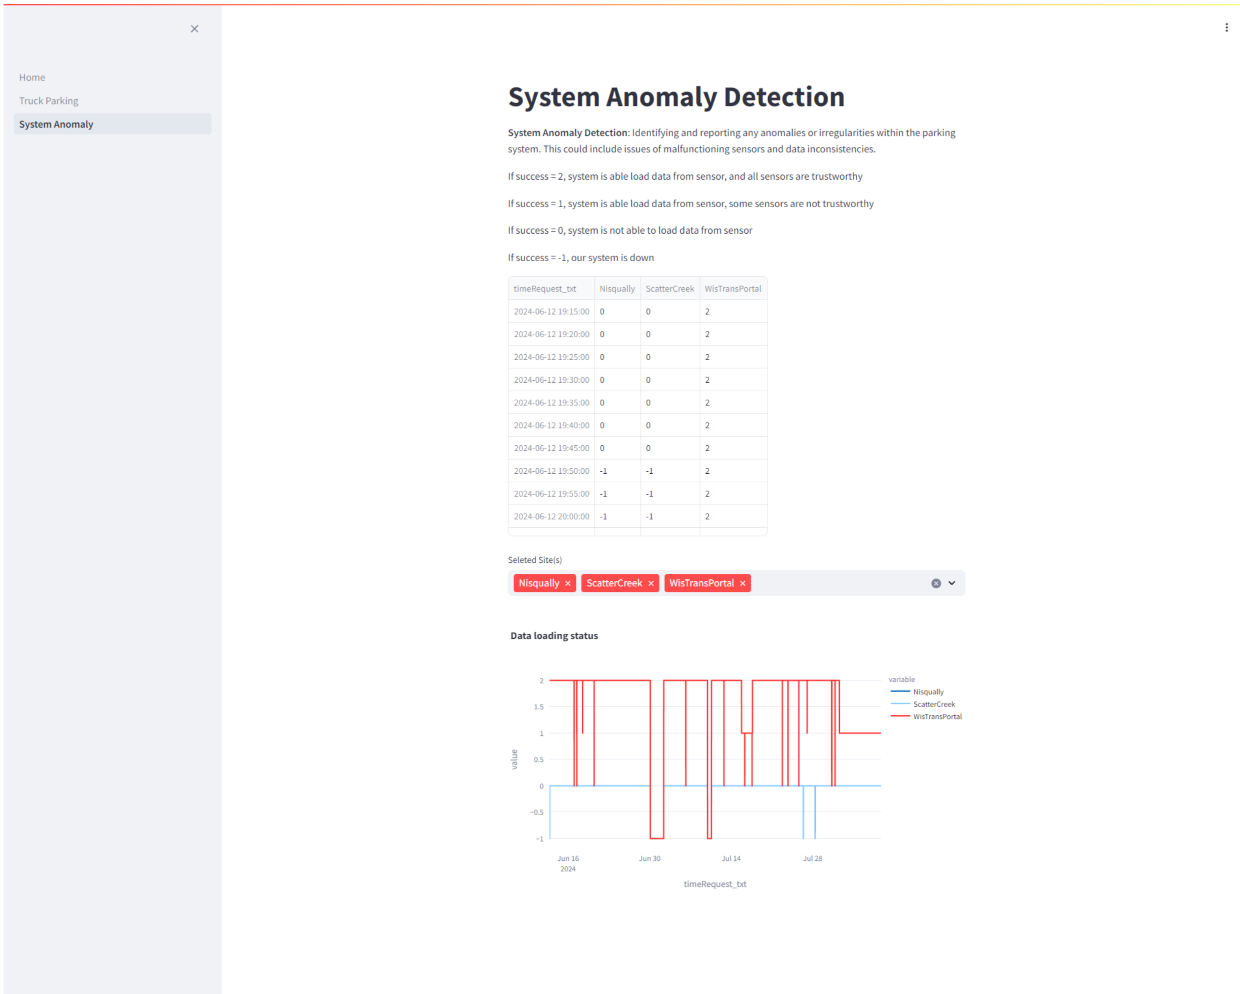

Figure 6-12 System anomaly detection module.

<u>Prediction Module</u>

The *Prediction* module allows operators to interactively explore the performance of the truck-parking prediction algorithm (Figure 6-13). Users can select any site from the system and choose a prediction horizon—from 10 minutes up to 8 hours ahead—to visualize how predicted availability compares with actual occupancy over time. In addition to the time-series plot, the dashboard displays the **latest accuracy values** for each prediction interval, along with a comparison to the accuracy calculated 5 minutes earlier. This provides staff with a quick, at-a-glance indicator of whether model performance is stable, improving, or deteriorating. Together, these tools support continuous monitoring of model behavior and help identify when retraining or recalibration may be necessary.

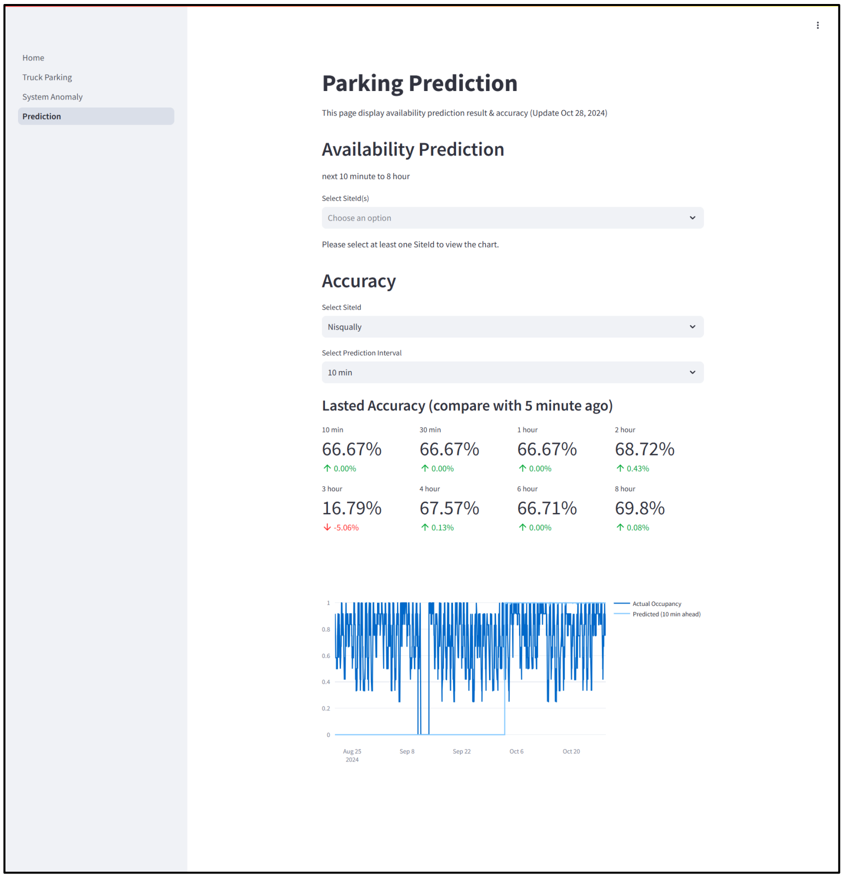

Figure 6-13 Prediction module.

7. USER FEEDBACK

7.1 USER FEEDBACK PLAN

The user feedback plan was designed to systematically capture truck drivers' experiences while interacting with the truck-parking information platform. Feedback will be collected throughout the two-month remote testing period during which participants use the system to check real-time availability, explore multi-timescale predictions, and filter parking locations based on amenities, distance, and availability. User interaction data—including screen recordings of navigation behavior—will support an in-depth assessment of usability, clarity of interface elements, and overall workflow efficiency. A post-interaction survey will gather qualitative insights regarding ease of use, feature relevance, and any difficulties encountered. Together, these feedback channels will help identify specific areas for improvement, prioritize enhancements for future development, and ensure that the system aligns with driver needs and operational contexts at the Nisqually site.

7.2 SURVEY

To complement behavioral data from user testing sessions, a structured online survey was developed to capture users' perceptions of the system's usability, clarity, and overall effectiveness (Figure 7-1). The survey is currently implemented using SurveyMonkey, which offers built-in analytics tools such as response summaries, item-level statistics, and trend visualization. These capabilities allow the team to efficiently monitor participation, identify common feedback themes, and compare responses across demographic groups. The questionnaire included sections on user demographics, ease of navigation, usefulness of real-time and prediction features, and any challenges encountered while using the application. Insights gathered from SurveyMonkey will directly inform refinement of the interface and prioritization of future feature enhancements.

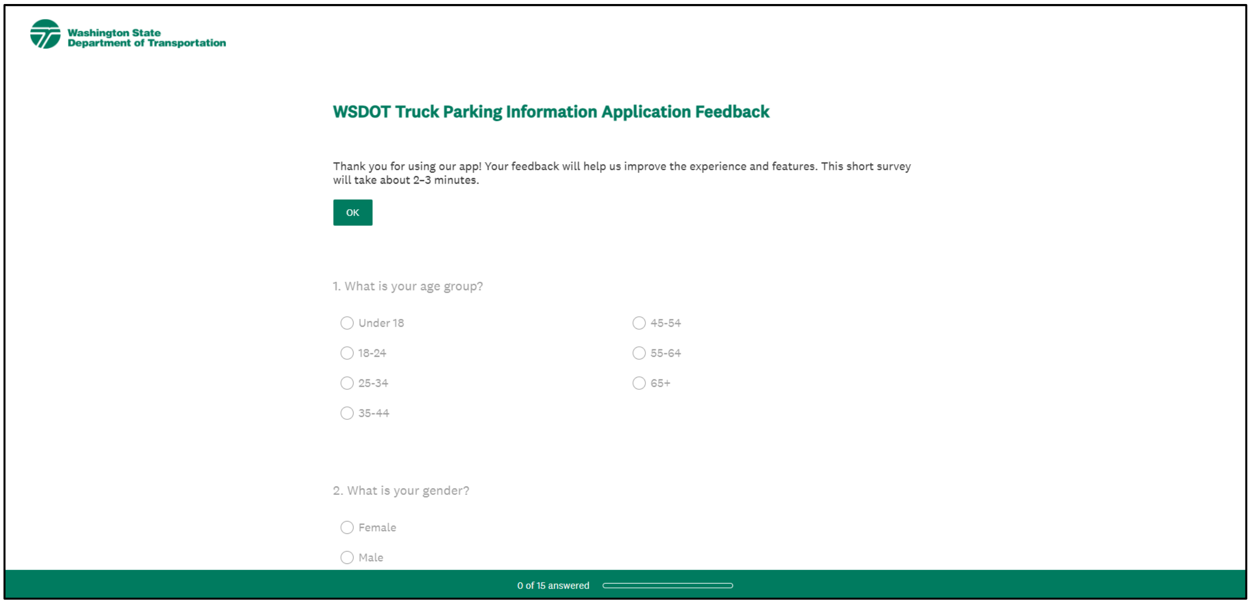

Figure 7-1 SurveyMonkey-based user feedback form used to collect demographic information and evaluate usability and feature satisfaction for the truck-parking application.

To follow up with users and gather feedback, we distributed a survey developed collaboratively by WSDOT and the working group. The questions covered four main areas: user background (e.g., primary role, how they heard about TPIMS), information sources and usage (how they currently obtain truck parking information, frequency of reliance on the system, and use of the prediction feature), perceived usefulness and accuracy (whether the information was helpful in finding parking and how accurate it was in comparison to actual experience), and impacts and feedback (changes to routes, schedules, or rest stop decisions, as well as additional comments or suggestions). Respondents were also invited to provide their email if they wished to receive future updates. Over the analysis period (September 2 – October 2, 2025), we received two responses.

The first respondent identified themselves as a regional planner working on freight policy. They reported accessing parking information primarily through the University of Washington/WSDOT website and first learned about TPIMS via a WSDOT email. While they considered the information not very useful, citing that the Gee Creek NB Rest Area was unavailable, they were uncertain about the overall accuracy of the data. The respondent indicated that they rely on the system only a few times per month, mainly to provide feedback for planning purposes, and noted that they had never changed routes or schedules since they did not operate a truck. They had not yet used the prediction feature, although they expressed appreciation for its availability. Overall, they were supportive of the project.

The second respondent was a long-haul truck driver who reported using Trucker Path as their primary source of parking information and first learned about TPIMS through a colleague. They found the system not at all useful, explaining that they were unable to view the map, which made the tool unusable for them. They were uncertain about the accuracy of the information and indicated they almost never relied on the system in its current form. They had never changed routes or schedules based on TPIMS data and had not used the prediction feature. Their main suggestion was to modify the user interface, as the current experience was not satisfactory given that they could not see the map.

From the two responses, we observed distinct perspectives: a regional planner valued the system as a concept but found its usefulness limited, while a long-haul truck driver expressed frustration with the user interface, particularly the inability to view the map. Both responses highlighted that the system was not yet reliable or practical for day-to-day parking decisions, with accuracy and usability remaining key concerns. These insights suggested that improving data reliability and enhancing the user interface were critical next steps to increase adoption and effectiveness.

7.3 ANALYTICS FROM THE UW TPIMS DASHBOARD

To monitor engagement with the dashboard, we tracked the number of unique IP addresses over the first month of operation (Figure 7-2). Before the system's go-live, 48 unique IP addresses were recorded. One week after launch, on September 9, 2025, this number rose sharply to 91. By the second week, the count had increased to 109, followed by 118 in the third week and 131 in the fourth week. These results indicated a steep rise in adoption during the initial phase, followed by a slower but steady increase thereafter.

It is important to note that these figures represent approximate unique visitors rather than exact individuals, as unique IP addresses do not always correspond to unique people; a single user may appear under multiple IPs because of dynamic addressing, virtual private networks (VPNs), or network changes, while multiple users may share the same IP through corporate networks.

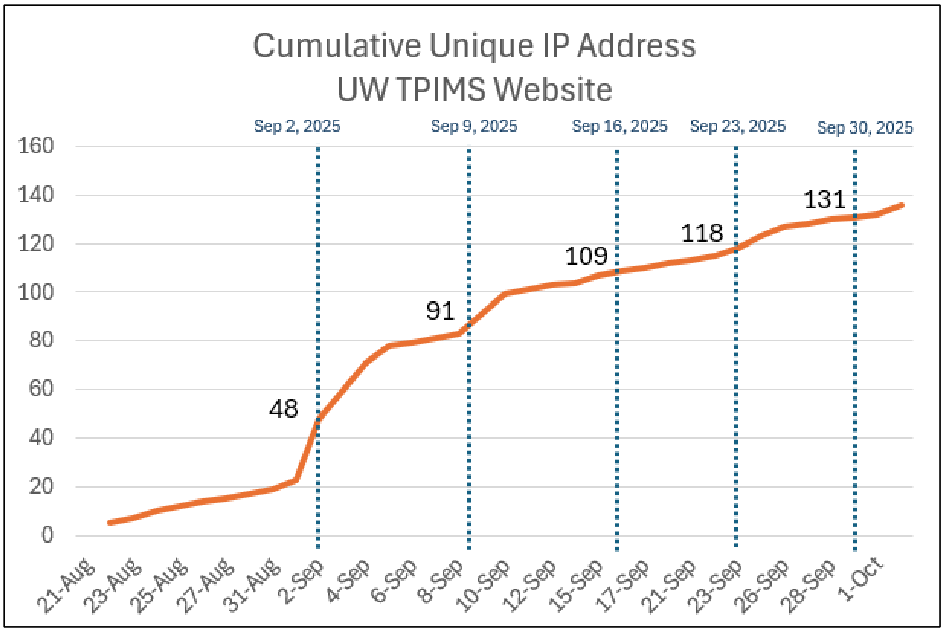

Figure 7-2 Cumulative unique IP address at UW TPIMS website

Using IP addresses, we were able to approximate the geographic location of website visitors. Figure 7-3 presents only results from the United States (94 out of 123 records). Among these, the majority were concentrated in Washington state (32 of 94 records), highlighting that a large portion of users originated from Washington.

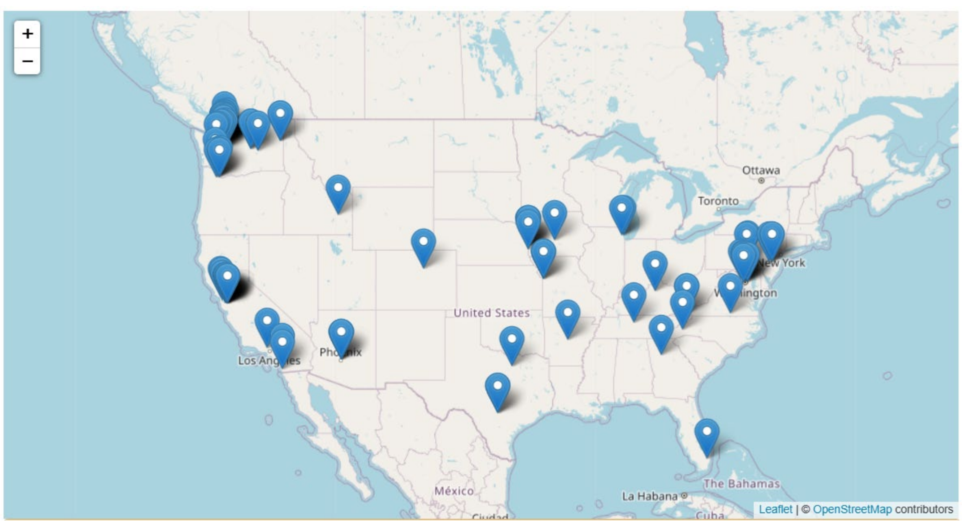

Figure 7-3 Distribution of IP addresses in USA visiting the UW TPIMS website (Aug 21, 2025 – Oct 2, 2025), based on data from https://www.showmyip.com

We performed additional data collection using Google Analytics embedded on the website. Data were collected over a seven-day period (September 25 – October 1, 2025). During this time, eleven users were identified, with an average engagement time of 34 seconds.

Device type analysis showed that eight users accessed via web browsers and three via mobile devices. This indicated that the dashboard was used more on desktop platforms than on mobile devices. Because the dashboard was not developed as a dedicated mobile application, it may not be fully optimized or widely recognized for mobile use. Developing a mobile-friendly version could expand adoption, particularly since truck drivers are more likely to access information on mobile devices for trip planning.

8. CONCLUSIONS

8.1 SUMMARY OF FINDINGS

The project successfully delivered a comprehensive, scalable framework for Washington State's Truck Parking Information Management System (TPIMS), integrating multi-sensor detection, real-time data aggregation, deep-learning prediction, and a standardized API for public and third-party access. Through extensive field evaluation, we confirmed that both in/out sensors and slot-by-slot systems can reliably measure truck presence when properly calibrated and validated against ground-truth video. The development of a unified data architecture and archival structure enables long-term monitoring of verification events, anomaly detection, and performance tracking across all sites. The prediction component—built using site-specific LSTM models—demonstrated strong accuracy in forecasting near-term parking availability and can be retrained as new statewide data become available. The API, designed to follow MASSTO standards, provides clean, consistent outputs while accommodating exception cases such as temporary site closures and sensor trust-state changes.

Operationally, the project highlighted the importance of continuous maintenance after deployment, including system monitoring, data-quality audits, and periodic model updates to adapt to evolving demand patterns. The statewide deployment scenario showed potential for learning shared patterns, improving model generalization, and supporting broader WSDOT operational decision-making. Overall, the project established a robust foundation for a fully integrated statewide truck parking information system, providing actionable intelligence for drivers, freight stakeholders, and transportation managers while supporting future expansion and innovation.

8.2 RECOMMENDATIONS FOR FUTURE DEPLOYMENTS AND MAINTENANCE

8.2.1 API Support

Future enhancements to the API will focus on improving service reliability, interoperability, and long-term maintainability. Expanding system-level health checks—such as automated server status reporting, endpoint availability monitoring, and load-handling diagnostics—will help ensure stable operation under increasing data volume and partner usage. Library and dependency management may also require periodic updates to maintain compatibility, security, and performance as the system evolves. Additional versioning support, more descriptive error handling, and expanded documentation will further streamline integration for third-party developers and promote broader adoption across public and private partners.

8.2.2 Detection System

Continuous improvement of the detection system will remain critical for maintaining data quality and ensuring that occupancy and prediction outputs remain trustworthy. Future work will emphasize strengthening sensor performance monitoring, refining trust-data logic, and integrating automated alerts when detection patterns deviate from expected baselines. Enhanced calibration pipelines, including more frequent or automated comparisons with video ground truth, will help track sensor drift and maintain long-term accuracy. As detection technologies evolve, the system may also incorporate more advanced sensing modalities or hybrid approaches to improve robustness across diverse site conditions.

8.2.3 Supporting Statewide Operational Decisions

As deployment expands, the system can move beyond real-time aggregation to providing actionable intelligence for WSDOT operations. By leveraging predictive analytics and cross-site pattern recognition, it will be able to flag emerging congestion, anticipate parking shortages along freight corridors, and detect abnormal sensor behavior. These insights will enable more proactive actions—such as targeted enforcement, improved traveler information, and timely maintenance intervention—driven by data collected across the statewide network.

8.2.4 System

After project completion, ongoing system-level maintenance will be essential to ensure uninterrupted operation of the TPIMS platform. This includes routine monitoring of server uptime, resource usage, API responsiveness, and database stability, as well as verifying sensor connectivity and data flow from each site. Identifying abnormalities—such as stalled sensors, communication failures, or server overload—will allow IT staff to intervene early and prevent data gaps that would affect downstream analytics and public-facing applications. Continuous oversight of the system infrastructure will ensure reliability, availability, and long-term operational resilience.

8.2.5 Data Analytics

Data analytics activities will focus on validating the accuracy and integrity of both sensor data and model outputs. Periodic audits comparing system outputs with ground-truth observations will help identify detection drift, misclassification, or systematic errors. When anomalies or exceptional patterns are observed—such as abrupt occupancy jumps or inconsistent trends—temporary algorithm adjustments or override mechanisms may be necessary. Regular review of parking patterns and flagged anomalies will ensure that the analytical layer continues to accurately represent real-world conditions and provide actionable insights to WSDOT.

8.2.6 Model

The prediction model will require periodic evaluation to maintain performance over time as new data accumulate and site conditions evolve. Monitoring prediction accuracy, identifying signs of model drift, and comparing predicted values against actual occupancy trends will help determine when retraining is necessary. Refreshing model weights using updated historical data—and updating normalization parameters as needed—will ensure that forecasts remain aligned with real operational behavior. Assessing cross-site generalization will also help determine whether shared or individualized models remain optimal as the statewide system expands. Through these activities, the model will continue to deliver meaningful and trustworthy predictions.

8.2.7 API Exception Case

Certain truck-parking locations—particularly weigh stations that allow parking only during specific hours—require the API to account for temporary or scheduled closures. In these situations, the system must override the normal occupancy reporting and explicitly indicate that the site is closed, following the MASSTO standard (open: true / false / null). When a closure occurs, the API returns open: false regardless of sensor-detected availability, ensuring that downstream systems and traveler information signs reflect the correct operational status. Because closure schedules may vary or change on short notice, future improvements may include tying closure notifications directly to the site's control system or sign infrastructure, allowing automatic real-time updates without manual intervention.

REFERENCES

CalTrans. (2022). California Statewide Truck Parking Study. https://dot.ca.gov/-/media/dot-media/programs/transportation-planning/documents/freight-planning/plan-accordion/catrkpkgstdy-finalreport-a11y.pdf

Gong, S., Qin, J., Xu, H., Cao, R., Liu, Y., Jing, C., ... & Yang, Y. (2023). Spatio-temporal parking occupancy forecasting integrating parking sensing records and street-level images. International Journal of Applied Earth Observation and Geoinformation, 118, 103290.

Gutmann, S., Maget, C., Spangler, M., & Bogenberger, K. (2021). Truck parking occupancy prediction: Xgboost-LSTM model fusion. Frontiers in Future Transportation, 2, 693708.

ODOT. (2020). Oregon Commercial Truck Parking Study. https://www.oregon.gov/odot/Planning/Documents/OCTPS_final_report_with_Appendices_and_exec_summary-Full_Report.pdf

ODOT. (2024). Truck parking on Interstate 5 will be safer and easier thanks to new grant-funded project. https://content.govdelivery.com/accounts/ORDOT/bulletins/387fa51

Slavova, S., Piest, J. P. S., & van Heeswijk, W. (2022). Predicting truck parking occupancy using machine learning. Procedia computer science, 201, 40-47.

Yang, H. et al., "Truck Parking Pattern Aggregation and Availability Prediction by Deep Learning," in IEEE Transactions on Intelligent Transportation Systems, vol. 23, no. 8, pp. 12778-12789, Aug. 2022, doi: 10.1109/TITS.2021.3117290.

Washington Truck 3 Parking Assessment, Appendix H (July 29, 2022 version)

WSDOT Geospatial Open Data Portal. https://gisdata-wsdot.opendata.arcgis.com/datasets/WSDOT::wsdot-highway-urban-and-urbanized-area-2013/explore?location=47.726963%2C-121.890924%2C10.76 (Accessed Oct 5, 2022)

BIBLIOGRAPHY

Washington State Joint Transportation Committee, Action Plan Supplement

2021 Washington State Truck Parking Workshop Overview and potential solutions for consideration

Boris, C. and Brewster, R., 2018. A comparative analysis of truck parking travel diary data. *Transportation Research Record*, *2672*(9), pp.242-248

Title VI Notice to Public

It is the Washington State Department of Transportation's (WSDOT's) policy to assure that no person shall, on the grounds of race, color, or national origin, as provided by Title VI of the Civil Rights Act of 1964, be excluded from participation in, be denied the benefits of, or be otherwise discriminated against under any of its programs and activities. Any person who believes his/her Title VI protection has been violated, may file a complaint with WSDOT's Office of Equity and Civil Rights (OECR). For additional information regarding Title VI complaint procedures and/or information regarding our non-discrimination obligations, please contact OECR's Title VI Coordinator at (360) 705-7090.

Americans with Disabilities Act (ADA) Information

This material can be made available in an alternate format by emailing the Office of Equity and Civil Rights at wsdotada@wsdot.wa.gov or by calling toll free, 855-362-4ADA(4232). Persons who are deaf or hard of hearing may make a request by calling the Washington State Relay at 711.