Improving Work Zone Management and Safety through AI-Powered Connected Vehicle Data Analysis

Final Report
April 2026

SWZDI
Smart Work Zone Deployment Initiative

Center for Transportation Research and Education

Iowa State University
Institute for Transportation

Sponsored by
Smart Work Zone Deployment Initiative
(Part of TPF-5(438))
Federal Highway Administration
(InTrans Project 24-888)

About the Smart Work Zone Deployment Initiative

The Smart Work Zone Deployment Initiative (SWZDI) is a transportation pooled fund that supports research investigations into better ways to improve the safety and efficiency of traffic operations and highway work in work zones. The primary objective is to promote and support research and outreach activities that focus on innovative policies, processes, tools, and products that enhance the implementation, safety, and mobility impacts of work zones.

About the Center for Transportation Research and Education

The mission of the Center for Transportation Research and Education (CTRE) at Iowa State University is to conduct basic and applied transportation research to help our partners improve safety, facilitate traffic operations, and enhance the management of infrastructure assets.

About the Institute for Transportation

The mission of the Institute for Transportation (InTrans) at Iowa State University is to save lives and improve economic vitality through discovery, research innovation, outreach, and the implementation of bold ideas.

Iowa State University Nondiscrimination Statement

Iowa State University does not discriminate on the basis of race, color, age, ethnicity, religion, national origin, pregnancy, sexual orientation, genetic information, sex, marital status, disability, or status as a U.S. Veteran. Inquiries regarding nondiscrimination policies may be directed to Office of Equal Opportunity, 2680 Beardshear Hall, 515 Morrill Road, Ames, Iowa 50011, telephone: 515-294-7612, email: eooffice@iastate.edu.

Disclaimer Notice

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. The opinions, findings and conclusions expressed in this publication are those of the authors and not necessarily those of the sponsors.

This document is disseminated under the sponsorship of the U.S. DOT in the interest of information exchange. The sponsors assume no liability for the contents or use of the information contained in this document. This report does not constitute a standard, specification, or regulation.

The sponsors do not endorse products or manufacturers. Trademarks or manufacturers' names appear in this report only because they are considered essential to the objective of the document.

Quality Assurance Statement

The Federal Highway Administration (FHWA) provides high-quality information to serve Government, industry, and the public in a manner that promotes public understanding. Standards and policies are used to ensure and maximize the quality, objectivity, utility, and integrity of its information. The FHWA periodically reviews quality issues and adjusts its programs and processes to ensure continuous quality improvement.

Iowa DOT Statements

The Iowa Department of Transportation (DOT) ensures non-discrimination in all programs and activities in accordance with Title VI of the Civil Rights Act of 1964. Any person who believes that they are being denied participation in a project, being denied benefits of a program, or otherwise being discriminated against because of race, color, national origin, gender, age, or disability, low income and limited English proficiency, or if needs more information or special assistance for persons with disabilities or limited English proficiency, please contact Iowa DOT Civil Rights at 515-239-7970 or by email at civil.rights@iowadot.us.

The preparation of this report was financed in part through funds provided by the Iowa DOT through its "Second Revised Agreement for Management of Research Conducted by Iowa State University for the Iowa Department of Transportation" and its amendments.

The opinions, findings, and conclusions expressed in this publication are those of the authors and not necessarily those of the Iowa DOT or the U.S. DOT FHWA.

Technical Report Documentation Page

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
Part of TPF-5(438)		

4. Title and Subtitle	**5. Report Date**
Improving Work Zone Management and Safety through AI-Powered Connected Vehicle Data Analysis	April 2026
	6. Performing Organization Code
7. Author(s)	**8. Performing Organization Report No.**
Sanjita Prajapati (orcid.org/0009-0004-7363-7771), Anuj Sharma (orcid.org/0000-0001-5929-5120), and Skylar Knickerbocker (orcid.org/0000-0002-0202-5872)	Part of InTrans Project 24-888
9. Performing Organization Name and Address	**10. Work Unit No. (TRAIS)**
Institute for Transportation Iowa State University 2711 South Loop Drive, Suite 4700 Ames, IA 50010-8664	**11. Contract or Grant No.**
12. Sponsoring Organization Name and Address	**13. Type of Report and Period Covered**
Smart Work Zone Deployment Initiative Iowa Department of Transportation 800 Lincoln Way Ames, IA 50010 Federal Highway Administration U.S. Department of Transportation 1200 New Jersey Avenue, SE Washington, DC 20590	Final Report
	14. Sponsoring Agency Code Part of TPF-5(438)

15. Supplementary Notes

Visit https://swzdi.intrans.iastate.edu/ for color pdfs of this and other research reports.

16. Abstract

This project designed, deployed, and validated a statewide platform for real-time incident detection, daily operations analysis, and data-driven support for smart work zone planning known as ReactorIQ. Unlike traditional systems that depend on fixed roadside sensors, manual contractor logs, or 511 updates, ReactorIQ is built around high-resolution connected vehicle (CV) telemetry as its primary data source. By processing these data at scale, ReactorIQ offers a granular and behaviorally grounded view of how drivers respond to lane closures, queues, and work zone configurations in both urban and rural environments. Insights are made available through an integrated web-based dashboard where users can access CV performance metrics, lane closure detection outputs, queue warning indicators, camera-based analytics, and machine learning results within a single decision-support environment.

This report details the development and key features of ReactorIQ, the challenges overcome and major accomplishments throughout the project, and directions for future work.

17. Key Words	**18. Distribution Statement**
artificial intelligence—cloud computing—connected vehicle—lane closure detection—queue detection—telemetry	No restrictions.

19. Security Classification (of this report)	20. Security Classification (of this page)	21. No. of Pages	22. Price
Unclassified.	Unclassified.	92	NA

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

IMPROVING WORK ZONE MANAGEMENT AND SAFETY THROUGH AI-POWERED CONNECTED VEHICLE DATA ANALYSIS

Final Report
April 2026

Principal Investigator
Anuj Sharma, Co-Director
Real-Time Analytics of Transportation Data (REACTOR) Lab, Iowa State University

Co-Principal Investigator
Skylar Knickerbocker, Researcher
Center for Transportation Research and Education, Iowa State University

Research Assistant
Sanjita Prajapati

Authors
Sanjita Prajapati, Anuj Sharma, Skylar Knickerbocker

Sponsored by
Smart Work Zone Deployment Initiative and
Federal Highway Administration Pooled Fund Study TPF-5(438):
Iowa (lead state), Illinois, Kansas, Michigan, Minnesota, Missouri, Nebraska, Texas, and Wisconsin

Preparation of this report was financed in part
through funds provided by the Iowa Department of Transportation
through its Research Management Agreement with the
Institute for Transportation
(InTrans Project 24-888)

A report from
Smart Work Zone Deployment Initiative
2711 South Loop Drive, Suite 4700
Ames, IA 50010-8664
Phone: 515-294-8103 / Fax: 515-294-0467
https://swzdi.intrans.iastate.edu/

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

This research was conducted under the Smart Work Zone Deployment Initiative (SWZDI) and Federal Highway Administration (FHWA) Pooled Fund Study TPF-5(438), involving the following state departments of transportation (DOTs):

- Iowa (lead state)
- Illinois
- Kansas
- Michigan
- Minnesota
- Missouri
- Nebraska
- Texas
- Wisconsin

The authors would like to thank the FHWA, the Iowa Department of Transportation (Iowa DOT), and the other pooled fund state partners for their financial support and technical assistance.

EXECUTIVE SUMMARY

This project designed, deployed, and validated a comprehensive statewide, real-time work zone intelligence system for Iowa capable of transforming raw connected vehicle (CV) telemetry, camera feeds, and cloud-based analytics into actionable operational insights. The resulting platform, ReactorIQ, integrates data ingestion pipelines, performance metric computation, machine learning components, and intuitive dashboards to provide transportation agencies with a continuous understanding of traffic conditions across active work zones. Through this work, the project demonstrated that CV data can serve as a strong foundation for modern traffic monitoring, achieving a level of spatial coverage and temporal resolution that traditional fixed sensors alone cannot provide.

A major accomplishment of the project was the creation of a cloud-native analytics architecture capable of processing millions of CV records per hour with near-real-time latency. Built on AWS Glue, Apache Spark, and Apache Sedona, the system filters and maps telemetry onto roadway geometries and generates key performance indicators such as queue length, delay, speed degradation, and slowdown frequency. The platform incorporates machine learning techniques to enhance situational awareness, including the detection of lane closures from subtle shifts in lateral driving behavior. These capabilities, combined with real-time alerting and an interactive dashboard, provide operators with improved visibility into merging turbulence, bottleneck formation, shockwave propagation, and behavioral responses to work zone configurations.

The platform supports both operational and planning needs. In real-time, transportation agencies can identify emerging congestion, respond to safety anomalies, and verify incidents through live camera feeds. Over longer time horizons, the system provides a rich archive of performance data that can inform work zone design, taper placement, contractor compliance, and seasonal traffic behavior. Overall, this project demonstrates that integrating CV data, cloud analytics, and AI-driven interpretation can significantly enhance work zone operations and provide a scalable blueprint for the next generation of intelligent transportation systems.

This report details the development and key features of ReactorIQ, the challenges overcome and major accomplishments throughout the project, and directions for future work.

CHAPTER 1. INTRODUCTION

1.1. Project Overview

This project focused on the design, deployment, and validation of a statewide, real-time work zone monitoring and traffic performance intelligence platform for Iowa known as ReactorIQ. Unlike traditional systems that depend on fixed roadside sensors, manual contractor logs, or delayed 511 updates, ReactorIQ is built around high-resolution connected vehicle (CV) telemetry as its primary data source. The platform continuously transforms naturalistic vehicle movement into detailed measures of mobility, safety, and operational performance across Iowa's uninterrupted flow network, demonstrating that CV telemetry data can support a complete work zone intelligence workflow.

Modern transportation networks are rapidly evolving as CVs become more prevalent and as agencies expand their adoption of smart infrastructure. Traditional fixed sensors, such as loops and radar units, often offer incomplete spatial coverage and have limited ability to capture emerging conditions across rural corridors and secondary roadways. The system developed in this project overcomes these limitations by leveraging high-resolution CV telemetry data that provide vehicle-level insight into speed, acceleration, lateral position, and trajectory patterns across the entire state of Iowa. By processing these data at scale, ReactorIQ offers a granular and behaviorally grounded view of how drivers respond to lane closures, queues, and work zone configurations in both urban and rural environments.

ReactorIQ was intentionally built around three core pillars: real-time incident detection, daily operations analysis to strengthen management, and data-driven support for smart work zone planning (Figure 1).

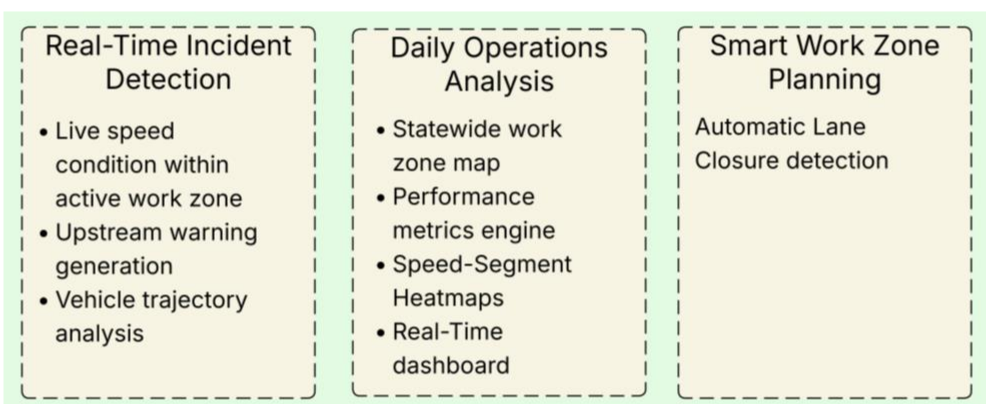

Figure 1.1. ReactorIQ's three core pillars supporting incident detection, daily operations analysis, and smart work zone planning

The first pillar, incident detection, focuses on revealing emerging slowdowns, queue buildups, speed drops, and turbulence patterns on the order of minutes. CV telemetry data are ingested into AWS Glue at frequent intervals and processed through Apache Spark and Apache Sedona pipelines to produce five-minute performance windows, which are then used to detect speed degradation, shockwave formation, and clusters of abrupt decelerations associated with elevated crash risk. These conditions are translated into JSON and XML messages that are made available for use in dynamic message signs (DMS), upstream traveler warnings, and operator alerts.

The second pillar, daily operations and management, provides the operational layer of the system. ReactorIQ continuously computes CV-based performance metrics such as average and percentile speeds, delay and journey-level travel time, queue length and duration, slowdown frequency, and approximate estimated volume. These metrics are exposed through an integrated web-based dashboard that visualizes statewide work zones, real-time speed segment heatmaps, daily time series plots, and vehicle trajectory diagrams. Trajectory visualizations are particularly powerful for revealing shockwaves, merging turbulence, and bottleneck formation patterns that cannot be accurately detected by point detectors.

The third pillar, long-term smart work zone planning, leverages historical CV telemetry data to evaluate future work zone scenarios. The analysis focuses on understanding how traffic conditions such as speed, queue length, and delay evolve under baseline and work zone conditions using observed data patterns. These insights are used to examine alternative deployment windows, lane closure configurations, and staging strategies.

As part of this platform, the dashboard also incorporates outputs from a CV-based automatic lane closure detection model, which identifies the location and timing of lane closures using weekly shifts in speed and lateral position distributions. This directly addresses a long-standing operational problem for the Iowa Department of Transportation (Iowa DOT), namely, identifying temporary work zones or moving lane closures. By inferring closure status from vehicle behavior itself, ReactorIQ improves the accuracy of work zone situational awareness and supports better integration with Work Zone Data Exchange (WZDx) feeds and contractor compliance checks.

As a result, the same CV telemetry data that support real-time monitoring and daily operations also enable planning-time evaluation of work zone impacts, including comparisons between hypothetical closures and historical behavior on similar corridors. Collectively, these three pillars establish a CV-driven framework that elevates ReactorIQ from a monitoring dashboard to an artificial intelligence (AI)-driven work zone management system.

1.2. Challenges Faced and Resolutions

Building a statewide, CV-driven, real-time monitoring ecosystem posed several technical, operational, and cloud infrastructure challenges, requiring multiple cycles of redesign and refinement to achieve both reliability and cost-effectiveness. One of the earliest challenges centered on the design of robust spatial queries capable of accurately joining massive volumes of CV telemetry to the underlying roadway network. Early versions of the pipeline frequently encountered spatial mismatches, leading to incorrect assignment of waypoints to segments.

A second major challenge was detecting queues and slowdowns in low-volume rural corridors with sparse CV penetration. In these environments, naive algorithms that rely on high sample densities or strict thresholds on mean speed can underestimate queues or produce fragmented detections. The system mitigated this by moving toward percentile-based speed metrics and temporal persistence rather than raw averages alone. Five-minute bins were used to stabilize speed estimates, and queue detection relied on the presence of slow-speed clusters that persist over consecutive bins, even when sample counts are low. This design allowed the platform to detect meaningful slowdowns while avoiding overreaction to isolated slow vehicles.

Lane closure detection posed its own set of difficulties. Early approaches depended heavily on the standard deviation of lateral position and other low-order distribution moments, which became unstable in segments with limited vehicle passes or atypical geometry. These instabilities led to spurious flags in locations where traffic was normal but sample sizes were small. The team shifted to a more robust representation based on empirical cumulative distribution functions (CDFs) of speed and lateral position and used Kolmogorov-Smirnov statistics to quantify weekly distributional shifts. Multi-week aggregation was introduced to stabilize features, and lateral movement was constrained to realistic bounds derived from cross-section geometry. These changes substantially reduced false positives and yielded a more reliable closure detection pipeline suitable for statewide deployment.

Another critical challenge was to optimize the cloud compute strategy and compute costs. Early prototypes relied on AWS Athena queries over large CV tables, which often required 8 to 10 minutes to process the most recent batches of data because of repeated full-table scans. Migrating core transformations into AWS Glue improved performance but introduced a different problem: cost. With two to three Glue jobs running continuously around the clock to keep statewide metrics updated, monthly compute charges grew quickly and were not sustainable for long-term, always-on operations. To address this, the team transitioned the majority of heavy processing to AWS EC2-based Spark clusters. This approach reduced latency to roughly 5 to 10 seconds for processing statewide five-minute windows and lowered monthly compute costs on the order of 40% to 60% relative to a Glue-based configuration. The move to managed Spark clusters also provided finer control over scaling, allowing the system to scale up during peak periods and downscale overnight or during low-demand periods.

Ensuring end-to-end reliability of the real-time pipeline required additional operational refinements. Some failures arose when AWS Lambda-based transformation services attempted to consume files that had already been deleted by cleanup processes or overwritten by newer outputs. To resolve this, dual-write strategies were adopted so that intermediate products were stored in both transient and longer-lived Amazon S3 locations, retention windows were extended, and pipeline stages were synchronized using S3 event triggers rather than fixed time offsets. On the front end, dashboard responsiveness and map rendering performance initially suffered when thousands of CV points and polygons were drawn at full resolution. This was mitigated by introducing map tile caching, asynchronous data loading, and zoom-dependent filtering so that the dashboard displays only the level of detail appropriate for the current view. These enhancements significantly improved user experience and made statewide visualization practical in operational settings.

1.3. Successes Achieved

Despite the complexity and scale of the problem, the project achieved several significant technical and operational milestones that collectively demonstrate the feasibility and value of a CV-based statewide work zone intelligence system. One of the most important accomplishments was the deployment of a fully automated real-time pipeline capable of ingesting, processing, and analyzing CV data on a continuous basis. The system now routinely processes CV telemetry on the order of every 30 seconds, aggregates it into five-minute bins, and generates statewide traffic performance metrics that capture queue formation, lane closures, and sudden speed drops with minimal latency. This performance lays a strong foundation for the short-term detection and prediction of emerging disruptions, fulfilling a core element of the project scope.

The project also delivered a unified operational dashboard that integrates these analytics into a single, cohesive interface for transportation operators. Users can view statewide work zones on an interactive map, inspect real-time speed heatmaps and queue metrics, review time series plots for speed and delay, and examine high-resolution trajectory plots that reveal microscopic driving behavior around work zones. This combination of spatial and temporal visualization offers a level of continuity and detail that legacy systems based on isolated detectors and manual logs cannot provide. The dashboard further exposes lane closure detection overlays and, when available, live camera feeds for visual verification, although cameras are not required for core computation.

A particularly impactful success was the accurate detection of lane closures using only CV data. Traditional lane closure reporting methods such as manual field logs, project databases, and sensor-based systems are often inconsistent, incomplete, or delayed, which makes it difficult to maintain reliable situational awareness. Prior research has noted that these conventional sources can be error prone, inefficient, and limited in their ability to capture dynamic or moving operations, especially when closures shift over time or when no dedicated detection devices are present [1]. By continuously analyzing CV lateral position and speed distributions at a 0.1-mile resolution, ReactorIQ's lane closure module provides an independent behavior-based source of truth that identifies when and where closures are actually in effect. This improves day-to-day situational awareness and offers a more accurate basis for evaluating work zone performance and contractor compliance over time.

Another key milestone was the integration of real-time queue warning and safety-related alerts. The system generates Traffic Management Data Dictionary (TMDD)-compliant XML messages describing the onset and extent of queues, as well as severe slowdowns in and around work zones. These messages are designed to feed directly into agency systems for DMS activation and traveler information dissemination. By closing the loop between detection and communication, the platform moves beyond passive monitoring and supports active, automated responses to emerging traffic conditions.

The project further demonstrated the diagnostic value of trajectory plots and speed heatmaps for operational decision-making. These visual tools have already been used to identify shockwave propagation patterns, evaluate the impacts of different taper designs, and analyze merging turbulence caused by lane reductions. In several case studies, they provided insight into the exact segment where breakdown occurred, how far and how fast the resulting queues propagated upstream, and how long it took for traffic to recover after work activities ended. This level of diagnostic detail directly supports both immediate operational decisions and longer-term work zone design improvements.

Finally, the architecture itself represents a major success. By shifting heavy computation from Glue to EC2-based Spark clusters and designing modular services for metric computation, alert generation, and visualization, the system achieves statewide scalability while maintaining low latency and reasonable operating costs. The modular design also makes it feasible to extend ReactorIQ to additional use cases, such as non-work-zone incidents or weather-related disruptions, without requiring a complete redesign.

1.4. Future Improvements

A major area of future development involves deeper fusion of CV telemetry data with camera-based analytics. While camera feeds are currently used primarily for visual verification, the longer-term goal is to deploy AI-based video models capable of identifying stalled vehicles, estimating lane occupancy, and detecting near-miss interactions in and around work zones. These visual analytics, when combined with CV-based speed and trajectory patterns, could significantly improve the detection of complex incidents and subtle safety issues that may not be evident from CV telemetry data alone, particularly in areas with lower penetration rates.

The lane closure detection framework itself can also be extended. Current models are designed to detect single-lane closures and relatively stable configurations. Future work aims to handle dynamically moving closures, partial shoulder closures, multi-stage construction layouts, and highly variable nighttime operations. Achieving this will require more granular modeling of lateral position distributions, time-of-day effects, and lane-specific geometric constraints, as well as tighter integration between automatic closure detection and planned closure records.

On the algorithmic side, there is room to refine trajectory smoothing, enhance shockwave detection algorithms, and formalize turbulence classification into clear operational categories such as recurrent stop-and-go waves, isolated breakdowns, and merge-related disturbances. These improvements would help operators quickly interpret complex patterns and prioritize responses. From an infrastructure standpoint, continued work is planned on strengthening data redundancy, incorporating message queues for more resilient inter-service communication, and further optimizing storage and compute costs. Additional use of EC2 autoscaling policies and spot instances may provide further cost savings without compromising reliability.

Finally, several user-facing enhancements are envisioned for the dashboard. These include multi–work zone comparison views that allow operators to assess performance across corridors, historical playback modes for replaying past days of operation, and automated daily or weekly performance reports that summarize key metrics and anomalies. Together, these improvements would make ReactorIQ even more useful as both a real-time operations tool and a planning and evaluation resource.

1.5. Conclusion

This project represents a significant advancement in real-time traffic monitoring and work zone management by demonstrating that CV data, when combined with cloud-native processing and tailored analytics, can serve as the foundation for a comprehensive statewide intelligence platform. Through an iterative development process, the team overcame challenges related to computational latency, cost management, and model stability. The resulting system provides transportation agencies with powerful tools for detecting congestion, accurately identifying lane closures, monitoring safety-related behavior, and understanding traffic dynamics at both the macroscopic and microscopic levels.

The successful deployment of ReactorIQ confirms the viability and importance of CV-driven traffic intelligence systems. By supporting short-horizon detection and prediction, real-time operational monitoring, and long-term smart work zone planning within a single architecture, the platform aligns closely with the scope defined in the original proposal for this project. With continued enhancements, ReactorIQ can serve as a foundation for statewide automated operations, predictive analytics, and integrated work zone management, laying the groundwork for the next generation of smart transportation infrastructure in Iowa and beyond.

1.6. Report Outline

The remainder of this report presents contextual information on advanced work zone monitoring and management (Chapter 2) and details each component of Reactor IQ: real-time incident detection (Chapter 3), performance monitoring and metric design (Chapter 4), smart planning and lane closure detection (Chapter 5), and the deployment of the integrated dashboard (Chapter 6). The closing chapter (Chapter 7) summarizes the key strengths and limitations of ReactorIQ and outlines potential improvements for the platform.

CHAPTER 2. BACKGROUND

2.1. Motivation

Work zones have emerged as increasingly important components of modern roadway infrastructure management, especially as the United States continues its efforts to rehabilitate aging highways and bridges under recent federal investment programs such as the Infrastructure Investment and Jobs Act. These temporary but essential work zones support long-term mobility and asset preservation, yet they introduce operational challenges that can significantly disrupt normal traffic flow. Drivers entering work zones encounter altered roadway conditions such as lane shifts, shoulder narrowings, uneven pavement surfaces, tapered lane drops, and the ingress and egress of construction vehicles. Each of these elements influences driver behavior and can lead to speed variability, frequent lane changes, and reduced stability in traffic flow. If not monitored and managed carefully, these disruptions can escalate into safety and mobility problems.

Recent national statistics from the Federal Highway Administration (FHWA) indicate that work zones continue to account for a substantial number of traffic fatalities each year, with several hundred to nearly one thousand deaths annually in the early 2020s [2]. Many of these crashes involve rear-end collisions and occur in the presence of sudden speed reductions and queue formation near lane drops and constrained geometries [3]. As metropolitan areas experience rising travel demand, many freeway corridors already operate near their effective capacity during peak periods. When a work zone is introduced, even modest reductions in available roadway capacity can sharply increase the likelihood of breakdown and delay. Prior research has shown that work zone capacity is sensitive to lane closure configuration, geometric design, driver familiarity, and the distribution of speeds and densities in the approach region [4,5,6]. Under these conditions, relatively small disturbances such as merging maneuvers or brief slowdowns near a taper can propagate upstream as shockwaves and evolve into sustained queues. Classical input-output style queueing models and hybrid macroscopic formulations further demonstrate that small differences between arrival and departure rates at bottlenecks can translate into significant delays and growing queue lengths at signalized and lane-restricted locations [7,8,9].

CV and probe data studies have documented important limitations of traditional fixed infrastructure sensing when used to monitor congestion and queues. Evaluations of wide-area probe data and loop detector systems have identified issues such as speed bias, incomplete spatial coverage, and sensor outages that can obscure the true extent of congestion near bottlenecks, including work zones [10–14]. Analyses of commercial probe products such as INRIX have highlighted both the opportunities and the challenges associated with using third-party data streams for real-time performance monitoring and historical trend assessment, emphasizing the importance of understanding sampling variability and latency when interpreting congestion patterns [14,15,16]. These findings motivate closer integration between CV data and agency monitoring systems so that queues can be followed as they form, grow, and dissipate along the corridor instead of being inferred only at isolated sensor locations.

Historically, the management of work zones has relied on fixed-point sensors and human observation, both of which face intrinsic limitations in environments with dynamic boundaries and moving bottlenecks. Fixed-sensor infrastructure such as loop detectors and side-fired radar units provides detailed measurements of speed, volume, and occupancy at the instrumented locations, but their spatial granularity is inherently limited to the specific points where sensors are installed. In many freeway systems, detectors are spaced on the order of one half-mile to one mile in urban areas, with even wider spacing on rural corridors, which can leave gaps between measurement locations [14,17,18]. When work zones are established or adjusted, the active disruption area may shift upstream or downstream of existing sensor sites, leading to blind spots that reduce situational awareness [17,19]. In addition, detector health is a persistent concern. Agency-level studies have reported nontrivial proportions of detector stations with communication failures, stuck values, or other malfunctions at any given time, reinforcing the need for continuous sensor health monitoring and redundancy in performance measurement systems [13,20].

Closed-circuit television (CCTV) cameras provide important visual context, particularly when staff at traffic management centers (TMCs) actively monitor feeds. However, camera effectiveness can degrade under low-visibility conditions such as nighttime operations, fog, snow, heavy rain, or glare. Even when video quality is adequate, the utility of CCTV depends on sustained human attention, which is difficult to maintain across numerous concurrent work zones. Recent advances in computer vision detectors, including architectures related to YOLO, have improved the automation of vehicle and object detection, but performance can still be sensitive to factors such as distance, scale, and adverse environmental conditions, and most existing efforts have not been tailored specifically to the full range of work zone scenarios [21–24]. As a result, camera-based methods often complement rather than replace other data sources for continuous, network-wide monitoring.

Temporary sensors such as portable radar trailers, temporary loops, and Bluetooth readers offer additional flexibility for long-duration projects, yet their deployment requires labor, calibration, and ongoing maintenance. These devices must be placed within or near the work zone, which introduces exposure to construction activity, power interruptions, and occasional misalignment or damage. For short-duration activities such as grinding operations, mobile striping, pothole repair, or moving maintenance convoys, the time and effort required to install temporary sensing equipment often outweighs the expected benefit, so many such lane closures receive little or no direct fixed-sensor coverage [25,26,27]. Prior work on portable and connected work zone devices has sought to improve the availability and standardization of work zone data feeds, but these systems are not yet universal and still require consistent field deployment and maintenance [27–30].

Human-based reporting remains an important component of agency workflows but introduces variability and potential delay. Field personnel typically report observations intermittently, and the onset of slowdowns or queues may not be recognized until conditions become pronounced. Communication between field staff and TMCs can also lag during busy periods, and subjective judgment can lead to differences in how severity and impact are classified [16,27,31,32]. These factors can limit the timeliness and consistency of information available for both real-time operations and retrospective evaluations.

Taken together, these limitations indicate that traditional fixed and temporary monitoring systems alone are often insufficient to provide the continuous, spatially rich, and adaptive coverage needed for complex work zone operations. The variability in work zone layouts, the presence of recurrent and nonrecurrent congestion, and the sensitivity of driver behavior to geometric and control changes all motivate the use of CV data and advanced analytical tools that can track traffic states along the roadway with higher resolution than point detectors can offer. Recent research has demonstrated the potential of large-scale CV data for real-time traffic applications, including queue detection, safety analytics, and work zone performance assessment [10,18,23,33,34,35]. These developments provide a foundation for integrated platforms that combine CV data with traditional sensors, probe products, and emerging field devices to support more robust and proactive work zone management.

2.2. Emergence and Importance of Connected Vehicle Data

CV technologies have expanded rapidly in recent years, resulting in greater availability of anonymized location, speed, and heading information from passenger vehicles and commercial fleets. Automakers and telematics providers now supply continuous probe data streams through platforms such as GM OnStar, Ford Safety Insights, Audi Traffic Light Information, and commercial aggregators including Wejo and INRIX. These data sources have enabled state transportation agencies to incorporate CV information into their traffic monitoring workflows, particularly for applications that benefit from wide spatial coverage and high temporal resolution [36,37,38]. Although penetration rates vary by region and roadway type, prior studies have shown that even moderate CV availability can provide meaningful insight into corridor performance and traffic behavior.

CV telemetry data offer several advantages for work zone monitoring. Unlike fixed detectors that capture traffic conditions at a single point, CV probes act as mobile sensors that traverse the entire work zone footprint. This allows analysts to observe speed changes, deceleration waves, and the onset of congestion along the full work zone length rather than only at detector locations. Because each vehicle contributes a time-ordered sequence of GPS observations, CV trajectories support analysis of how drivers respond to lane shifts, tapers, and other temporary traffic control elements. Prior work has demonstrated that trajectory information can complement traditional infrastructure by capturing queue dynamics and bottleneck formation in real-world conditions [33,39].

Research has shown that CV-based queue detection methods can perform reliably when a sufficient number of probes are present, with several studies reporting strong agreement between CV-derived queue estimates and those obtained from reference sensors or manual observations [33,40]. CV telemetry data are also less sensitive to factors that can degrade camera or detector performance, such as nighttime lighting, weather conditions, or temporary equipment failures. This robustness makes them particularly useful for monitoring work zones where traditional sensors may have limited spatial placement or operational uptime [9,41,42].

Overall, the combination of wide-area coverage, continuous sampling, and driver behavior insights positions CV telemetry data as a valuable resource for modern work zone performance management. When integrated with cloud-based processing and analytics, CV telemetry data support applications such as identifying speed reductions, monitoring queue formation and dissipation, estimating travel time and delay, and generating performance measures for agency decision-making [15,33,43]. The following sections describe how these data streams are processed and how artificial intelligence techniques can further enhance work zone monitoring and prediction.

2.3. Need for Advanced Analytics and AI-Enabled Interpretation

The availability of CV telemetry has created new opportunities for monitoring work zone operations, but the raw data streams alone are not directly suitable for operational decision-making. CV telemetry data are generated at high frequency and large scale, often producing detailed latitude-longitude coordinates, timestamps, speeds, and headings for millions of records across a statewide network. These observations must be spatially aligned to roadway networks, filtered for noise, and aggregated in ways that preserve the underlying traffic dynamics. CV telemetry datasets are also inherently sparse within the traffic stream and exhibit variable penetration rates over time and space, which requires analytical procedures that are robust to uneven sampling. Extracting meaningful performance measures from such datasets therefore requires analytical approaches capable of recognizing patterns that are not readily apparent through simple thresholding or point-level speed summaries. As a result, the integration of traditional traffic engineering principles with modern analytical and AI-based methods has become increasingly important.

AI-enabled analytics provide a structured way to interpret spatial and temporal patterns within CV trajectories. For example, sustained lateral shifts in trajectory paths, reductions in lane-specific speeds, or localized clusters of deceleration can indicate lane restrictions or active merging behavior within a work zone. Prior work has shown that trajectory distributions and density patterns can reveal operational conditions such as queue buildup, merging turbulence, and disruptions associated with temporary traffic control [33,39]. Unsupervised learning approaches, including clustering and segmentation techniques, have been used to group trajectories exhibiting similar behaviors, enabling identification of unusual movements near taper areas or separation of normal flow patterns from anomalous driving behavior. These analytical procedures complement traditional work zone monitoring by providing lane-level and segment-level behavioral insights that are not available from point detectors alone.

Queue detection has also benefited from advanced analytical techniques applied to CV telemetry data. Traditional queue identification relies heavily on fixed-detector occupancy or aggregated speeds, which may not capture short-lived disturbances or localized congestion. In contrast, CV trajectories provide continuous measurements of speed, acceleration, and spacing along the roadway, enabling more detailed characterization of queue onset, propagation, and dissipation. Studies have demonstrated that speed variance, deceleration patterns, and changes in trajectory spacing can serve as reliable indicators of queue fronts and stop-and-go conditions [33,40]. Other queue length and delay estimation frameworks developed for signalized and unsignalized locations show that CV telemetry data can be incorporated into input-output models and hybrid estimation techniques to quantify queue dynamics when traditional detectors are unavailable [8].

As transportation agencies expand their use of CV telemetry data, the role of advanced analytics and AI-based interpretation is becoming increasingly central. Without such tools, agencies would face significant challenges in detecting emerging congestion, identifying lane restrictions, estimating queue length and delay, or generating accurate flow and density measures from large-scale probe datasets. AI-enabled analytics help automate these tasks, reduce manual processing burdens, and support timely operational decision-making. This combination of data-driven analytical methods and established traffic engineering practices forms the foundation for scalable and resilient work zone monitoring systems [15,33,41,43].

2.4. Cloud-Native Processing Framework for Statewide Scalability

Implementing CV analytics at statewide scale requires a processing architecture capable of ingesting, filtering, and summarizing large volumes of data with minimal latency. Cloud-native platforms offer the computational flexibility and distributed processing capacity necessary to support real-time work zone monitoring across large geographic regions. By combining cloud services with distributed frameworks such as Apache Spark and Apache Sedona, agencies can efficiently perform geospatial filtering, segmentation, and aggregation on CV telemetry datasets [45].

A key advantage of cloud-native architectures is their elasticity. CV telemetry data volume varies substantially across corridors, construction seasons, and times of day, and cloud platforms can scale computational resources up or down in response to these fluctuations. During periods of elevated traffic or intensive construction activity, clusters can expand to maintain timely updates; during low-volume periods, they can scale down to reduce cost. Such dynamic resource allocation aligns with best practices recommended in recent cloud-based intelligent transportation systems (ITS) and transportation cyberinfrastructure studies that emphasize scalable and resilient architectures for big data analytics [13,46].

Geospatial processing is one of the most computationally intensive aspects of CV telemetry data analysis, particularly when mapping thousands of points to numerous work zone polygons or roadway geometries. Apache Sedona supports distributed spatial queries such as point-in-polygon filtering, trajectory mapping, and distance projection. Work zone boundaries can be stored as geometric polygons, and each CV point can be assigned to the appropriate roadway segment, enabling fine-grained computation of speed, density, and queue conditions. Prior work in large-scale traffic surveillance and sensor health monitoring highlights the importance of distributed spatial engines for enabling near-real-time analysis of transportation datasets [42,13].

Temporal aggregation also plays an important role in processing CV telemetry data. Many agencies rely on five-minute intervals to balance responsiveness with stability in performance measures. Aggregating CV points into short time windows supports consistent computation of average speeds, penetration rates, and other operational metrics. Studies examining probe data aggregation for freeway and arterial performance reporting support these short interval windows as appropriate for real-time operational monitoring [47].

Cloud-native designs additionally provide fault tolerance and operational robustness. Distributing processing tasks across multiple nodes mitigates the effects of individual node failures, and cloud storage services offer redundancy and high availability. This ensures consistent access to CV telemetry data streams, which is particularly important during active construction, when congestion and queue formation may change rapidly.

Integrating CV telemetry data with cloud-native architectures provides agencies with a scalable and flexible foundation for monitoring work zones statewide. New or modified work zones can be incorporated simply by updating geometric boundaries, enabling rapid deployment without field hardware. These capabilities build on prior work in transportation cyberinfrastructure and statewide work zone databases that emphasize standardized, scalable data integration frameworks [46,48].

2.5. Importance of Real-Time Performance Metrics for Work Zone Operations

Real-time performance metrics are essential for effective work zone management because they provide timely insight into developing congestion, mobility impacts, and potential safety concerns. Without real-time visibility, agencies must rely on delayed or reactive information sources such as crash notifications, motorist complaints, or intermittent field observations. This limits their ability to adjust traffic control strategies, communicate delays to travelers, or mitigate emerging queues before they affect upstream interchanges.

Key performance measures include queue length, queue duration, speed distributions, and vehicle delay. Queue length indicates how far congestion extends upstream of the work zone and whether it threatens nearby access points. Queue duration reflects the persistence of bottlenecks. Speed distributions reveal turbulence, instability, and driver behavior changes in response to lane reductions. Delay quantifies the additional travel time imposed on users and is widely used in performance-based evaluation frameworks. Classical input-output approaches and model-based estimation methods, originally developed for signalized intersections, provide a useful foundation for estimating queue length and delay and have been adapted for CV-based analyses [7,8,9].

Recent studies have shown that CV telemetry data can be used to estimate these performance metrics with reasonable accuracy. Trajectory data provide detailed information on speed reductions, stop-go patterns, and the spatial extent of congestion. Several empirical evaluations have demonstrated strong agreement between CV-derived queue measures and ground truth or sensor-based observations when sufficient probe data are available [10,15]. CV trajectories have also been used to analyze speed variability and driving behavior inside temporary lane reductions, offering insights into operational conditions and potential safety risks [49].

Real-time performance metrics also support automated warning systems. When speed drops or queue formation is detected upstream of a work zone, agencies can activate DMS to warn drivers of sudden slowdowns. Studies evaluating these systems have reported reductions in congestion-related safety risks when warnings are triggered based on real-time traffic information [42,50].

Integrating probe and CV telemetry data into these systems improves the timeliness and reliability of activation logic.

More broadly, standardized and consistent real-time metrics across all active work zones enable state DOTs to evaluate contractor compliance, adjust work hours, optimize staging decisions, and allocate operations resources effectively. As agencies increasingly adopt performance-based management practices, reliable real-time indicators derived from CV telemetry data become critical for both day-to-day operations and long-term planning [31,47,48].

2.6. Integrating AI-Driven Analytics Into Agency Workflows

The usefulness of an advanced work zone monitoring system depends not only on the quality of its data and analytics but also on how effectively the resulting information integrates with existing agency workflows. Transportation agencies operate within established operational structures that include TMCs, maintenance of traffic (MOT) plans, DMS systems, 511 traveler information platforms, and communication channels with field personnel and contractors. For AI-driven CV analytics to be operationally impactful, they must be delivered in formats that align with these workflows and support decision-making without introducing additional complexity.

Visual dashboards play a central role in translating large volumes of CV trajectories into insights that operators can interpret quickly. Time-space diagrams, speed profiles, queue visualizations, and geospatial heatmaps allow staff to recognize developing congestion patterns and operational changes at a glance. Prior work in traffic operations visualization shows that trajectory- and density-based displays derived from probe and CV telemetry data can improve operators' ability to detect speed drops, instability, and variability compared with traditional point-sensor displays. Within work zones, these visualizations help operators understand how traffic responds to lane reductions, taper layouts, and construction staging. Experiences from statewide probe-data platforms and traffic operations open data services further demonstrate that such visual tools can support both real-time decision-making and post-event reviews across multiple agency units [46,47].

AI-enabled analytics can also assist in automating routine operational tasks. Queue detection algorithms, for example, can generate alerts when speeds fall below expected thresholds or when spatial patterns suggest that congestion is forming. These alerts may be incorporated into DMS workflows through the forwarding of recommended messages to operators for review. Studies on automated DMS activation and probe-based surveillance systems indicate that integrating real-time CV indicators into operational workflows can improve timeliness and consistency in messaging [42]. Such automation does not replace operator judgment, but it provides a structured decision support layer that reduces manual monitoring workloads.

AI analytics may also support verification of MOT plans. Work zones frequently undergo changes due to staging adjustments or temporary field conditions, causing deviations between planned and actual lane configurations. Trajectory-based analytical methods can help identify patterns that reflect lane closures or shifts in lane availability. When discrepancies are detected, such systems can notify operators or inspectors, providing an additional layer of quality assurance. Similar approaches have been explored in the broader context of probe data–based traffic surveillance and operational validation [42].

Nighttime operations represent another workflow area where CV analytics add value. Lighting and visibility limitations often reduce the effectiveness of traditional field-based monitoring tools, whereas CV telemetry data provide consistent information regardless of light conditions. Research examining day-night differences in driving behavior has noted greater variability in nighttime speeds and lateral positioning. Incorporating these insights into agency workflows can help inform targeted nighttime strategies such as enhanced lighting or modified MOT layouts.

Overall, integrating AI-driven analytics into agency workflows supports a transition from reactive to proactive operations. Rather than responding only after congestion or incidents occur, agencies can act when early indicators suggest developing instability. This shift highlights the importance of delivering analytics not as isolated metrics but as actionable intelligence embedded within daily operational practices [33,42,47,48].

2.7. Statewide Scalability and Operational Feasibility

A major advantage of CV analytics is their ability to scale across an entire roadway network without requiring extensive new infrastructure. This scalability improves the feasibility of comprehensive work zone monitoring by enabling agencies to assess both long-duration projects and short-duration or mobile operations that have traditionally lacked monitoring coverage.

The cloud-native processing framework described earlier allows large volumes of CV telemetry data to be ingested and analyzed for many active work zones simultaneously. Because CV telemetry spans nearly all freeway segments, monitoring a new work zone typically requires only adding its polygon boundary to the system's configuration, eliminating the need for field hardware deployment. This capability allows agencies to observe planned and unplanned lane closures with minimal lead time.

Statewide scalability also promotes consistency in performance measurement. When all work zones are evaluated using the same CV-based metrics, such as queue length, delay, and speed degradation, agencies can compare performance across regions, contractors, and work zone types. Guidance published by the Transportation Research Board (TRB) on performance-based work zone management emphasizes the value of standardized, repeatable metrics for supporting equitable evaluation and planning [31]. Prior research on probe and CV telemetry data reliability further supports the feasibility of generating consistent performance indicators on a wide-area basis [10,12,34]. Short-duration mobile work zones, such as pavement operations or striping activities, particularly benefit from CV monitoring. Traditional fixed detectors cannot track the moving bottleneck associated with mobile operations, but CV trajectories naturally capture these movements.

Studies examining probe-based speed patterns have demonstrated that CV telemetry data can reveal the temporal and spatial footprint of mobile operations, allowing agencies to infer the evolving location of the work activity [42,48]. These findings suggest broad potential for CV-based monitoring of operations that have historically fallen outside the scope of traditional instrumentation.

Another dimension of statewide feasibility relates to cost-effectiveness. Fixed roadside detectors require installation, maintenance, and periodic recalibration, whereas CV analytics leverage data generated by vehicles already in the traffic stream. Several agencies have reported that CV-based monitoring reduces the need for additional field hardware and lowers long-term maintenance burdens [51]. Complementary studies examining commercial probe and crowdsourced data products highlight their usefulness and limitations, reinforcing the value of combining multiple wide-area datasets with CV telemetry [11,12,52].

As CV penetration continues to increase, the reliability and resolution of statewide analytics will improve organically. The expansion of automaker data partnerships and greater adoption of telematics systems are expected to increase sample sizes over time, enhancing the long-term utility of CV-based performance monitoring. These trends strengthen the operational feasibility of systems that rely on CV telemetry data for real-time insights [33,39,47].

In summary, statewide scalability is a defining strength of CV analytics. By enabling broad monitoring coverage, standardized metrics, and compatibility with mobile and short-duration operations, CV telemetry data provide a foundation for transforming work zone management into a proactive, data-driven statewide practice [33,46,47,48].

2.8. Enhancing Safety Through Proactive Detection and Response

Improving work zone safety remains a central priority for transportation agencies, as work zones introduce abrupt changes in roadway geometry and traffic flow that increase the likelihood of sudden braking, rapid lane changes, and speed differentials. Rear-end collisions make up a substantial portion of work zone crashes, particularly in transition areas such as tapers and merge points, where driver expectations differ from free-flow conditions. Proactive detection of emerging congestion and queue formation is therefore essential for reducing crash risk and improving overall safety.

CV data support safety improvements by providing early indications of traffic instability. When speeds begin to decline upstream of a work zone, CV trajectories capture these changes immediately, offering earlier visibility than many fixed detectors. Prior work has demonstrated that CV-based metrics can identify the onset of congestion, speed reductions, and stop-and-go traffic patterns that signal increased safety risk [33]. These early indicators can support operational decisions such as activating warning messages or deploying enforcement resources.

In addition to detecting queues, CV telemetry data reveal traffic turbulence characterized by high speed variance, repeated acceleration and deceleration, and irregular lane changing behavior.

These patterns may indicate elevated crash potential even when full congestion has not developed. Work on safety surrogate measures in work zones has shown that trajectory data can provide information about conflict risk, especially where DMS and temporary traffic control devices influence driver behavior [53]. Similarly, analyses of driver behavior at the back of the queue have highlighted how CV trajectories capture interactions between approaching vehicles and slowing traffic, supporting improved design and placement of queue warning systems [33].

Real-time traffic state analysis also contributes to safety management. Elevated density levels within constrained work zone segments often signal unstable flow conditions that increase susceptibility to shockwaves and abrupt speed drops. Studies using CV telemetry data for monitoring turbulence and acceleration events have shown that CV-derived deceleration patterns can serve as practical indicators of emerging safety concerns [39]. Related research that integrates CV speeds with incident or harsh braking data provides further evidence that trajectory-based metrics can support safety assessments across a range of operating conditions [54].

Nighttime operations represent another domain where CV telemetry data enhance safety. Crash rates in work zones are often higher at night due to reduced visibility, yet many traditional sensors provide limited information during nighttime hours. CV telemetry data, unaffected by lighting conditions, offer consistent insight into speed patterns, lane positioning, and merging behavior, although, depending on the penetration rates, this might only be useful for work zones with high volumes during nighttime. Prior analyses have documented meaningful differences between daytime and nighttime driving behavior, including increased speed variability and different merging dynamics. These insights can inform nighttime strategies such as enhanced lighting, reflectorized channelization, or modified closure layouts. Additional work has shown that CV-based speed patterns can support safety decisions during winter operations, which often overlap with work zone activity [41].

Taken together, the early detection capabilities and behavioral insights provided by CV telemetry data enable agencies to adopt a more proactive approach to work zone safety. Rather than responding only after crashes occur, agencies can identify developing risk, implement countermeasures early, and continuously evaluate the impact of interventions [33,39,41,53,55].

2.9. Supporting Performance-Based Work Zone Management

The shift toward performance-based work zone management represents an important evolution in transportation operations, emphasizing measurable outcomes over prescriptive requirements. Agencies increasingly rely on objective performance indicators to evaluate the effectiveness of lane closures, construction staging, and contractor compliance. This approach aligns with national guidance encouraging data-driven evaluation of mobility and safety impacts across active work zones [31]. CV data provide a strong foundation for these frameworks because they offer consistent, repeatable metrics that can be applied across diverse work zone types and locations.

Metrics such as queue length, delay, speed degradation, and turbulence allow agencies to compare work zone performance across corridors, contractor teams, and time periods. These comparisons support identification of recurring bottlenecks, opportunities to refine traffic control plans, and potential adjustments to staging schedules. CV-based performance evaluation builds on earlier work using model-based queue and delay estimation methods, which emphasize systematic indicators for operational assessment [7,8,9]. CV telemetry data extend these principles to freeway and high-speed environments by providing continuous, corridor-wide coverage.

Post-operation analysis is another essential component of performance-based work zone management. After a project concludes, agencies can review historical CV-derived metrics to assess overall operational performance. Such retrospective evaluations help determine whether closure durations were appropriate, whether taper lengths were effective, and whether the implemented MOT plans resulted in avoidable congestion. TRB guidance highlights the value of these reviews for improving future project planning and strengthening predictive modeling frameworks [31]. Large-scale evaluations of CV and probe data reliability and congestion detection accuracy provide additional context for interpreting performance metrics in statewide applications [10,12,34].

Performance-based management also supports improved communication and transparency. Agencies can publish objective summaries derived from CV telemetry data to inform stakeholders and the public about mobility and safety conditions during construction. These summaries help build trust, justify operational decisions, and support evaluation of alternative traffic control strategies. As agencies continue to adopt standardized performance indicators, CV telemetry data are expected to play an increasingly central role, especially when integrated with research-grade work zone databases and open data platforms that consolidate performance information [46,47,48].

2.10. Vision for Future Expansion and Integration

The AI-powered CV analytics platform presented in this report establishes a foundation for continued expansion as CV penetration increases and analytical methods evolve. Advancements in machine learning and the growing availability of high-resolution trajectory data offer opportunities to enhance predictive capabilities, integrate additional data sources, and support more automated operational workflows.

One promising direction for future development is real-time incident detection modeling. While the current system focuses primarily on real-time detection, machine learning models such as recurrent neural networks and temporal graph-based methods have shown potential for forecasting near-term traffic conditions on freeways and arterials [56]. Applying these approaches to work zones could allow agencies to anticipate queue formation or speed degradation and proactively adjust lane closure timing or staging strategies.

Another opportunity is deeper integration with incident detection frameworks. Sudden localized speed reductions or irregular lateral movements in CV trajectories may indicate crashes or other disruptions. Work on real-time crash detection and harsh braking analysis using CV telemetry data demonstrates the feasibility of identifying such events through trajectory features and acceleration patterns [39,55]. Incorporating these capabilities into work zone monitoring would create a unified analytic environment in which both mobility- and safety-related anomalies are identified.

Automated verification of work zone configurations also represents a potential enhancement. MOT plans specify expected lane closures and taper locations; however, field conditions may deviate from planned layouts. Trajectory-based analytics could be used to detect changes in lane availability or merging behavior that signal differences between planned and actual conditions. These concepts align with broader research on probe-based operational monitoring and spatial consistency verification in traffic surveillance systems [42].

Integration with traveler information systems is another area for expansion. Real-time metrics such as delay, travel time, and queue length could be incorporated into 511 systems, navigation applications, and advanced traveler information signage. Prior work on open data services and probe-based performance reporting platforms demonstrates the feasibility of disseminating such metrics through public APIs and operator dashboards [47,46].

Future systems may also incorporate multimodal data fusion. Combining CV telemetry data with probe truck telematics, crowdsourced incident reports, Bluetooth tracking, or machine vision outputs could improve the completeness and reliability of monitoring. Existing evaluations of wide-area probe datasets and incident reporting platforms provide guidance on how multiple data sources may be compared or fused to enhance situational awareness [10,52,15,12].

Ultimately, the long-term vision is a fully integrated, AI-enabled work zone monitoring ecosystem that supports predictive analytics, multimodal data fusion, automated verification, and the seamless dissemination of information to both operators and road users. The CV platform described in this chapter provides a scalable and robust foundation for realizing this vision [33, 47,42,46,48].

2.11. Chapter Summary

This chapter provided a comprehensive overview of the motivation, challenges, and analytical opportunities associated with AI-powered CV data for work zone management. The discussion began by highlighting the increasing prevalence of work zones nationwide and the associated operational and safety challenges, including sudden speed changes, merging friction, and dynamic bottlenecks. Traditional monitoring tools such as fixed detectors, CCTV cameras, portable sensors, and field observations were shown to have limitations in capturing the full spatial and temporal variability of work zone conditions [17,20,21,25].

CV telemetry data have emerged as a powerful complement to traditional sensing technologies, offering continuous, segment-wide visibility into driver behavior and traffic dynamics. With CV penetration increasing in recent years, these data now support real-time monitoring of speed, queue formation, density, and turbulence across large roadway networks [36,37,38]. The chapter described how AI-driven analytics transform raw CV records into actionable performance metrics, drawing on techniques for identifying lane closures, queues, and speed variations, as well as model-based estimation frameworks for computing queue length and delay [7,8,9,4,5,50,53,33,34].

Cloud-native processing architectures were shown to enable statewide scalability by supporting distributed geospatial filtering, segmentation, and temporal aggregation of millions of CV points in near real-time [45,57,47,13,46,48]. The chapter also emphasized the importance of real-time performance metrics for agency decision-making and their integration into emerging performance-based management frameworks [58,49,26,31].

Subsequent sections described how AI-enabled dashboards, automated alerts, and MOT verification tools fit within existing agency workflows, improving situational awareness and operational efficiency [50,53,42]. Finally, the chapter highlighted opportunities for future expansion, including predictive modeling, multimodal data integration, incident detection, and expanded dissemination of work zone performance metrics [56,47,39,55,46,48].

Together, these elements establish the foundation for the chapters that follow, which detail the system architecture, performance metric computation, performance estimation methodologies, and the integrated dashboard that supports statewide work zone monitoring.

CHAPTER 3. REAL-TIME INCIDENT DETECTION

3.1. Introduction

Highway work zones have become increasingly complex operational environments as agencies respond to aging infrastructure, increased freight volumes, and rising traffic demand. Temporary lane closures, lane shifts, shoulder restrictions, and reduced speed zones change roadway geometry and operating conditions, which influences driver behavior and traffic stability. Studies have shown that drivers approaching work zones often encounter turbulence in the traffic stream, irregular deceleration, frequent lane changes near tapers, and sudden queue formation during peak periods [27,60,61]. When demand exceeds the reduced work zone capacity, queues can develop upstream of the taper and create safety concerns for drivers, who face abrupt speed drops.

Queue buildup near work zones is one of the most important operational hazards. Free-flowing vehicles can encounter slow-moving or stopped queues with very large speed differences, which leaves limited time for reaction and increases the risk of rear-end and secondary crashes [2,3,62]. Many work zone crashes have been reported in the queue tail or transition zone just upstream of the lane closure, and these are often severe because of the speed reduction involved [2,3]. Safety risks are higher at night or during adverse weather when visibility is reduced and drivers depend more on upstream warnings [62].

Traditional queue monitoring has relied on point detectors such as inductive loops, radar trailers, microwave sensors, and portable units. These devices provide useful information but are limited by their fixed locations, hardware failures, and difficulty matching detector positions to the true queue boundaries. Performance declines when detector spacing is large or when queues form outside the monitored region [63,64]. Deployment, power, calibration, and maintenance needs also limit the number of locations that can be monitored.

The growing availability of CV telemetry provides a more flexible and scalable alternative. CV data contain frequent GPS locations, speeds, headings, timestamps, and trajectories that describe natural vehicle movement along a given corridor. Because probe vehicles travel through the full spatial extent of the network, they reveal where queues form, how far they extend, how long they last, and how they dissipate. Studies have shown that probe speeds can approximate traffic state indicators and detect queues with performance comparable to or better than many fixed-detector systems when penetration is sufficient [65,23,66,18]. Probe data offer broad spatial coverage without field hardware.

This chapter describes a cloud-native real-time system that uses CV telemetry to monitor average speeds, detect queues, generate standardized queue warnings, and produce detailed trajectory plots around work zones. The system continuously ingests CV data, aggregates them using distributed computing tools, identifies slow-moving conditions that indicate queues, and publishes alerts in both JSON and TMDD-compliant XML formats for use by TMCs. An AWS EC2-based pipeline produces high-resolution trajectory plots that provide insight into merging behavior, turbulence, and queue formation near work zones [60,67]. Together, these components support statewide short-term planning and operations.

3.2. Data Inputs and Processing Requirements

The real-time system relies on CV telemetry as its primary input. The live CSV feed contains spatial coordinates (latitude, longitude, and centroid), sensor identifiers, timestamps, and operational variables such as speed, direction, vehicle presence, and occupancy-style values derived from upstream telematics platforms. Files arrive continuously in Amazon S3 and are partitioned by date for traceability and efficient retrieval.

A monitoring function scans recent S3 objects using metadata queries and identifies new arrivals based on their modification time. This design separates upstream data delivery from downstream processing and allows the pipeline to tolerate variability in arrival patterns. Detected files are loaded into Apache Spark as DataFrames, which support vectorized transformations and grouping. Early preprocessing includes converting capture times into a standardized timestamp format so that records are interpreted consistently across time zones.

Spatial preprocessing uses Apache Sedona to convert latitude and longitude into geometric point objects and to run point in polygon tests. This step keeps only those vehicle points that lie within work zone polygons or near monitored sensor locations, which reduces computational load while preserving relevant observations. A simple presence and occupancy-style metric is computed, with occupancy set to five when vehicles are present and zero otherwise, creating a binary indicator for sensor activity.

Speeds are discretized into three operational categories: severe slowdown (less than 25 mph), transitional flow (25 to 45 mph), and near free-flow (greater than 45 mph). This categorization reduces random variation and makes traffic states easier to interpret. Route and centroid metadata are retained and passed to downstream alert generation so that each message can be clearly linked to a roadway segment.

3.3. Live Speed Conditions within Active Work Zones

This section describes how the system converts CV data into live speed metrics around work zones and exposes them through JSON-based alerts.

3.3.1. Aggregation Logic

Within AWS Glue, the aggregation logic transforms raw CV points into structured metrics that represent the operational state at each monitored location. CV penetration varies by region and time of day, so observations are grouped into fixed five-minute intervals to obtain stable estimates.

For each interval, Spark groups records by sensor identifier and computes earliest and latest timestamps, capture duration, average speed, and total vehicle count. When average speeds fall below a low-speed threshold or when a large share of observations fall into the slowest speed category, the location is classified as operating under queued or near-queued conditions. The binary presence metric provides additional context on sustained vehicle activity in the observation window.

Speed thresholds alone do not capture every queue scenario, but they remain widely used because they are simple and easy to implement [63]. Probe-based speed studies have shown that low-speed threshold methods can detect queues reliably when CV penetration is adequate [66]. For work zones with reduced capacity, speed-based thresholds provide a practical indicator of congestion onset.

The architecture also supports spatial refinement using Sedona so that queue classifications can be limited to activity that occurs within or near active construction areas. Aligning speed drops with work zone geometry helps reduce false positives and improves the operational value of the alerts.

3.3.2. Generation of JSON-Based Queue Alerts

After aggregation, the system generates structured JSON alerts that summarize conditions at each sensor. Alerts consist of two primary objects: an inventory object and a data object. The inventory object includes sensor identifier, route, centroid coordinates, and sensor name, so that downstream systems can map the alert to a specific location. The data object reports average speed, presence indicator, vehicle count, and timestamps in both Coordinated Universal Time and Central Standard Time, which supports integration with a range of logging and scheduling systems.

Alerts are written to S3 in two locations. A short-term staging path triggers downstream dissemination, while a long-term archive retains alerts for auditing, performance evaluation, and reporting. A consistent file naming convention (for example, `detector_YYYYMMDD_HHMMSS.json`) supports systematic tracking and troubleshooting.

Real-time traffic research emphasizes the value of structured, machine-readable messages that can be automatically parsed and integrated into situational awareness platforms [68]. JSON provides a lightweight, widely supported format and can be queried at scale using tools such as AWS Athena, which makes it suitable for both operations and retrospective analysis.

3.4. Queue Detection and Upstream Warning Generation

The main objective of the system is to detect developing queues near work zones and provide early warnings to upstream drivers through DMS and other traffic management systems. Aggregated speed- and presence-based measures derived from CV data act as primary queue indicators. When conditions meet specified thresholds, the system issues alerts that can be transformed into TMDD-compliant XML and pushed to external systems.

The overall architecture is built on the AWS ecosystem and follows common cloud-native design practices (Figure 3.1). Connected vehicle CSV files arrive in S3, a monitoring script identifies recent files on a short refresh cycle, and an AWS Glue job processes them in distributed fashion. Apache Sedona performs spatial filtering to focus on relevant work zones, and JSON alerts are written to S3 for real-time and archival use. Prior work on cloud-based traffic systems has shown that distributed frameworks such as Spark can process streaming probe data at scale and support real-time applications [69].

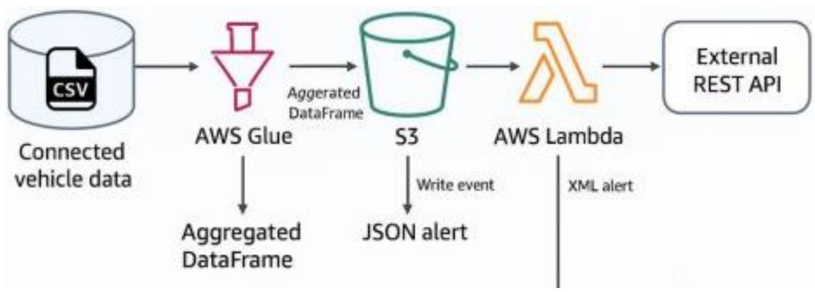

Figure 3.1. Cloud-based CV data processing pipeline illustrating AWS Glue aggregation, Amazon S3 storage, AWS Lambda-based alert generation, and integration with external REST APIs

Operationally, the system uses a 30-second refresh cycle, which is consistent with typical update intervals used in many real-time lane closure and congestion monitoring systems [18]. End-to-end latency from vehicle passage to alert availability is kept low through frequent S3 checks, distributed Spark processing, and immediate Lambda execution. Spark and Glue support horizontal scaling as CV penetration grows, in line with broader guidance on scalable cloud-based transportation architectures [69,70]. TMDD-compliant XML formatting ensures interoperability with existing active traffic management system (ATMS) and DMS infrastructures and aligns with established data exchange standards [30,71]. The architecture is also extensible; additional analytics and visualizations can be connected using Amazon Athena, QuickSight, or Redshift [72], and the logic can be migrated to EC2 or hybrid infrastructures for cost and performance optimization [73].

3.4.1. XML Transformation and API-Based Dissemination

The final alert dissemination stage converts JSON alerts into TMDD-compliant XML and sends them to an external REST API endpoint. An AWS Lambda function is invoked whenever a new JSON alert appears in the staging path.

Lambda loads the JSON, extracts the inventory and data fields, and maps them into the hierarchical structure specified by the TMDD. TMDD is widely used for representing detector data, incident records, and DMS-related information in agency systems [30]. Constructed XML messages include station metadata, timestamps, average speeds, and presence or occupancy indicators.

After the XML document is assembled, Lambda sends an HTTP POST request to the endpoint at `reactorfeeds.org`. The endpoint validates the message and integrates it into downstream workflows such as DMS updates or ATMS dashboards. Lambda includes appropriate content-type headers and basic error handling for failed requests.

Studies on data interoperability in traffic management highlight TMDD as a robust standard for interagency data exchange [74]. Using TMDD-compliant XML allows queue alerts to be incorporated into existing systems with minimal custom integration work.

3.4.2. File Retention, Cleanup, and Reliability Measures

To maintain performance and data quality, the system includes automated file retention and cleanup functions. A module in the Spark and Glue environment periodically scans JSON and Parquet directories and removes files older than a configured 10-minute retention window. This sliding window approach reflects common practice in real-time ITS pipelines, where stale data can misrepresent rapidly changing conditions, especially near congested work zones [75].

Limiting directory size also improves metadata listing performance in object storage. Studies of cloud storage behavior have shown that very large directory structures can increase listing times when directories are scanned often [76]. Regular cleanup reduces this risk and helps maintain low latency for the queue warning pipeline.

Reliability is enhanced through redundant JSON storage paths: one for real-time dissemination and one for long-term archiving. If a staging alert fails to be processed, the archived copy can be used for diagnostics. The stateless nature of Lambda and the retry options for Glue jobs further support resilience during fluctuating loads. These design choices are consistent with guidance on fault tolerance and elastic scaling for cloud-based transportation analytics [70].

3.5. Vehicle Trajectory Analysis for Operational Monitoring

Aggregated speed and queue alerts provide concise indicators of traffic state, but detailed trajectory analysis offers deeper insight into how congestion develops and dissipates around work zones. This section summarizes the EC2-based pipeline that generates trajectory plots and explains how these plots support operational monitoring and real-time incident detection.

3.5.1. EC2-Based Trajectory Processing Pipeline

A dedicated analytic pipeline on Amazon EC2 generates high-resolution trajectory plots in near real-time (Figure 3.2). This pipeline runs in parallel with the Glue-based aggregation process. It reads the same CV files from S3, but instead of aggregating data, it focuses on individual vehicle traces.

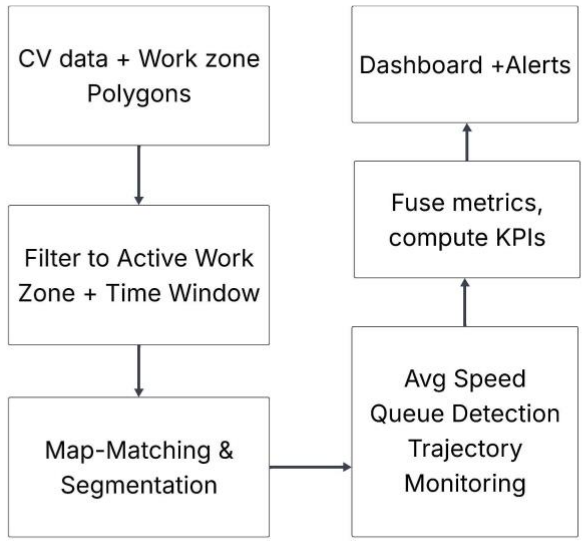

Figure 3.2. Real-time incident detection pipeline illustrating data ingestion, segmentation, metric computation, and key performance indicator (KPI) fusion for real-time work zone monitoring

The pipeline uses Apache Sedona or Shapely to intersect vehicle points with work zone polygons stored in S3, extracting only trajectories that pass through the active work zone. Localized filtering around bottleneck areas can provide clear insight into queue onset and merging turbulence [86]. Filtered points are then ordered by journey identifier and timestamp to reconstruct continuous trajectories. GPS coordinates are converted into distance along route or segment indices, which is a common step in trajectory-based congestion studies [87]. Python libraries such as Matplotlib or Plotly are used to build time-space diagrams, with time on the x axis and distance or segment index on the y axis. Speed is encoded as a color gradient, which highlights slow-moving or stopped vehicles. The pipeline saves PNG images and optional summary statistics (such as minimum speeds, vehicles per segment, and distance to taper) back to S3, making them available for dashboards and offline analysis.

Trajectory visualization has been widely applied in traffic flow research to analyze breakdown locations, capacity drops, and turbulence [64,79].

3.5.2. Trajectory Plots Around Work Zones

Trajectory plots provide an intuitive picture of queue formation and dissipation. Each vehicle appears as a slanted line in time space, and the slope reflects its speed. Dense clusters of nearly horizontal lines indicate very low speeds or stopped vehicles, which shows the presence of queues and shockwaves [80]. As vehicles decelerate while approaching the taper, trajectories flatten; when merging occurs, the shape and spacing of trajectories change, revealing turbulence near the lane drop. Deceleration waves appear as changes in slope across many vehicles and propagate upstream at characteristic velocities described in traffic flow literature [81].

By examining these plots, operators can study how different work zone configurations influence traffic stability. Short tapers may result in abrupt merges and visible turbulence, while longer tapers may lead to smoother merging. Recovery patterns after a queue forms appear as trajectories gradually steepening as vehicles return to higher speeds. Persistent instability downstream of a bottleneck, reflected in irregular acceleration patterns, has been linked to elevated crash risk [61].

Trajectory plots also help distinguish demand-driven congestion, which forms and dissipates quickly, from geometry-driven congestion caused by lane reductions or restrictive tapers, which tends to produce more persistent low-speed bands. EC2-generated plots provide visual validation of aggregated speed metrics and queue alerts, allowing operators to confirm slow traffic, investigate underlying causes, and check whether TMDD-based queue warnings reflect observed conditions.

3.6. Conclusion

This chapter presented a cloud-native system for monitoring average speeds, detecting queues, issuing real-time queue warnings, and generating detailed trajectory plots around highway work zones using CV telemetry. Continuous ingestion through Amazon S3, distributed processing with AWS Glue and Spark, spatial filtering using Sedona, and standardized dissemination via JSON and TMDD-compliant XML together provide a scalable, interoperable, and low-latency platform.

Queue detection in the system is based on aggregated speed and presence-style metrics computed over short time windows. Prior work has shown that such metrics are practical indicators of congestion onset when probe penetration rates are sufficient [63,66]. Real-time JSON and XML outputs provide information that can support DMS and other ATMS components, which in turn can help reduce the likelihood of secondary crashes. Automated retention and reliability mechanisms support continuous operation under changing loads [75,76,70].

The EC2-based trajectory pipeline adds a complementary layer of detail, capturing shockwaves, merging behavior, and queue dynamics around work zones. These plots help agencies understand the mechanisms of congestion, assess temporary traffic control plans, and plan improvements for future work zones [77,78,64,79,80,81,61].

As CV penetration increases and agencies expand their cloud-based ITS deployments, systems of this type are likely to play an important role in improving work zone safety and mobility. Future work may incorporate predictive queue models, automatic work zone boundary detection, and integration with camera-based analytics to support multimodal situational awareness, building on the foundation established by this CV-driven platform.

CHAPTER 4. DAILY OPERATIONS AND MANAGEMENT

4.1. Introduction

Monitoring work zone performance has become increasingly important as roadway conditions grow more dynamic and traffic volumes continue to rise. Temporary work zones modify the roadway cross section, reduce capacity, and require drivers to merge or shift laterally, which introduces turbulence in the traffic stream and increases crash risk, particularly during peak demand. FHWA guidance on work zone performance measurement emphasizes that agencies should track mobility and safety impacts in real-time and move toward performance-based management of lane closures and construction activities [14,82].

Traditional monitoring tools such as fixed loop detectors, radar trailers, and portable cameras provide valuable information, but they rarely align perfectly with work zone extents and often lack complete temporal coverage. Sensors are usually installed at a limited set of locations, which means that queues may extend beyond the instrumented area and that short-duration disruptions may be missed. Many state departments of transportation (DOTs) do not have dense freeway detector coverage, especially on rural or lower-volume corridors, which limits their ability to monitor work zones consistently across the network [16].

CV data, when combined with cloud computing tools, provide an alternative for continuous performance monitoring. Crowdsourced vehicle trajectories and segment speeds have been shown to capture congestion onset and recovery with a latency comparable to that of some fixed sensors, particularly for detecting queue formation near work zones when penetration rates are sufficient [14,83]. Because CVs report speeds and locations along the entire corridor, they allow agencies to observe traffic behavior along the work zone rather than only at discrete detector points.

The performance monitoring system described in this chapter processes CV data every five minutes and converts them into a suite of metrics that describe current operating conditions within each predefined work zone polygon. These metrics quantify how traffic is moving, where congestion is forming, how long queues persist within each time slice, and how much additional delay is imposed relative to nominal speeds. The chapter explains what each metric represents, why it matters from an operations perspective, and how it is derived from CV telemetry. It also documents the preprocessing steps, spatial segmentation framework, and logic used to identify queues, calculate delays, and aggregate results over rolling 30-minute windows.

By the end of the chapter, the reader will understand both the practical relevance of each metric and the analytical pipeline that generates it. The chapter also demonstrates the metrics using several example work zones, where speed segment heatmaps, speed-time scatterplots, and queue length time series have been generated from Iowa CV data. This builds on prior research that evaluated work zone conditions using probe-based datasets [14,83,82].

4.2. Background and Importance of Work Zone Performance Metrics

4.2.1. Role of Performance Measures in Work Zone Management

Performance measures translate raw data into indicators that support operations. FHWA distinguishes between measures that describe mobility, such as travel time, queue length, and delay, and measures that describe safety, such as crash frequency, speed variability, or exposure to stop-and-go traffic [82]. For work zones, these two dimensions are closely related. A lane closure that causes long queues can also increase crash risk, while an effective configuration maintains acceptable mobility and improves conditions for workers and drivers. National scans have documented how agencies use real-time and archived indicators to manage construction activities and evaluate work zone plans [16]. Many indicators focus on queues because unexpected stopped or slow traffic at work zone approaches has been linked with rear-end crashes. FHWA guidance recommends that agencies monitor queue length, maximum queue, queue duration, and user delay as core measures for freeway work zones [14,82].

Historically, these measures were often estimated from point detectors using shockwave methods and speed thresholds or by manual field observations. While useful in some locations, these approaches are difficult to scale and are sensitive to detector placement and reliability. As a result, agencies have faced challenges in implementing performance-based work zone management consistently.

4.2.2. Emerging Use of Probe and Connected Vehicle Data

The growth of probe and CV data has created new options for work zone performance measurement. Probe-based speed products aggregate crowdsourced data from navigation services, while CV trajectory datasets contain anonymized waypoints from individual vehicles. Several studies have shown that these data sources can reproduce congestion patterns observed by fixed sensors and can be used to estimate travel times, identify bottlenecks, and measure queue lengths when penetration levels are adequate [14,84,85].

Recent work under the Smart Work Zone Deployment Initiative (SWZDI) compared several probe-based performance measures and reported that CV trajectory data can detect work zone queues with low missed call rates and reasonable latency, even when overnight coverage is sparse [83]. These findings support the use of CV data as a primary input for automated queue warning systems and performance dashboards. They also emphasize the need for robust spatial filtering, trajectory aggregation, and threshold selection so that performance measures remain stable and easy to interpret.

The present system builds on these findings and on the lane closure detection framework described in Chapter 3 [1]. It uses CV trajectories to compute segment-level speeds, classify reduced speed conditions, determine when and where queues form, and translate these patterns into standardized metrics that align with FHWA-recommended measures [14,82].

4.3. Connected Vehicle Processing Pipeline (Spark-Sedona-S3 Architecture)

The real-time work zone monitoring system uses a cloud-based pipeline that ingests CV data from Amazon S3, applies geospatial filtering and segmentation using Apache Sedona, aggregates the data with Apache Spark, and produces analysis-ready inputs every five minutes. This section describes the architecture that supports the metric framework described in Section 4.4.

4.3.1. Ingestion of Recent Connected Vehicle Telemetry

CV data are delivered to an Amazon S3 bucket as Parquet files throughout the day. A Boto3-based scanner identifies all files whose `LastModified` timestamp falls within the previous 30 minutes. This approach ensures that delayed uploads are still captured and that the pipeline uses the most recent available telemetry.

Spark's `s3a` connector is used to read these files in parallel, which supports distributed input output and large hourly volumes. Each record includes a journey identifier, latitude and longitude, instantaneous speed, work zone identifier, route identifier, and a capture timestamp.

4.3.2. Geospatial Filtering Using Apache Sedona

After ingestion, the pipeline applies geospatial filtering to retain only points that fall inside predefined work zone polygons. All waypoints are converted into Sedona point geometries, and filtering is performed using `ST_Contains(polygon, point)`.

For each filtered point, Sedona computes several geometric attributes (Table 4.1). This produces a consistent spatial segmentation across all work zones and time periods.

Table 4.1. Geometric attributes computed for each CV point during spatial preprocessing

Attribute	Description
Exterior ring	Boundary of the work zone polygon used for reference geometry.
Fractional position along reference line	Computed using `ST_LineLocatePoint` to obtain the relative location of the point along the polygon's reference line.
Distance along corridor (ft)	Linear measure from the start of the work zone reference line to the point, expressed in feet.
Segment identifier	Derived by dividing the corridor distance by 528 ft (0.1 mile), producing consistent spatial segmentation.

4.3.3. Timestamp Normalization and Five-Minute Binning

CV timestamps arrive in mixed formats, including epoch seconds, epoch milliseconds, and strings. A unified conversion function converts each timestamp to a common UTC `timestamp` type. Each point is then assigned to a five-minute analysis window by flooring the UTC time to the nearest 300 seconds:

$$t_{bin} = \left\lfloor \frac{t_{UTC}}{300} \right\rfloor \times 300.$$

Bin start times and event timestamps are also expressed in local time (America/Chicago) for reporting and dashboard visualization.

4.3.4. Work Zone Level Percentiles

To support delay estimation and reference speed normalization, all observations belonging to the same work zone, route, and time bin are aggregated to compute approximate 85th and 35th percentile speeds using Spark's `percentile_approx`:

$$v_{85,t}, \quad v_{35,t}.$$

The 85th percentile speed serves as a free-flow reference for delay calculations, while the lower percentile characterizes intermediate and congested conditions.

4.3.5. Segment-Level Aggregation

For each work zone, route, segment, and five-minute window, the system records four quantities: the number of CV observations, the number of unique journeys, the mean segment speed $\bar{v}_{s,t}$ in mph, and the first and last local timestamps observed within the bin.

These segment-level aggregates, combined with the work zone–level percentiles, form the inputs for the metric computation framework described in Section 4.4.

4.4. Metric Computation Framework

The metric computation framework operates on a rolling stream of CV telemetry that has been spatially filtered to work zone polygons and segmented along the work zone centerline. Each record includes a work zone identifier, route identifier, capture timestamp, instantaneous speed, and segment index defined in 0.1-mile increments.

Data are grouped into non-overlapping five-minute analysis windows in coordinated universal time, with local timestamps stored for reporting. Within each window, the system first forms segment-level and work zone–level aggregates, then identifies queued segments and queue groups, computes journey-level travel time and delay, and finally collapses all information into a five-minute summary record per work zone. A rolling 30-minute snapshot for operations is obtained from the six most recent summary records.

4.4.1. Segment-Level Speed Aggregation and Queue Thresholds

For each work zone, route, segment, and five-minute window, the system computes the number of CV records and distinct journeys, the mean segment speed $\bar{v}_{s,t}$ in mph, and the temporal coverage within the bin.

At the work zone level, the 85th and 35th percentile speeds,

$$v_{85,t}, \quad v_{35,t},$$

are computed using approximate percentile routines.

A segment is classified as queued when

$$\text{unique\_journeys}_{s,t} \geq 6 \text{ and } \bar{v}_{s,t} \leq v_{35,t}. \tag{4.1}$$

The minimum distinct journey requirement reduces the chance that isolated slow vehicles produce spurious queues. The queue classification defined in equation (4.1) is recomputed every five minutes.

4.4.2. Segment-Level Speed Profiles and Visualization

After segment-level aggregation, the pipeline maintains a matrix of segment indices versus time bins where each cell contains $\bar{v}_{s,t}$ and, when sample sizes allow, segment-level percentiles. This structure supports speed segment heatmaps, where color encodes mean or percentile speeds for each segment and time bin. These visualizations show the spatial extent of reduced speeds and the timing of congestion.

4.4.3. Queue Identification and Queue Length

Once segment-level queue flags have been assigned, the system groups contiguous queued segments within each work zone and time bin. For a queue that spans segments s_{start} through s_{end}, with segment length L_{seg} = 0.1 mi, the queue length is

$$L_{queue} = (s_{end} - s_{start} + 1) \times L_{seg}. \tag{4.2}$$

This converts the count of queued segments into a spatial measure. In line with FHWA queue estimation guidance, the system reports the length of each queue and summary statistics across all queues in the bin [82].

Within each five-minute bin, the system tracks maximum queue length, minimum and average queue length across detected queues, and number of distinct queues (Table 4.2).

Table 4.2. Journey-level aggregates computed within each work zone and five-minute analysis window

Quantity	Symbol	Description
First local timestamp	t_{first}	Earliest local timestamp for the journey within the analysis bin.
Last local timestamp	t_{last}	Latest local timestamp for the journey within the analysis bin.
Record count	–	Number of CV records associated with the journey in the bin.
Minimum segment index	s_{min}	Smallest segment index visited by the journey in the bin.
Maximum segment index	s_{max}	Largest segment index visited by the journey in the bin.
Average observed speed	v_{obs}	Mean of instantaneous speeds observed for the journey in the bin.
Minimum observed speed	v_{min}	Lowest instantaneous speed observed for the journey in the bin.

4.4.4. Per Bin Queue Duration and Temporal Evolution

The analysis uses non-overlapping five-minute bins. Each queue identified within a bin is associated with a nominal duration of

$$T_{queue,t} = 5\ minutes. \tag{4.3}$$

For real-time operations, the system summarizes the queues within each bin using this fixed duration. Longer-term temporal evolution can be reconstructed by linking queues across consecutive bins based on spatial overlap and computing total episode duration as

$$T_{episode} = N_{bins} \times 5\ minutes, \tag{4.4}$$

where N_{bins} is the number of consecutive bins with at least one queue present [91].

4.4.5. Journey-Level Travel Time and Delay

Delay represents the additional travel time experienced relative to a reference speed. In this framework, delay is computed at the journey level.

For each work zone, route, journey identifier, and five-minute bin, the system aggregates all records and computes the following:

Observed travel time is calculated as

$$t_{actual} = t_{last} - t_{first}, \tag{4.5}$$

and path length within the work zone is approximated as

$$L_{path} = (s_{max} - s_{min} + 1) \times L_{seg}, \tag{4.6}$$

where L_{seg} = 0.1 mi. Journeys with nonpositive path length are excluded.

Expected travel time is calculated using the work zone 85th percentile speed:

$$t_{expected} = \begin{cases} \frac{L_{path}}{v_{85,t}} \times 3600, & v_{85,t} > 0, \\ undefined, & otherwise. \end{cases} \tag{4.7}$$

Journey-level delay is calculated as

$$\Delta t_{sec} = \max(0, t_{actual} - t_{expected}). \tag{4.8}$$

Journey-level delays are summarized at the work zone and bin level using mean, minimum, and maximum delay values.

4.4.6. Speed Event Frequency and Turbulence

Average speeds and delays describe typical conditions but may not capture short periods of instability. To characterize these regimes, the system defines three journey-level event flags:

- **Overspeed events**: Set when the journey's average observed speed exceeds $v_{85,t}$.
- **Speed drop events**: Set when the average speed at the start of the journey path exceeds the average speed at the end by at least 30 mph.
- **Hard braking events**: Set when the minimum observed speed along the journey path is less than or equal to 15 mph.

At the work zone and bin level, counts of overspeed, speed drop, and hard braking events are computed. Higher counts indicate more unstable traffic and higher exposure to sudden decelerations [19,83].

4.4.7. Throughput and Vehicle Counts

Throughput is estimated by counting the number of distinct journeys observed within each work zone and five-minute bin,

$$\text{total\_vehicles}_t = |\{j: j \text{ traverses the work zone in bin } t\}|.$$

CV counts represent a sample of the total traffic stream. Penetration rate studies have shown that CV counts can be scaled using factors derived at locations with both CV data and ground truth volume measurements [85].

Throughput estimates, combined with delay and queue length, help distinguish whether poor performance is driven by high demand, reduced capacity, or both.

4.4.8. Daily Time Series Visualization for Operations

Although metrics are stored at a five-minute resolution, the visualization layer shows the evolution of conditions over the entire current day. For each work zone, the dashboard retrieves all summary records with bin start times between local midnight and the current time. This produces a continuous time series of five-minute points.

These daily time series are used to plot average speed, maximum and average queue length, estimated delay, and counts of overspeed, speed drop, and hard braking events (Table 4.3). Presenting the full day on a single timeline allows operators to see how conditions have evolved and to compare current performance against earlier periods of the same day.

Table 4.3. Summary of core work zone performance metrics, computation logic, and aggregation level

Metric	How It Is Calculated	Formula / Condition	Time Bin / Level
Segment Mean Speed	Average of all CV speeds within segment *s* in time bin *t*	$\bar{v}_{s,t} = \frac{1}{N_{s,t}} \sum_{i=1}^{N_{s,t}} v_i$	5-min, segment
Work Zone Speed Percentiles	85th and 35th percentile speeds from all CV observations in bin *t*	$v_{85,t}, v_{35,t} = percentiles(\{v_i\}_t)$	5-min, work zone
Queue Segment Identification	Segment flagged as queued when mean speed is low and distinct journeys are sufficient	$\text{unique\_journeys}_{s,t} \geq 6$ and $\bar{v}_{s,t} \leq v_{35,t}$	5-min, segment
Queue Length	Length from contiguous queued segments times segment length (0.1 mi)	$L_{queue} = (s_{end} - s_{start} + 1) \times 0.1$	5-min, work zone
Delay	Difference between actual and expected travel time at reference speed	$\Delta t = \max(0, t_{actual} - t_{expected})$	5-min, journey
Throughput	Count of distinct CV journeys in bin *t*	$\text{throughput}_t = \lvert\{j : j \text{ traverses WZ at } t\}\rvert$	5-min, work zone
Overspeed Events	Journey flagged if average observed speed exceeds $v_{85,t}$	$overspeed_j = \mathbb{I}(v_{obs,j} > v_{85,t})$	5-min, work zone
Speed-Drop Events	Journey flagged if speed drops by at least 30 mph between start and end	$speed\_drop_j = \mathbb{I}(v_{start,j} - v_{end,j} \geq 30)$	5-min, work zone

For plots, all bins since local midnight are shown. For real-time summaries and alerts, only the most recent six bins (30 minutes) are examined.

4.5. Case Studies: Visual Interpretation of Traffic Dynamics at Selected Work Zones

This section presents two example work zones that show how speed segment heatmaps, speed-time scatterplots, and queue length profiles together describe congestion formation and dissipation. The chosen sites exhibit two patterns of morning peak behavior: a short moderate slowdown (4DA – Sideroad JCP EB 2) and a sharper slowdown (1DW – Mainline WB 1).

4.5.1. Case Study 1: Work Zone 4DA – Sideroad JCP EB 2

At work zone 4DA, the speed segment heatmap in Figure 4.1 shows a localized region of reduced speeds centered between approximately 07:30 and 08:10. A vertical band of orange and red cells appears across a short segment of the corridor, indicating speeds dropping to approximately 20 to 30 mph. Outside this interval, speeds remain largely in the green range, suggesting near free-flow conditions. The spatial extent of the slowdown is limited, indicating that congestion is confined to a small portion of the work zone. This pattern suggests a short-duration, demand-driven slowdown during the morning peak.

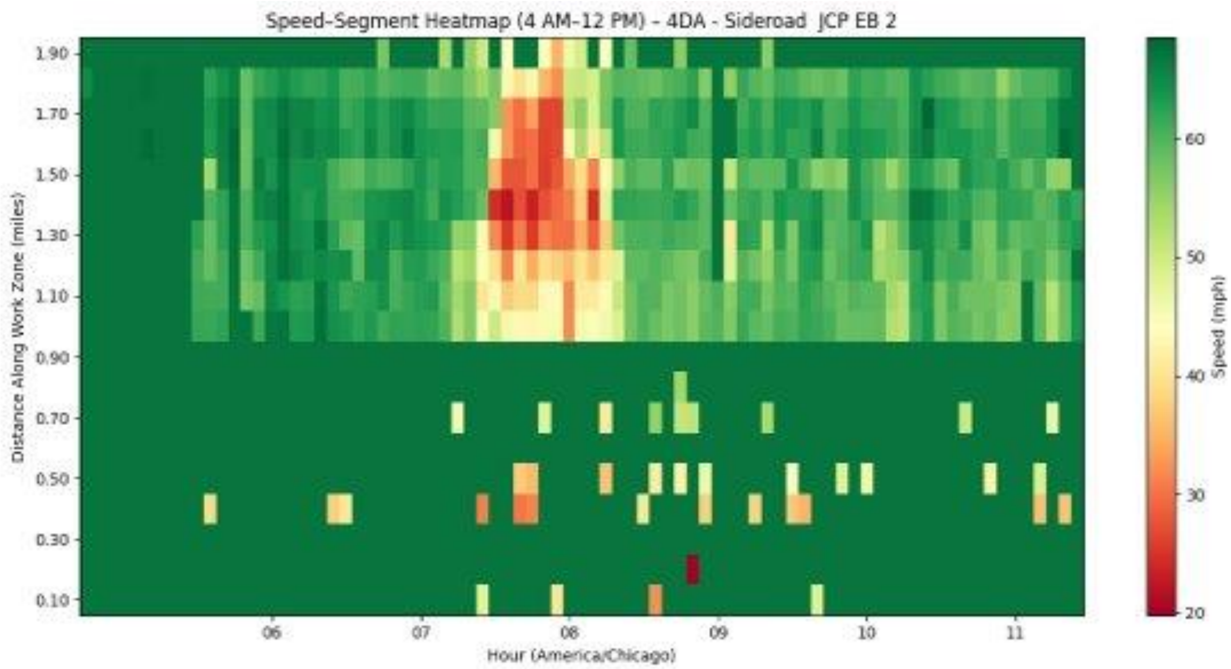

Figure 4.1. Speed-segment heatmap for work zone 4DA – Sideroad JCP EB 2

The speed-time scatterplot in Figure 4.2 shows a dense cluster of low-speed points primarily between 07:30 and 08:10, with speeds dropping below 35 mph. Before and after this interval, most observations are above 60 mph, indicating free-flow conditions. The limited temporal spread of low-speed points confirms that the slowdown is short-lived and concentrated within the peak period.

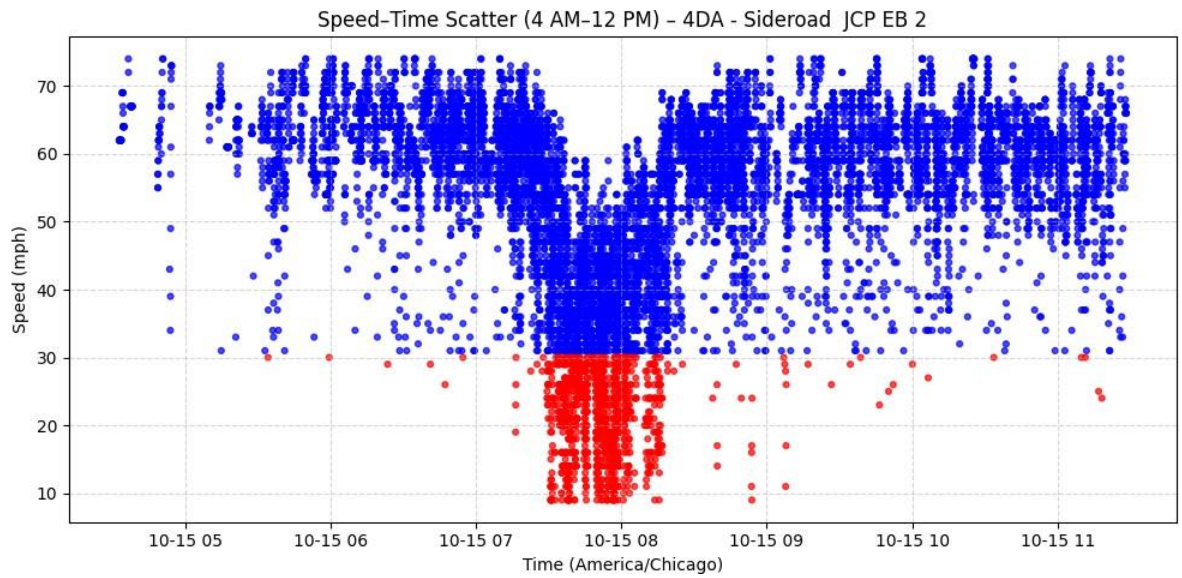

Figure 4.2. Speed-time scatterplot for work zone 4DA – Sideroad JCP EB 2

Figure 4.3 shows the queue length profile. Queue length remains near zero before 07:30, rises to nearly 0.6 mile, and then returns toward zero after about 08:10. The alignment between the slowdown band in the heatmap, the cluster of low-speed points, and the peak queue length indicates a short-duration demand-driven queue. The alignment between the localized slowdown in the heatmap, the concentrated cluster of low-speed observations in the scatterplot, and the short-lived peak in queue length indicates a compact, demand-driven queue with limited spatial extent and rapid dissipation.

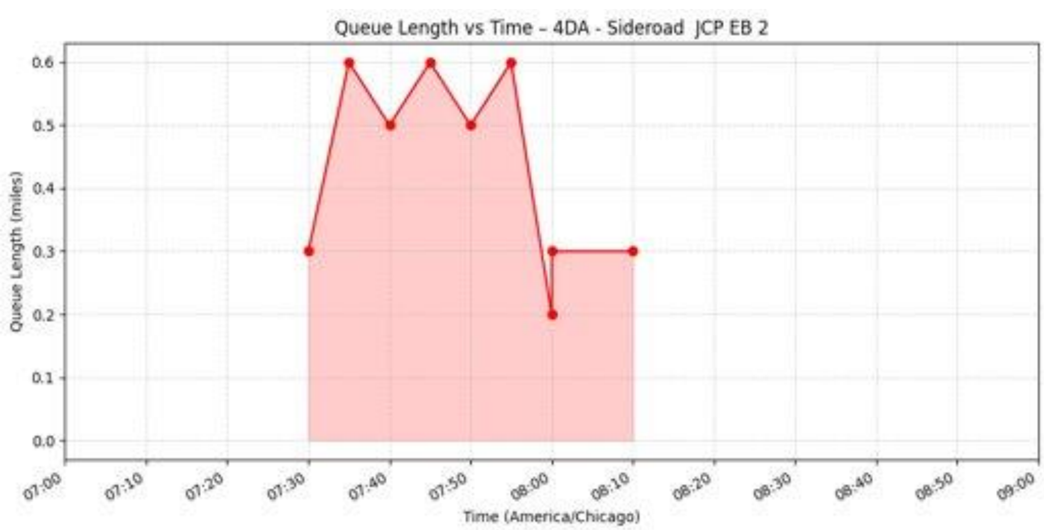

Figure 4.3. Queue length versus time for work zone 4DA – Sideroad JCP EB 2

4.5.2. Case Study 2: Work Zone 1DW – Mainline WB 1

Work zone 1DW shows a more intense but still limited slowdown. The speed-time scatterplot in Figure 4.4 shows a tight cluster of low-speed points emerging just before 07:50. Outside this window, speeds are close to free-flow.

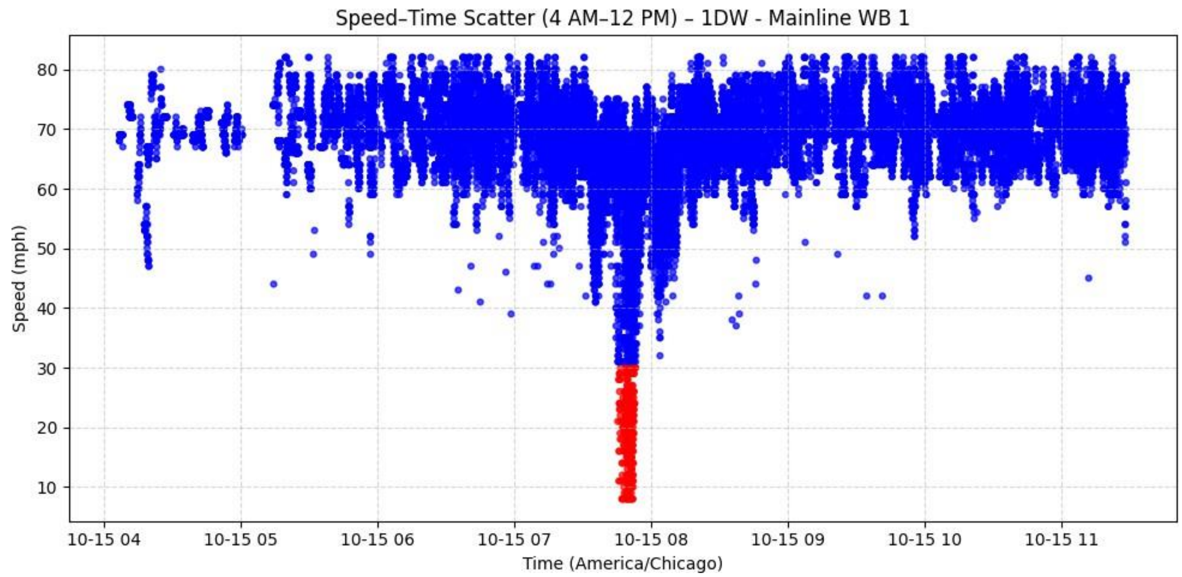

Figure 4.4. Speed-time scatterplot for work zone 1DW – Mainline WB 1

The heatmap in Figure 4.5 shows a vertical band of red cells across several segments just before 07:50, representing a short period of very low speeds. The slowdown disappears quickly, suggesting a short-duration congestion episode.

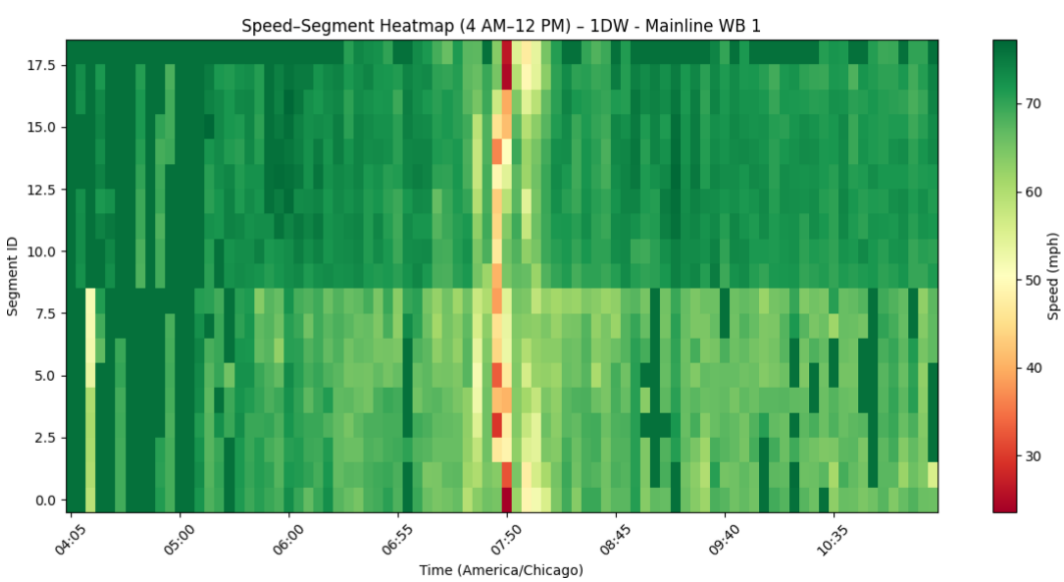

Figure 4.5. Speed-segment heatmap for work zone 1DW – Mainline WB 1

Figure 4.6 shows the queue length profile. Queue length remains near zero for most of the period, with a single sharp spike around 07:50 reaching approximately 0.5 miles. The queue forms briefly and dissipates almost immediately, returning to near-zero levels shortly after. Unlike the previous case, there is no sustained queue or gradual decline, indicating that the congestion event is highly transient.

Although the slowdown is intense, the short duration and immediate recovery suggest a brief, localized disturbance rather than a prolonged congestion buildup, and the overall queue length remains smaller and less persistent than in the 4DA case.

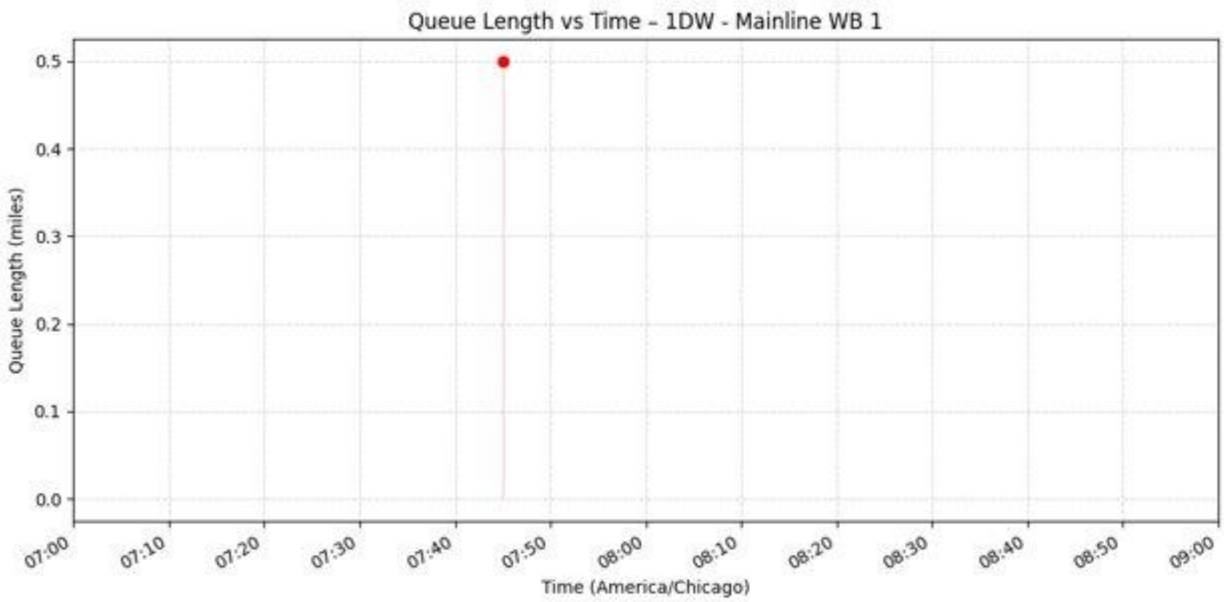

Figure 4.6. Queue length versus time for work zone 1DW – Mainline WB 1

4.5.3. Comparative Summary of Both Cases

Both case studies show morning peak slowdowns, but the patterns differ. Work zone 4DA experiences a moderate queue that forms and dissipates over a longer window, whereas work zone 1DW experiences a shorter but more severe slowdown. In both examples, the heatmaps, scatterplots, and queue profiles provide a consistent picture of congestion formation and recovery.

4.6. Real-Time Operational Insights from the Live Dashboard

The performance monitoring system described in this chapter features a live dashboard that provides a real-time view of work zone traffic conditions by combining spatial context, CV trajectory feeds, and performance indicators. This section illustrates the dashboard's functionality using an example work zone (1EA – Mainline WB 1).

The top-level map interface allows operators to select an active work zone and view its geospatial extent, surrounding roadway network, and the location of CV activity, as shown in Figure 4.7. This helps relate detected queues or speed drops to ramps, intersections, or local geometric features.

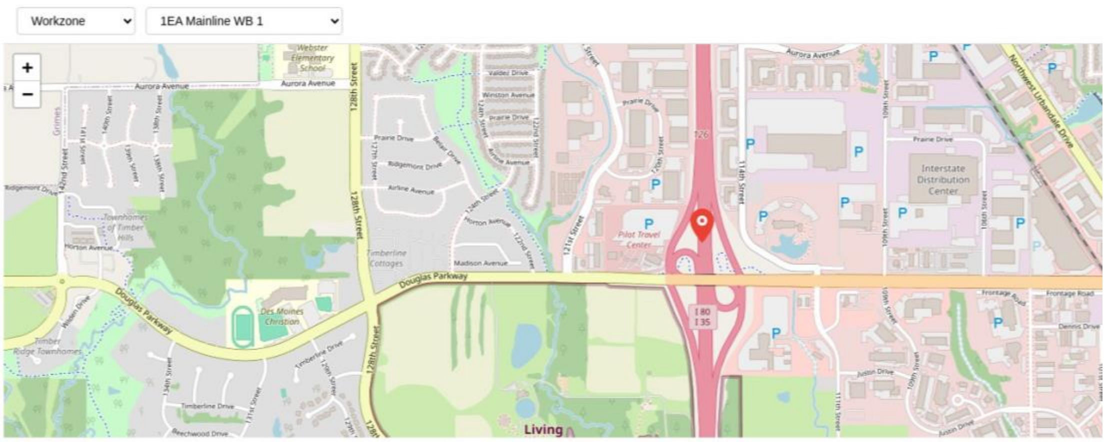

Figure 4.7. Live dashboard map interface showing work zone 1EA – Mainline WB 1

4.6.1. Queue Length Behavior

The queue length time series in Figure 4.8 shows two congestion periods during the day. In the morning peak, between approximately 07:30 and 08:15, the maximum queue length briefly exceeds 2.5 miles. A second congestion window occurs in the late afternoon between about 16:00 and 17:30, with queues generally remaining below 1.0 mile. Outside these periods, queue length remains close to zero.

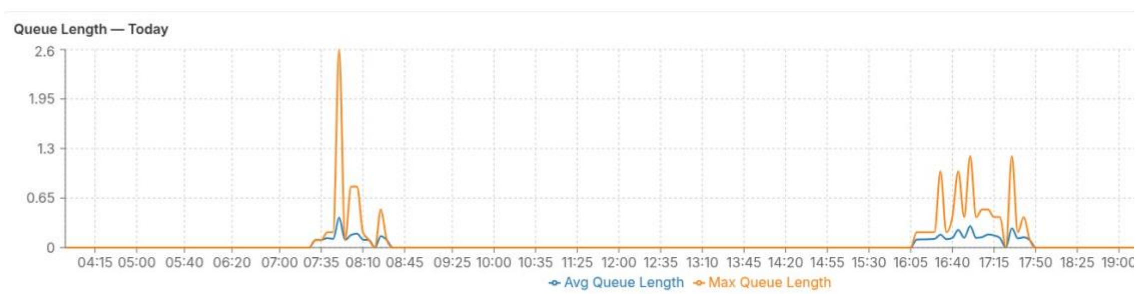

Figure 4.8. Real-time queue length throughout the day for work zone 1EA – Mainline WB 1

4.6.2. Queue Duration Characteristics

The queue duration plot in Figure 4.9 shows the persistence of queues. During morning and evening peaks, queue duration rises to the upper limit of the display (300 seconds). During midday and late night hours, queue duration is near zero. This indicates that congestion is confined to a limited set of high demand periods.

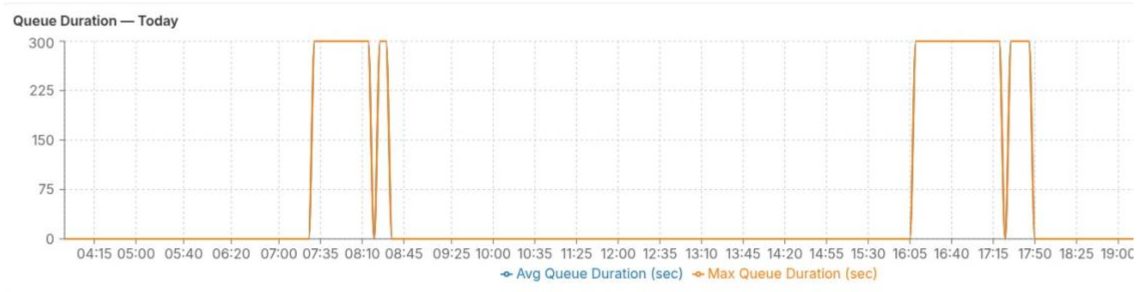

Figure 4.9. Queue duration time series showing the persistence of queue events for work zone 1EA – Mainline WB 1

4.6.3. Delay Trends

Figure 4.10 shows delay estimates. Delay remains small for most of the day, with brief spikes during the evening peak. Maximum delay reaches higher values for a short interval, while average delay remains relatively low. This indicates that a subset of vehicles experience higher delay when shockwaves develop, while many others pass through with smaller impacts.

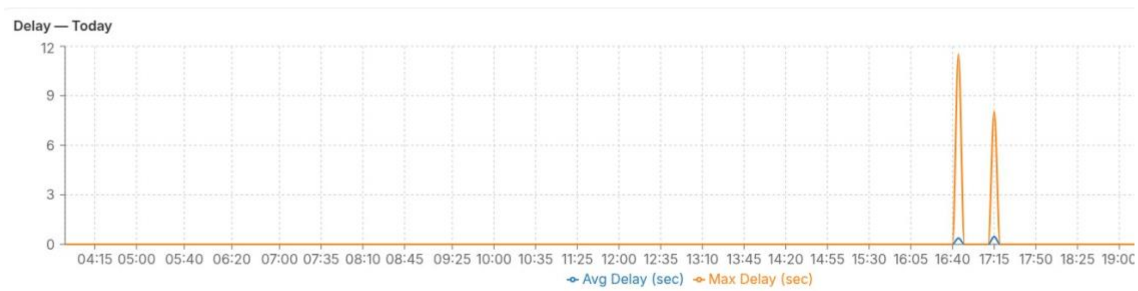

Figure 4.10. Delay estimates for work zone 1EA – Mainline WB 1

4.6.4. Connected Vehicle Trajectory Volumes

Figure 4.11 shows CV trajectory counts over the day. Counts rise from early morning into the 07:00 to 09:00 period, flatten during midday, and rise again between about 16:00 and 17:30. Higher volumes during the peaks support reliable speed and queue estimates. Even during off-peak hours, counts are sufficient to maintain monitoring.

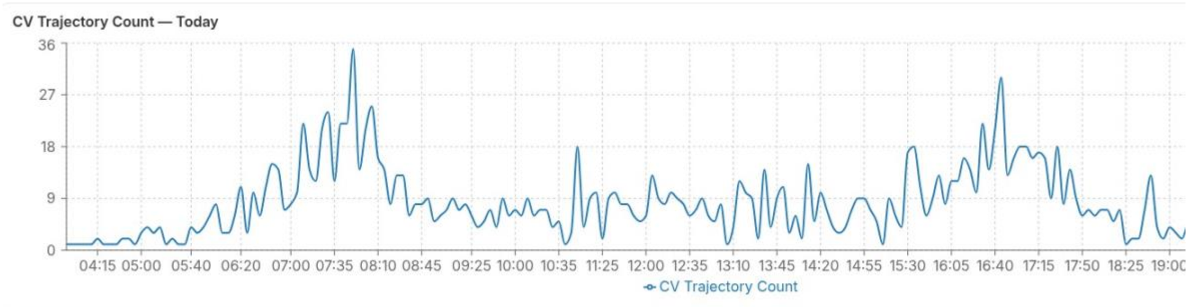

Figure 4.11. CV trajectory counts over the day for work zone 1EA – Mainline WB 1

4.6.5. Average Speed Variations

Average speed trends, shown in Figure 4.12, match the queue metrics. Speeds drop to around 25 to 35 mph during the morning congestion period, with the lowest speeds near 07:50. During the evening peak, speeds drop toward 20 mph before recovering as demand decreases after 18:00. Midday speeds remain between about 55 and 70 mph, indicating stable free-flow conditions.

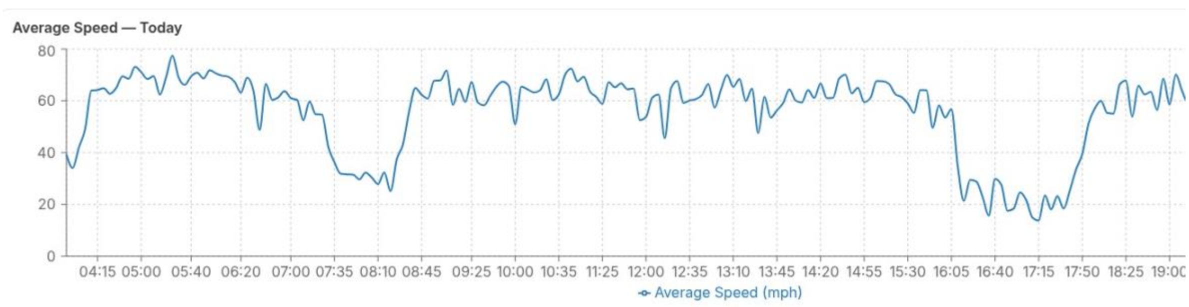

Figure 4.12. Average speed time series for work zone 1EA – Mainline WB 1

4.6.6. Operational Summary

The dashboard results indicate that congestion at this work zone is concentrated in the morning and evening peaks, while midday and overnight periods show limited degradation. The morning peak produces longer queues but relatively low delay per vehicle, consistent with slow but moving traffic. Evening congestion generates shorter queues but higher delay spikes, which suggests more stop-and-go behavior. Together, the metrics and visualizations allow practitioners to diagnose conditions, assess safety risks associated with queue formation, and identify opportunities for targeted mitigation strategies.

4.7. Chapter Summary

This chapter presented a framework for measuring work zone performance using CV data and cloud-based analytics. Building on FHWA guidance and recent work with probe-based congestion monitoring, the chapter described how CV telemetry can be transformed into indicators of queue length, queue duration, delay, speed event frequency, and throughput [14,19,83,82].

The data pipeline developed for the framework integrates near-real-time ingestion of CV records, temporal standardization, and spatial filtering using work zone polygons. Each work zone is segmented into 0.1-mile units so that localized disruptions can be captured, and five-minute time bins are used to balance responsiveness with stability. Reduced-speed conditions are identified using percentile-based thresholds, and queues are detected by grouping contiguous segments that meet these conditions. From these patterns, the system computes the spatial extent and temporal duration of congestion, as well as the additional travel time imposed on drivers.

Visualization is a key part of the framework. Speed segment heatmaps, speed-time scatterplots, and queue length profiles were developed for two example work zones. The first site exhibited a short demand-related slowdown during the morning peak, while the second site showed a shorter but more intense congestion period. In each case, the agreement between visual diagnostics and computed metrics demonstrated that CV data can capture the formation and dissipation of work zone queues.

By combining numerical metrics with interpretable visualizations, the system provides a practical assessment of work zone conditions that supports real-time monitoring by TMC staff. These metrics can support queue warnings, traveler information, and performance-based evaluations of lane closures. As CV penetration increases and cloud computing tools continue to advance, such systems will help agencies monitor many work zones without additional roadside hardware [85,83].

Future work may extend these metrics to include incident-related closures, recurring bottlenecks, and special events. Further research on short-term queue prediction and data fusion methods that combine probe data with selective fixed sensors may enable forecasting and early warning capabilities that support work zone operations [84,86,32]. Subsequent chapters will describe flow estimation and the integration of these metrics into the broader ReactorIQ dashboard environment.

CHAPTER 5. SMART WORK ZONE PLANNING

This chapter presents methods for monitoring and evaluating freeway work zones using connected vehicle data and data-driven analytics. The primary focus is on automatic lane closure detection, which identifies the location and timing of lane closures using changes in vehicle speed and lateral position patterns. This behavior-driven approach provides continuous, scalable detection of work zone activity without reliance on manual reporting or fixed roadside sensors, improving situational awareness for transportation agencies

5.1. Automatic Lane Closure Detection

5.1.1. Introduction

Lane closures are among the most influential temporary changes that occur on high-speed highways during construction, maintenance, utility work, incident response, and rehabilitation activities. When a travel lane becomes unavailable, the effective capacity of the roadway decreases and drivers must shift laterally, merge, or navigate narrower paths. These adjustments alter traffic flow and can produce queues, speed reductions, and increased lane changing activity that extend upstream of the work area. The effects are most pronounced during peak periods, when demand is high and the roadway is already operating near capacity.

National assessments of work zones and lane closures have documented these impacts. Analyses of work zone-related crashes and nonrecurring congestion show that temporary lane reductions contribute meaningfully to delay and safety risk, including primary and secondary crashes and associated economic costs [1,3]. Recent FHWA summaries indicate that lane reductions in work zones are an ongoing safety concern on high-speed facilities, particularly for mixed traffic and commercial vehicles [2,3].

Because of these risks, lane closure management is a central task for state transportation agencies. Timely and accurate closure information supports law enforcement deployment, DMS updates, queue warning systems, contractor compliance checks, and evaluation of traffic control plans [2]. In practice, however, lane closure records have often been generated from manual sources such as contractor logs, 511 reports, and field verification by inspectors or TMC staff. These approaches can be delayed, incomplete, or inconsistent. Field sensors such as loops, radar trailers, Bluetooth readers, or cameras provide additional information but are costly to install and maintain and do not cover every location [87,28,29]. Short-duration and mobile closures further complicate reliable documentation when their exact location and timing change during a shift or across multiple days.

Several national efforts have sought to standardize and distribute lane closure information more systematically. The FHWA WZDx specification defines a common format for publishing planned and active work zone records, including closure locations, lane configurations, and time windows [28]. Studies on smart work zone devices and connected arrow boards show that onsite devices can provide high-quality lane closure data when properly deployed and maintained

[28,29]. These devices, however, require dedicated hardware, power, and communications, which can make universal deployment across all work zones challenging.

CV telemetry offers a complementary data source. CV telemetry data collected at one- to three-second intervals provide detailed trajectories that include coordinates, timestamps, speeds, headings, and journey identifiers [88,34]. These trajectories describe how drivers move under actual conditions and implicitly encode lateral positioning through the spread of waypoints across the roadway cross section. When a lane is closed and drivers shift into the remaining lanes, the distribution of waypoints changes in a systematic way. By aggregating many vehicles, these changes can be observed without any additional roadside hardware [1,35].

This section describes a methodology for automatic lane closure detection that uses CV movement data and machine learning. The approach follows the statewide lane closure detection framework proposed, which distinguishes closure conditions from normal operations using distributional changes in speed and lateral positioning [1]. Weekly CDFs of speed and distance from the reference line (DFRL) are compared with a baseline week, Kolmogorov-Smirnov–based features are extracted, and support vector machine (SVM) and k nearest neighbor (KNN) classifiers identify closure weeks for 0.1-mile segments.

In the remainder of this section, that methodology is summarized and integrated into the broader work zone analytics framework presented in this report. The following subsections describe the data preparation, spatial mapping, feature construction, development of a classification pipeline, cross-corridor validation, formulation of a statewide scaling strategy, postprocessing for stability, and visualization. The objective is to document the algorithmic steps and show how the lane closure detection module fits into the larger work zone monitoring system.

5.1.2. Literature Review

5.1.2.1. Work Zones, Lane Closures, and Nonrecurring Congestion

The safety and mobility effects of work zones and lane closures have been studied extensively in both research and agency guidance. Early work on recurring and nonrecurring congestion distinguished between congestion driven by predictable daily demand and congestion caused by incidents, weather, or temporary capacity restrictions such as lane closures [87,28]. Empirical studies have shown that lane reduction areas often experience higher crash rates, especially in transition and activity areas where drivers merge, weave, and travel in close proximity to workers and equipment [2,3].

Recent summaries from the Work Zone Safety Information Clearinghouse and FHWA report that although some long-term trends have improved, the absolute number of work zone-related fatalities remains high on high-speed facilities. Unexpected slowdowns and abrupt lane drops can create substantial speed differentials that contribute to severe crashes [2,3]. Work on the resilience of urban transportation networks has also highlighted that short-term nonrecurring congestion from incidents or lane closures can reduce network connectivity and increase travel time variability [87].

Within this context, reliable lane closure information serves both short-term and long-term purposes. It supports immediate operational responses such as queue warning and enforcement, and it enables longer-term evaluations of work zone performance, development of mitigation strategies, and planning of future maintenance programs.

5.1.2.2. Traditional Lane Closure Monitoring and Emerging Data Standards

Historically, lane closure information has been derived from planned maintenance schedules, permits, inspection reports, and TMC logs. These sources rely on manual updates and may not reflect real-time conditions when closures are delayed, extended, or moved without corresponding database updates [1,28].

The WZDx specification was developed to improve data consistency and shareability by defining an interoperable format for publishing work zone and lane closure events [28]. Pilot implementations have shown that standardized feeds can help navigation providers, researchers, and agencies use work zone information more easily, but they still depend on accurate upstream reporting. Investigations of connected work zone devices, including smart arrow boards and low-cost sensors, show that onsite devices can automatically broadcast closure status using standardized messages [28,29]. The main limitations are the hardware and maintenance requirements, which constrain full coverage, especially on short-duration or rural closures.

These constraints motivate additional approaches that infer lane closure conditions directly from traffic behavior. Such methods can provide an independent behavior-based view that complements device-based and manually reported data.

5.1.2.3. Connected Vehicle Data for Work Zone and Lane Closure Monitoring

CV probe data have evolved from coarse travel time samples to detailed trajectories with one- to three-second sampling. Large-scale evaluations have examined coverage, speed bias, congestion detection, and representativeness relative to traditional sensors [88,14,10,34]. Results indicate that, even at penetration rates on the order of several percent, CV telemetry data can reproduce freeway speed and congestion patterns with sufficient fidelity for operational use in many settings.

Work zone applications of CV telemetry data include performance measurement using probe speeds, estimation of delay and queue length, and analysis of CV-based metrics such as hard braking and speed profiles for safety and mobility evaluation [14,35,89]. The study in [35] demonstrated that CV data can support near-real-time monitoring of work zone traffic operations without relying on fixed detectors. Studies using naturalistic driving datasets have examined work zone mobility and safety and highlighted the behavioral detail observable in temporary traffic control environments [89,62]. Additional work has reviewed real-world CV applications in traffic state estimation, incident detection, speed harmonization, and safety analytics [88].

Lane closure detection is more specific than general traffic monitoring because closures may not always produce dramatic changes in average speed. Closure detection therefore benefits from features that capture both longitudinal speed and lateral positioning rather than speed alone.

5.1.2.4. Machine Learning and Deep Learning for Traffic State Estimation and Anomaly Detection

The growth of transportation data has encouraged the use of machine learning and deep learning for traffic state estimation, prediction, and anomaly detection. Physics-informed learning frameworks that incorporate macroscopic models such as the LWR model and the CTM have been applied to estimate density and flow under sparse and noisy measurements [90]. Spatial and temporal neural architectures including convolutional networks, graph neural networks, and recurrent networks have been studied for network-wide traffic prediction and for modeling spatial dependencies [91,92].

Fully data-driven deep learning methods can require large labeled datasets and significant computational resources [90,91,92]. For lane closure detection, where labeled closure information is relatively limited and interpretability and efficiency are important, hybrid approaches that use engineered features with classical machine learning methods are an attractive option [1].

The Kolmogorov-Smirnov test and related distributional distance measures are widely used to quantify differences between empirical CDFs and have been applied in anomaly detection and traffic analytics [93]. SVM and KNN classifiers remain strong baselines for low-dimensional nonlinear classification tasks when features are carefully designed [94,95,96].

5.1.2.5. Research Gap and Contribution

Despite progress in probe data analytics, relatively few studies have described a complete framework for automatic lane closure detection that operates at the spatial and temporal scales required for statewide deployment. Many prior efforts have focused on individual sites, used simulated data, or emphasized general speed-based congestion or incident detection rather than the explicit identification of lane closure activity [88,97].

A CV-based lane closure detection system is proposed that uses distributional statistics of speed and lateral position at a 0.1-mile resolution and integrates data processing, feature extraction, classification, validation, filtering, and visualization into a single pipeline [1]. The remainder of this section describes the adoption of that methodology as a lane closure detection module within the broader work zone monitoring system described in this report.

5.1.3. Connected Vehicle Data and Spatial Mapping

The lane closure detection system uses a high-frequency CV telemetry dataset containing anonymized vehicle waypoints. For this project, data were obtained from the commercial provider Wejo, which supplies one- to three-second movement records for a substantial portion of the vehicle fleet on Iowa freeways [1,34]. Each record includes latitude, longitude, timestamp, speed, heading, and a journey identifier that links points from the same trip.

To interpret these points in the context of the roadway, the trajectories are aligned with the Iowa DOT LRS [34]. The LRS includes route identifiers, start and end mileposts, bearings, speed limits, number of lanes, and county information. Many LRS segments are longer than two miles, which is too coarse for lane closure detection because closures are often localized near tapers, bridges, or short maintenance areas.

To increase spatial resolution, each long LRS segment is subdivided into analysis segments of lengths of at most 0.1 miles while preserving attributes (Figure 5.15). This subdivision creates a dense set of short segments that cover the freeway network. Each CV waypoint is projected onto the nearest reference line segment and assigned to one of these 0.1-mile units.

Figure 5.15. Color-coded representation of the segments present on the road from the LRS data on US 30

This mapping step links each CV record to a specific segment, attaches static attributes, and defines the spatial framework for computing weekly behavioral distributions. The 0.1-mile units are short relative to typical work zone extents and are therefore suitable for detecting localized lane closure effects within longer corridors.

Preprocessing and Lateral Position Estimation

Lateral positioning is a key signal for lane closure detection because it reflects how drivers distribute themselves across available lanes. When the rightmost lane is closed, for example, vehicles shift toward interior or left lanes and the lateral distribution of waypoints changes. Speed reductions alone can be caused by various factors, but a systematic lateral shift is more directly related to changes in lane availability.

To quantify lateral behavior, each CV waypoint is assigned a DFRL. The reference line is the LRS centerline, which approximates the geometric center of the roadway (Figure 5.16). For each waypoint, the shortest perpendicular distance to the reference line is computed. By convention, positive values indicate points on one side of the reference line and negative values indicate points on the opposite side (Figure 5.17) [1].

Figure 5.16. Colored lines representing the reference line on a divided highway, which was used for the DFRL calculation

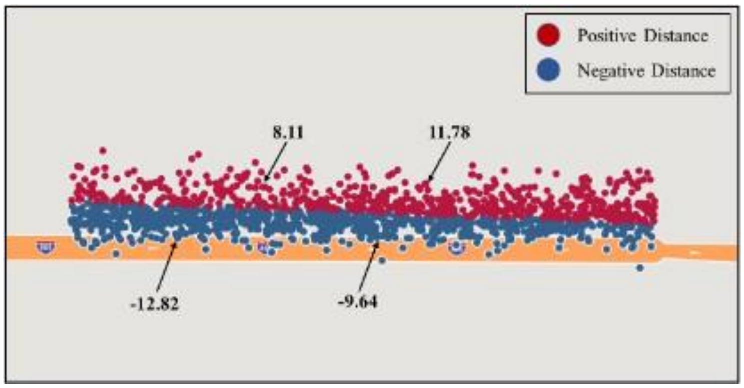

Figure 5.17. Waypoints representing DFRL, with red points indicating presence on the right side of the reference line and blue points indicating presence on the left side

The DFRL computation uses geometric operations on polylines. Each CV point is projected onto the nearest polyline segment, and the perpendicular offset is measured to obtain DFRL. Points that fall well outside the roadway envelope, such as those for on-ramps or frontage roads, can be removed using bounds on DFRL based on known cross section widths. This filtering focuses the analysis on mainline behavior.

After cleaning, DFRL values and speeds are aggregated by week and by 0.1-mile segment. Weekly aggregation provides enough observations to construct stable empirical distributions in segments with moderate probe penetration while providing a short enough timeframe that lane closures lasting several days appear as distinct changes relative to baseline weeks.

5.1.4. Weekly Traffic Behavior Through Distributional Signatures

After weekly speed and DFRL values are computed for each segment, their distributions can be examined over time. Under normal operations, weekly distributions of speed and lateral position for a given segment tend to be similar from week to week because drivers use available lanes and maintain free-flow speeds in a consistent manner [1].

To capture this behavior, empirical CDFs of speed and DFRL are constructed for a baseline week and for each subsequent week. The baseline week corresponds to a period with no known lane closures. When a lane closure is present in a later week, both CDFs change. The DFRL CDF shifts as vehicles move laterally into the remaining lanes, and the speed CDF shifts toward lower speeds as drivers slow down near the closure.

Under closure conditions, these changes are consistent for the affected segments. Weeks that do not contain closures show CDFs that closely match the baseline apart from sampling variation, while closure weeks show noticeable divergence in both speed and DFRL. The pattern is demonstrated by comparing baseline and closure week CDFs on US 30 westbound (Figure 5.18, with a shift in waypoints illustrated in Figure 5.19) [1]. To convert these observations into numerical features, the Kolmogorov-Smirnov *D* statistic is used. For each segment, the *D* statistic is computed between the baseline week CDF and the current week CDF for speed and for DFRL. The *D* statistic is the maximum vertical difference between two CDFs and is a standard measure of distributional change [93]. This process yields a pair of weekly *D* statistics per segment: one for speed and one for DFRL. Under normal conditions, both values remain small. During closure weeks, both tend to increase, reflecting the combined effect of speed reduction and lateral position shifts.

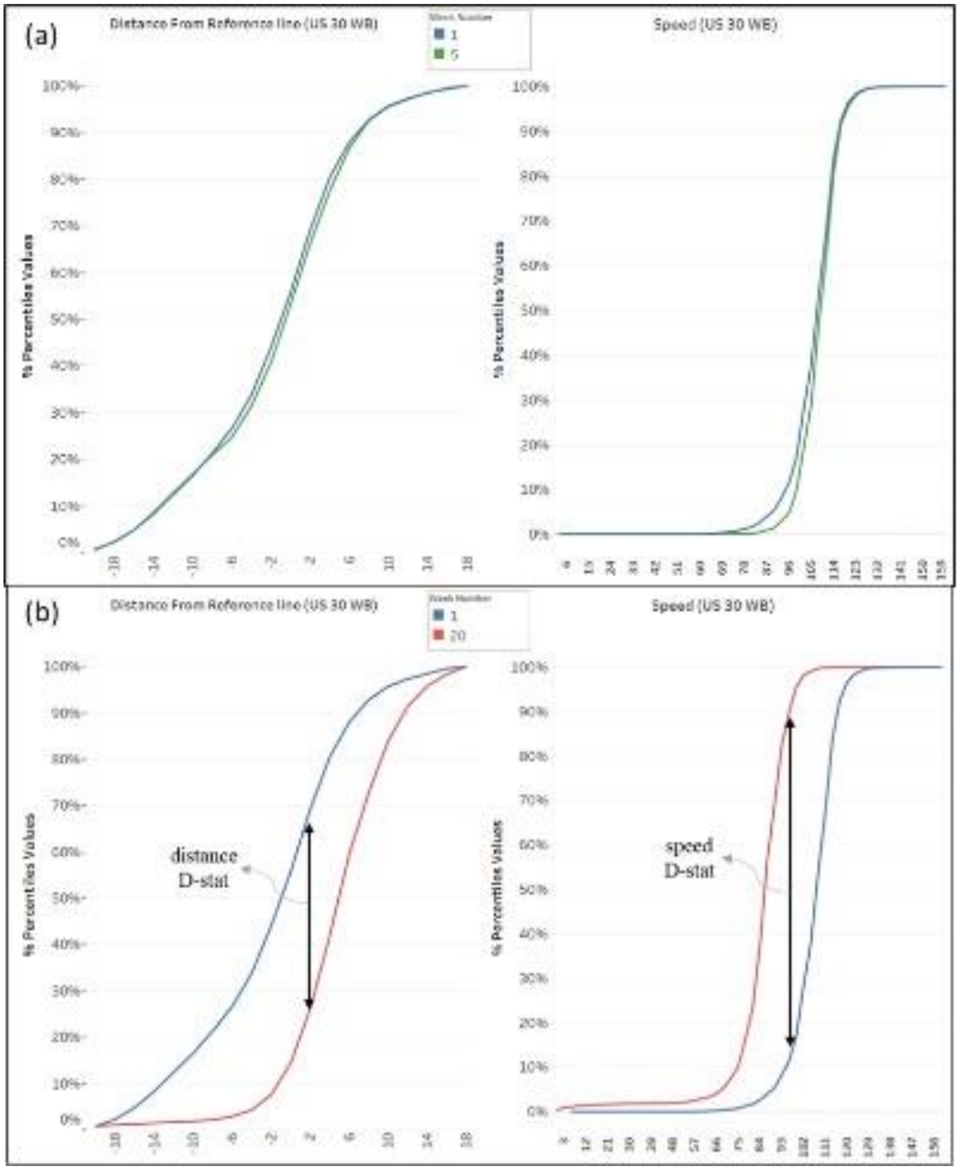

Figure 5.18. CDF plot of DFRL and speed for US 30 WB: (a) comparison of weeks 1 and 5 when both weeks have no closure, (b) plot for weeks 1 and 20 when only week 20 has a closure

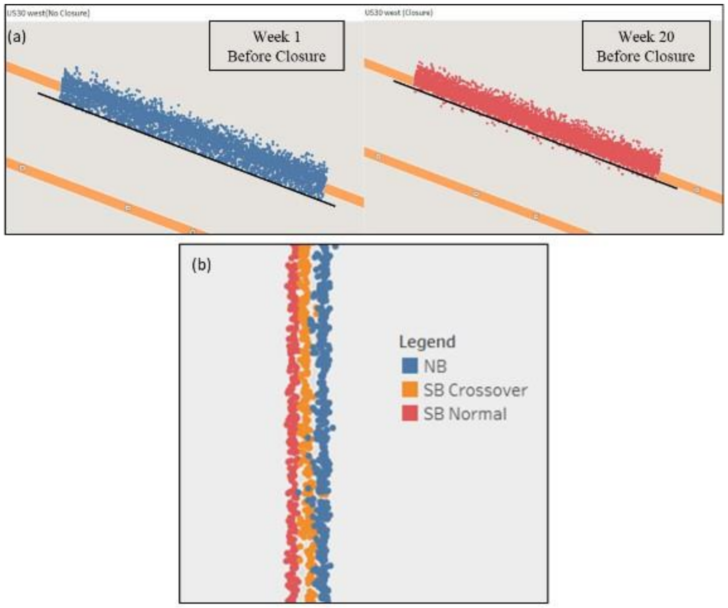

Figure 5.19. Shift of waypoints on road segments before and during closure: (a) Week 20 closure waypoints on US 30 westbound (red) compared with a baseline week, (b) US 65 southbound waypoints highlighting shift during closure

5.1.5. Machine Learning Classification of Lane Closures

With weekly *D* statistics computed for each segment, lane closure detection is framed as a binary classification problem. Each segment week is labeled as closure or non-closure using ground truth derived from Iowa 511 lane closure records.

When the weekly speed and DFRL *D* statistics are plotted in a two-dimensional feature space, closure and non-closure weeks occupy different regions (Figure 5.20). Weeks with confirmed lane closures cluster where both statistics are relatively large. Weeks without closures cluster near the origin, where both statistics are small. This separation is clearly visible for the US 30 training corridor [1].

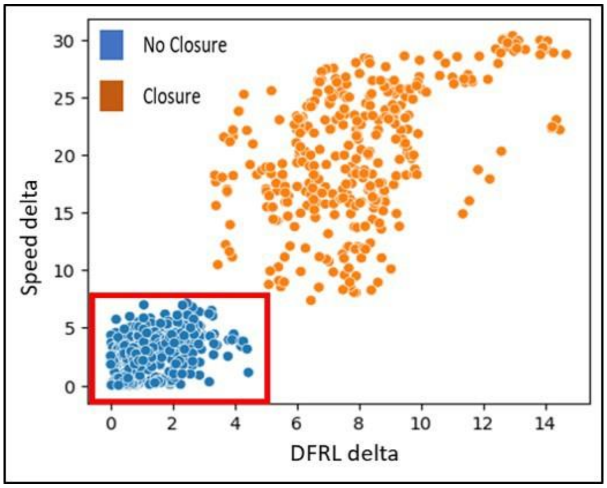

Figure 5.20. Scatterplot between speed *D* statistic and distance *D* statistic illustrating separation between closure and non-closure weeks

Two supervised learning methods are used: SVM with a radial basis function kernel and KNN (Figure 5.21). SVM constructs a nonlinear decision boundary in the feature space by finding a separating surface with maximum margin between classes [94,95]. KNN assigns labels based on the majority class among the nearest neighbors and is suitable when classes form coherent local clusters [96].

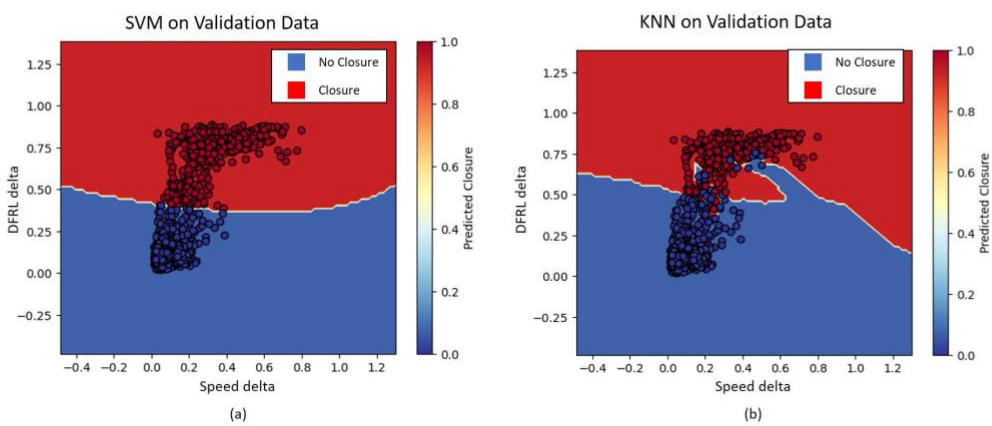

Figure 5.21. Decision boundary plots for validation datasets: (a) SVM decision boundary, (b) KNN decision boundary

The training dataset includes 21 weeks of data for 205 short segments on US 30, for a total of 4,305 segment week observations [1]. Lane closure weeks are less frequent than non-closure weeks, which creates class imbalance. To reduce bias toward the majority class, the Synthetic Minority Oversampling Technique (SMOTE) is applied to increase the representation of closure weeks by interpolating new samples from existing closure points in the feature space [98]. Model hyperparameters for SVM and KNN are tuned using a grid search with *k* fold cross validation. Performance metrics including accuracy, precision, recall, and F1 score are monitored to identify parameter settings with a reasonable balance between false positives and false negatives.

Validation accuracies of approximately 97.7% for SVM and 95.8% for KNN were reported, indicating that the weekly speed and DFRL *D* statistics form an effective feature set for lane closure detection on the training corridor [1].

5.1.6. Validation Across Additional Highway Corridors

To check that the method is not specific to a single roadway, the trained models are applied to two additional corridors: I-880 westbound and US 65 northbound. These corridors differ from US 30 in lane configuration, geometry, speed limits, and demand levels. Applying the models without retraining tests whether closure signatures in the feature space are consistent across different sites.

Ground truth lane closure data for I-880 and US 65 are again obtained from Iowa 511. The same preprocessing pipeline is used to map CV points to 0.1-mile segments, compute weekly speed and DFRL distributions, and derive *D* statistics [1,34].

When the validation corridor observations are plotted in the speed and DFRL *D* statistic space, closure and non-closure weeks show patterns similar to those on US 30. Closure weeks cluster in regions with larger *D* statistics, and non-closure weeks remain near the origin. The decision boundaries learned from US 30 generalize well, with only a small number of misclassified segment weeks on the validation corridors [1].

This validation suggests that the behavioral response of drivers to lane closures, as reflected in changes in lateral position and speed distributions, is consistent enough across freeways to support corridor-independent lane closure detection.

5.1.7. Statewide Upscaling and Computational Optimization

Extending the methodology to a statewide scale introduces computational challenges. The Iowa freeway network has thousands of LRS segments, which translate into many tens of thousands of 0.1-mile analysis units. When combined with multiple weeks of data, the number of segment-week combinations can reach into the hundreds of millions [1]. Computing full CDFs and Kolmogorov-Smirnov statistics for every segment week is expensive.

To address this, the methodology in [1] uses a two-stage filtering approach. In the first stage, the system computes differences at selected percentiles of the weekly speed and DFRL distributions relative to a baseline week, such as changes at the 10th, 50th, or 90th percentiles. These percentile deltas are faster to compute than full CDFs and provide approximate indicators of distributional change.

Empirical analysis shows that closure weeks exhibit larger percentile deltas than non-closure weeks for both speed and DFRL. Thresholds are defined for these percentile differences. Segment weeks that do not exceed any thresholds are unlikely to contain closures and are filtered out. Only segment weeks with larger percentile changes proceed to full CDF construction and Kolmogorov-Smirnov *D* statistic computation.

Large-scale data processing tools such as Amazon Athena and AWS SageMaker are used to implement this pipeline across statewide data [1]. This architecture provides a scalable environment that can allocate more compute resources when needed while keeping costs manageable.

5.1.8. Postprocessing for Temporal Stability and Spatial Consistency

Even with a well-performing classifier, some segment-level predictions can be noisy, especially in low-volume areas or near the decision boundary. This noise can appear as isolated closure predictions or frequent switches between closure and non-closure labels that do not match typical work zone practice. Postprocessing is therefore used to improve temporal stability and spatial consistency.

Segments are grouped into spatial clusters using Geohash or similar spatial indexing. Within each cluster, closure predictions are examined over time. Actual multi-week closures should appear as continuous or nearly continuous sequences of closure labels along adjacent segments. Isolated closure weeks surrounded in space and time by non-closure weeks are more likely to be artifacts of data sparsity or classification noise.

The methodology in [1] defines an instability score that measures how often closure labels change over time within each spatial cluster. Clusters with high instability are smoothed by applying continuity constraints. Short, isolated closure bursts may be relabeled as non-closures, and brief non-closure gaps within longer closure periods may be filled to reflect a continuous closure. This smoothing reduces instability scores and yields more realistic closure patterns across space and time.

5.1.9. Visualization and Operational Integration

The final step is to present lane closure detections in a form that supports agency operations. An interactive Tableau dashboard was developed in [1] to display closure predictions across Iowa's statewide freeway network (Figure 5.22). Each predicted closure is plotted along the network, and users can select a location to view metadata, weekly *D* statistics, model outputs, and closure timelines.

Figure 5.22. Tableau dashboard showing lane closure locations, with red dots indicating detected lane closure locations: (a) snapshot of traffic camera video feed near Omaha, (b) snapshot of traffic camera feed near Iowa City, (c) zoomed-in view of segment-level closure effects

For transportation agencies, such dashboards support situational awareness. Operators can scan for active closures, compare detected closures with planned maintenance schedules, and identify locations where closure activity is not documented in existing records. When dashboards are linked with camera feeds, staff can visually confirm closures detected by the CV-based system and refine thresholds as needed.

Closure detection outputs also provide inputs to other modules of the broader work zone monitoring system described in this report. Detected closure extents and durations can be used by the queue length estimation and speed heatmap tools, which rely on the same CV telemetry data. Closure start and end times can be used to align crash and naturalistic driving data for detailed safety analysis near closure locations [62,89]. In future applications, closure detections could be linked to WZDx feeds or used to support real-time traveler information.

5.1.10. Conclusion and Future Research Directions

This section outlined an automatic lane closure detection framework that uses CV trajectories, distributional statistics, and machine learning to infer closure activity at a 0.1-mile resolution on a statewide freeway network. The method described in [1] combines spatial mapping using LRS-based segmentation, DFRL-based lateral positioning, weekly CDF modeling of speed and lateral positioning distributions, Kolmogorov-Smirnov based feature extraction, supervised classification with SVM and KNN, validation on multiple corridors, statewide filtering, postprocessing, and dashboard visualization.

The results reported in [1] show that lane closures create consistent signatures in both lateral positioning and speed distributions and that these signatures can be captured with relatively simple distributional metrics. The use of CV data avoids the need for additional roadside hardware at every site and is therefore suitable for large-scale monitoring. The two-stage filtering and cloud-based processing framework keeps computation manageable while maintaining high detection performance.

Future research directions include examining shorter aggregation windows, such as daily or hourly, to improve responsiveness for short-duration or moving closures. Additional behavioral features such as hard braking, acceleration variability, and lane change frequency could complement speed and lateral positioning distributions and help distinguish among different types of lane restrictions [35,89]. Spatial and temporal deep learning models that process trajectories or gridded representations may capture more complex patterns, but their computational cost and data requirements will need to be evaluated relative to operational needs [90,91,92].

Another research direction involves integrating lane closure detection with incident detection and queue warning systems to create a combined operational picture. Linking closure signatures with camera analytics, WZDx feeds, and naturalistic driving studies could support joint analysis of closure implementation, driver response, and safety outcomes [89,62]. As CV penetration increases and additional probe datasets become available, the sensitivity and robustness of closure detection methods are likely to improve, allowing finer spatial detail and closer alignment with specific work zone features such as taper starts and barrier transitions.

When automatic lane closure detection is incorporated into the AI-based work zone monitoring system described in this report, agencies gain a behavior-based tool that can complement planned closure records and device-based data. The method provides a scalable way to monitor lane closure activity across a wide freeway network using information already available from CVs.

CHAPTER 6. INTEGRATED TOOL DASHBOARD

6.1. Introduction

The emergence of high-resolution CV data, cloud computing frameworks, and large-scale geospatial analytics has created new opportunities to modernize how work zones are monitored and managed. Traditional monitoring approaches have relied heavily on fixed detectors and manual observations, both of which provide only partial visibility into the changing conditions within active construction zones. With the continued rise of mobile operations, short-duration closures, and variable contractor schedules, agencies increasingly require systems that deliver real-time, spatially accurate, and behaviorally rich information. The integrated tool dashboard developed in this project (Figure 6.1) was designed to meet that need by consolidating multiple analytical pipelines into a single operational platform.

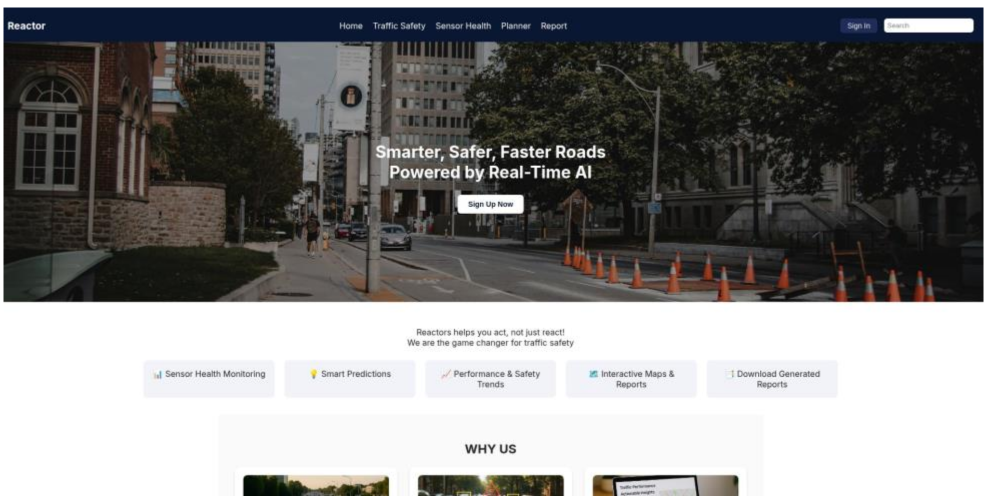

Figure 6.1. ReactorIQ dashboard landing page showing traffic safety modules, predictive tools, and real-time monitoring interfaces

The dashboard functions as the main user interface for accessing CV performance metrics, lane closure detection outputs, queue warning indicators, camera-based analytics, and machine learning results. Rather than treating these components as separate systems, the dashboard brings them together in a single decision-support environment. The aim is not only to visualize data but to convert CV telemetry into information that supports work zone safety, mobility, and operational decisions. The dashboard reflects a broader shift toward data-driven transportation operations, where decisions are informed by continuous streams of vehicle movement data that capture the behavior of road users [99].

Integrated dashboards are now common in TMCs, where operators must synthesize information from many sensors, data feeds, and predictive tools. Prior research has shown that dashboards can improve situational awareness and reduce operator workload, especially during periods of congestion or incidents [100]. The dashboard presented here extends that concept to work zone monitoring, where construction-related bottlenecks and lane closures introduce additional challenges. By displaying multiple analytic outputs in one interface, the dashboard helps users understand not only what conditions are present but also the traffic patterns that led to those conditions.

6.2. Conceptual Foundations of the Dashboard

The integrated tool dashboard is based on three main design principles.

The first principle is the fusion of geospatial context with CV telemetry. Unlike fixed detectors, CV data provide continuous measurements along the full roadway alignment. This allows the dashboard to represent performance metrics directly on the roadway rather than at a small number of sensor locations. Earlier studies have shown that spatially distributed probe data can improve the fidelity of congestion monitoring compared with traditional sensor-based systems [65].

The second principle is the integration of analytical outputs with visualization. The dashboard overlays the results from automatic lane closure detection, queue identification logic, and speed turbulence indicators directly on a map and related plots. Users do not need to move between multiple tools to interpret model outputs, which is consistent with recommendations that analytical capabilities should be embedded in visual platforms for transportation analytics [68].

The third principle is alignment with operational practice. Work zones change frequently as lanes open and close, crews move tapers, and traffic demand fluctuates throughout the day. The dashboard is designed to ingest updates at short intervals and present them in formats that are consistent with human-centered interface principles used in traffic management environments [101]. By matching data refresh cycles, layout, and interaction patterns with common workflows in operations centers (Figure 6.2), the dashboard is intended to function as a practical tool rather than a stand-alone research prototype.

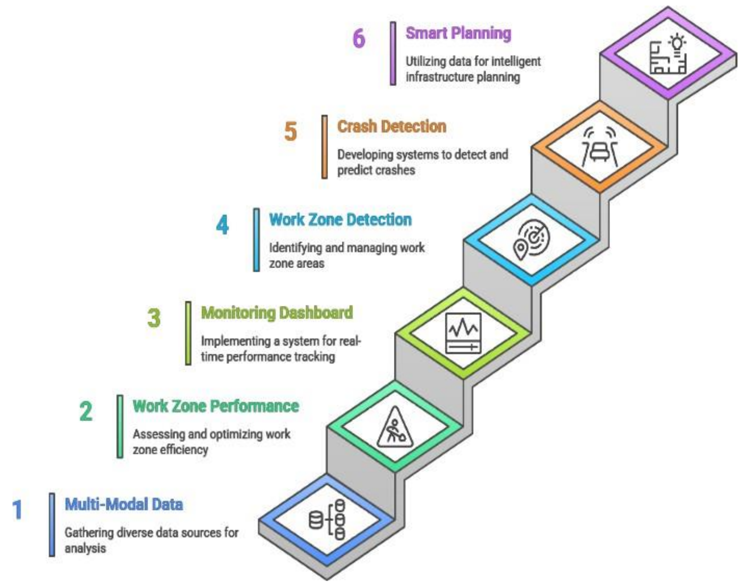

Figure 6.2. Conceptual framework illustrating the multi-stage data-driven workflow used in this study, with a staircase model highlighting six components: (1) multimodal data acquisition, (2) work zone performance assessment, (3) real-time monitoring dashboard development, (4) automated work zone detection, (5) crash detection and prediction, and (6) planning support for future infrastructure operations

6.3. Structure and Layout of the Integrated Dashboard

The dashboard uses a map-based layout to provide immediate situational awareness across statewide work zone activities. Active work zones appear as polygons overlaid on a basemap rendered with aerial or street-map tiles. Around each polygon, a color-coded performance ribbon summarizes the most recent 30 minutes of CV-based speed metrics. Green, yellow, orange, and red segments indicate approximate regimes of free-flow, moderate slowdown, severe slowdown, and queueing conditions, respectively [64]. This visual summary allows operators to quickly identify locations where conditions are deteriorating.

A left-hand panel provides text-based and drop-down filters for district, corridor, direction, and project attributes. Selecting a work zone from this list or clicking a polygon on the map triggers a zoom to that location and reveals more detailed performance statistics, including time series plots, heatmaps, and summary metrics. This details-on-demand interaction model follows established usability guidelines for map-based analytics platforms, where high-level information is always available and more detailed views can be accessed through simple, context-aware actions [102].

6.4. Integration of Connected Vehicle Performance Metrics

The dashboard is driven by CV telemetry processed through a Spark-Sedona-S3 cloud pipeline. Each CV point includes speed, timestamp, location, and journey identifiers. These points are mapped to roadway segments, cleaned, grouped into five-minute windows, and aggregated into performance measures.

Within each work zone polygon and for each five-minute bin, the dashboard displays core metrics derived from these aggregated data. Key measures include average speed, the share of vehicles operating in reduced-speed regimes, the frequency of slowdowns, estimated queue length, patterns of performance change over time, and counts of distinct journeys traversing the segment. These indicators together describe current operating conditions.

Speed-segment heatmaps provide spatial insight into bottleneck formation and congestion dynamics. In these plots, the horizontal axis represents time, the vertical axis represents distance or segment index, and color encodes mean segment speed or a selected percentile. As the heatmaps update throughout the day, operators can observe when congestion begins, how far it extends, and when recovery occurs, and can identify segments where slowdowns occur repeatedly [103].

6.5. Dashboard Interpretation of Work Zone Congestion

To support interpretation of congestion dynamics, the dashboard combines several complementary visualizations. The top panel typically shows a real-time speed time series for the selected work zone or segment (Figure 6.3). This plot helps identify breakdown onset, recovery timing, and recurring congestion across the day. Previous work has emphasized the diagnostic value of such time series plots for understanding traffic state transitions and variability [104].

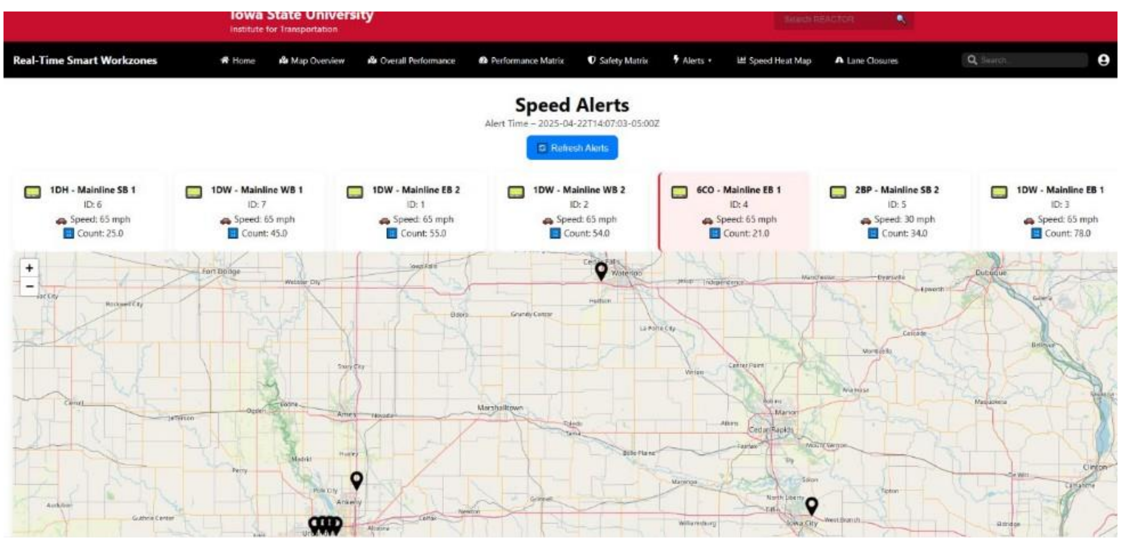

Figure 6.3. Real-time speed alerts and CV activity

Below the time series, a CV scatterplot shows microscopic behavior. In this plot, time is on the horizontal axis and speed is on the vertical axis, with each point representing an individual CV observation. Dense vertical clusters of low-speed points correspond to queues or stopped traffic, whereas more diffuse point clouds indicate transitional or free-flow states. The combination of aggregate time series and microscopic scatter distributions helps operators distinguish between short, localized slowdowns and sustained congestion [77].

6.6. Integration of Vehicle Trajectory Plots

An additional feature of the dashboard is the integration of vehicle trajectory plots (Figure 6.4) constructed on Amazon EC2. These plots depict microscopic space-time traffic dynamics, with time on the horizontal axis and distance along the corridor on the vertical axis. Each individual vehicle appears as a slanted line. Steep trajectories indicate higher speeds, and flattening trajectories indicate slowdowns, queues, or stop-and-go behavior [24].

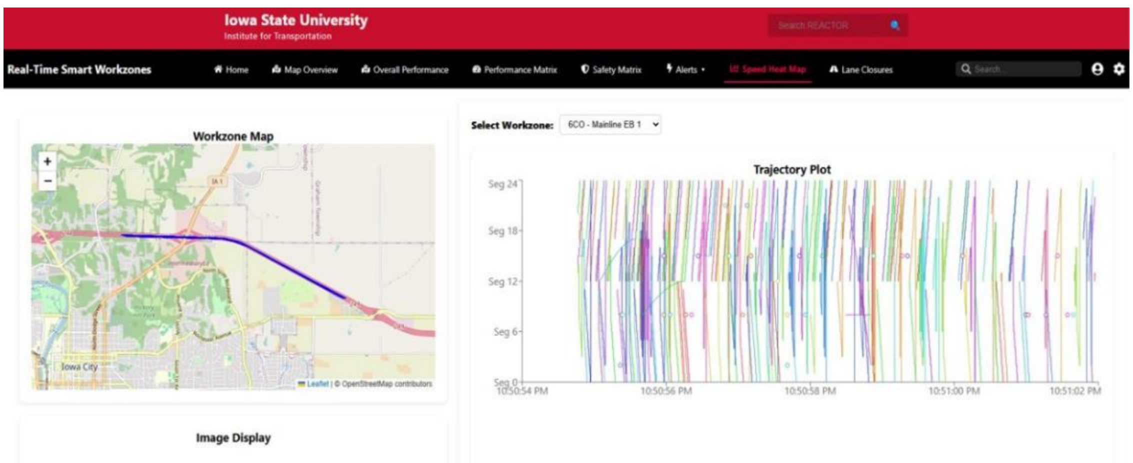

Figure 6.4. Real-time vehicle trajectories within a work zone

Trajectory patterns reveal several operational phenomena. Merging activity near tapers appears as abrupt changes in trajectory spacing and crossing patterns. Bottleneck activation is visible when many trajectories flatten near a specific location, indicating the formation of a stationary queue. Shockwave propagation appears as diagonal bands of slow trajectories moving upstream, and the duration of congestion is reflected in how long these bands persist. Empirical studies have shown that shockwaves often propagate upstream at characteristic speeds in typical freeway congestion [6]. Trajectory plots allow operators to visually confirm these patterns.

6.7. Overlay of Lane Closure Detection Results

Lane closure predictions from the machine learning classifier described in Chapter 5 are integrated into the dashboard. Detected closures appear as highlighted 0.1-mile segments overlaid on the work zone polygons and base map. When a closure is active, the affected segments are emphasized with distinct styling so that they are easy to locate.

This overlay allows operators to verify whether observed congestion patterns align with closure activity, to identify unreported or extended closures that may not yet appear in contractor logs or 511 systems, and to compare contractor-reported schedules with closure extents inferred from driver behavior. When closure predictions are viewed alongside speed heatmaps and trajectory plots, operators can more easily link geometric changes to mobility impacts.

6.8. Work Zone Diagnostics and Performance Interpretation

To help interpret the metrics, the dashboard includes a diagnostic component that summarizes the current state of each work zone in short narrative statements. Using CV-based speed patterns, queue metrics, and indicators of turbulence, the system identifies episodes such as unstable flow, rapid queue growth, emerging shockwaves, or recurring speed oscillations. These behaviors are then translated into brief text, for example, noting that a queue has formed upstream of the taper, that merging is creating speed variability, or that conditions have returned to stable flow.

Visual elements such as queue length timelines, delay statistics, and slowdown counts support this narrative. For each work zone, operators can view how maximum and average queue lengths change over time, how total delay accumulates, and how often slowdowns occur in different periods of the day. The combination of numerical metrics, plots, and narrative summaries provides a coherent picture of performance and helps operators decide whether observed patterns are consistent with expectations or indicate emerging issues.

6.9. Operator Workflow and Interactive Features

The dashboard supports an operator-centric workflow similar to that used in TMCs. Users typically begin with a statewide overview that shows all active work zones on the map with their performance ribbons. They can then filter by corridor, district, or route to focus on a smaller region and drill down into individual work zones through simple point-and-click interactions.

Within a selected work zone, operators can choose time windows of interest, such as a morning peak or a period containing an incident, and examine the corresponding metrics and visualizations. Tabs or panels allow switching between views that emphasize speed, queues, delay, or trajectories. Where camera feeds are available, they appear alongside plots so operators can visually confirm traffic conditions, queues, and events (Figure 6.5). This combination of analytics and video reduces the need to manage multiple separate systems.

Figure 6.5. Live traffic camera integration

6.10. Data Refresh Logic and Real-Time Synchronization

The dashboard operates on a continuous refresh cycle that is synchronized with the cloud processing outputs. CV telemetry is processed in near real-time, and aggregated performance summaries are written to Amazon S3 at regular intervals. The dashboard either polls these summaries or receives notifications and updates its displays every few minutes, depending on configuration and system load.

To avoid inconsistencies, versioned JSON files are used so that partially written data are not read by the dashboard. Each update produces a new, fully written object that is only loaded after the write operation finishes. Prior work on real-time dashboards recommends refresh intervals ranging from tens of seconds to a few minutes, depending on the balance between timeliness and computational overhead [105]. Versioned writes and careful file handling help maintain consistency and prevent operators from seeing incomplete snapshots [69].

6.11. Cloud-to-Dashboard Communication Layer

The dashboard communicates with AWS-hosted metrics through secure API endpoints that retrieve structured, timestamped performance summaries stored in S3 or related services (Figure 6.6). These endpoints expose aggregated CV metrics, lane closure predictions, trajectory plot references, and diagnostic summaries in compact JSON formats suitable for web applications.

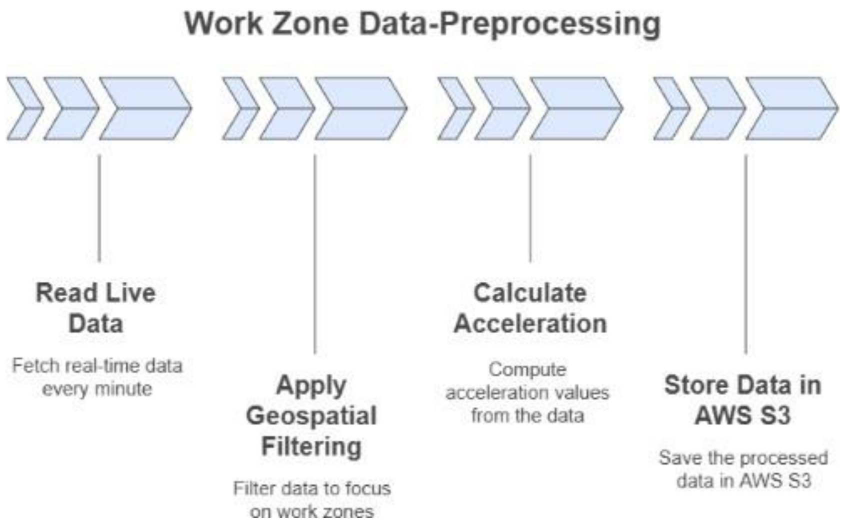

Figure 6.6. Cloud-to-dashboard communication architecture

The front-end dashboard is hosted on AWS Amplify, which supports scalability, integration with authentication services, and deployment from a backing code repository. Amplify manages static asset hosting, environment configuration, and continuous integration. Combined with services such as API Gateway, Lambda, and S3, this architecture forms a loosely coupled system in which compute, storage, and visualization components can develop independently while still communicating reliably.

6.12. Expanded User Interaction Features

Beyond basic filtering and drill-down, the dashboard includes several additional interactive features. Users can select custom time ranges such as specific incident windows, construction shifts, or peak periods. They can also view multiple work zones side by side to compare performance across districts and identify corridors with higher delays or queues.

The system supports user-defined performance thresholds. Operators can specify thresholds for metrics such as maximum queue length, average delay, or slowdown frequency. When real-time metrics exceed these values, the dashboard highlights the affected work zones, and automated alerts or notifications can be issued where configured. This allows agencies to encode performance-based management policies directly into the dashboard and receive immediate feedback when those policies are not met.

6.13. Advanced Data Analytics Integration

The dashboard is designed as a platform for integrating more advanced data analytics modules. In addition to the core metrics and lane closure detection, it can incorporate anomaly detection routines that flag unexpected deviations from typical speed or queue patterns. Predictive models for queue duration, trained on historical CV data, can provide short-term forecasts of how long current queues may persist. Lane utilization diagnostics estimate how traffic distributes across lanes, and pattern mining tools can highlight recurring congestion locations and times.

These models are trained on archival CV data using combinations of supervised and unsupervised learning to capture the structure of traffic behavior [75]. Once trained, they operate in near real-time, producing predictions and anomaly scores that are embedded in the dashboard visualizations. For example, if a work zone exhibits a queue pattern that differs from historical expectations, the anomaly detection module can highlight this condition and prompt additional review.

6.14. Real-Time Performance Metrics and KPI Tracking

KPIs are tracked within the dashboard to provide a concise assessment of work zone performance. Mobility-focused KPIs include queue length, average delay, and slowdown frequency. Safety-related KPIs can include counts or rates of events such as hard braking or severe speed drops when those data are available.

Performance summaries are displayed using color-coded indicators to support rapid interpretation. Green indicates conditions within defined limits, yellow indicates emerging concerns, and red indicates that one or more metrics have exceeded thresholds. In addition to real-time status, historical KPI time series are available so that users can review performance over an entire day or compare multiple days for the same work zone. This supports after-action reviews and planning activities.

6.15. Cloud-Based Scalability and System Redundancy

The dashboard uses a cloud-native design to handle statewide CV data streams and multiple concurrent users. Distributed ingestion of real-time CV data is managed through services such as AWS Glue and Spark, which can scale horizontally as volumes increase. Autoscaling groups adjust compute resources based on workload so that processing remains responsive while costs are controlled during off-peak periods.

Redundancy is supported through multi-availability-zone deployments and high-availability configurations for services such as S3, Lambda, and CloudWatch. Prior work on distributed systems and cloud-based ITS has shown that frameworks like Spark and Glue are suitable for large-scale analytics tasks of this type [106]. The dashboard builds on these capabilities to provide reliable monitoring under varying demand.

6.16. Data Security and Privacy Considerations

Data security and privacy are central to the dashboard design. All data transmitted between back-end services and the front-end application are encrypted using SSL/TLS. Data stored in cloud services such as S3 or managed databases are encrypted at rest using established standards such as AES-256. Authentication and role-based access control ensure that only authorized users can access sensitive dashboards or underlying metrics.

CV data are anonymized at ingestion. Journey identifiers are pseudonymized or hashed, and no personally identifiable information is retained. Data handling practices are designed to be consistent with current privacy frameworks such as the General Data Protection Regulation and the California Consumer Privacy Act, even if the deployment context is not directly governed by these regulations. Treating privacy protection as a core requirement helps build trust with agencies and the traveling public.

6.17. User Interface and Experience

The user interface follows user-centered design principles that emphasize clarity, simplicity, and responsiveness. The layout uses a central map view, concise metric panels, and context-aware overlays. Color-coded status indicators communicate conditions at a glance, while drill-down interactions allow users to explore specific locations and periods without being overwhelmed by detail.

Interactions are structured to reduce cognitive load. Common actions such as changing the selected work zone, switching between metrics, or adjusting the time window are performed through familiar controls such as drop-down menus, sliders, and clickable legends. Real-time responsiveness ensures that operators see current metrics and that the interface responds promptly to filter changes.

6.18. Data Visualization Best Practices

Visualization elements in the dashboard follow established best practices. Color gradients represent performance states, for example, mapping speed or delay to ordered scales. Time series plots convey temporal trends, scatterplots highlight microscopic behavior, and heatmaps summarize spatial patterns along corridors.

Interactive controls allow users to filter by metric, time, and space and to toggle overlays such as lane closure predictions or safety event markers. By following visualization principles that emphasize readability, consistent encoding, and careful use of color and scale, the dashboard supports exploratory analysis while remaining interpretable. Users can move between high-level summaries and detailed diagnostics as needed.

6.19. Future Enhancements and Integration with Other Systems

The integrated dashboard is intended as a flexible foundation for future enhancements and integration with other systems. One direction is the incorporation of predictive queue forecasting, where short-term traffic prediction models provide forward-looking estimates of queue length and delay. Another is incident detection based on trajectory anomalies, where abrupt speed changes or unusual lateral movements trigger incident flags.

Additional enhancements may include expanded camera-based analytics, such as automated detection of stopped vehicles, wrong-way movements, or near-crash events, and closer integration with downstream systems such as DMS, adaptive signal control, and WZDx feeds. In this configuration, the dashboard could serve as a coordination layer that not only visualizes conditions but also supports semi-automated responses. Literature on integrated corridor management and adaptive traffic control suggests that such multi-system integration can support more responsive traffic management strategies [40].

6.20. Conclusion

The integrated tool dashboard provides a unified platform for real-time work zone monitoring, performance visualization, and decision support. By combining CV data, machine learning outputs, geospatial analytics, and cloud infrastructure, the system supports operational awareness for agencies that manage work zone environments.

As probe penetration grows and analytical methods develop further, the dashboard can incorporate more advanced predictive models, automated alerts, and expanded integration with other systems. These enhancements can support more proactive work zone management, where decisions are informed by both current conditions and near-term forecasts. In this way, the dashboard serves as a practical operational tool and a foundation for future data-driven traffic operations.

CHAPTER 7. CONCLUSION, LIMITATIONS, AND FUTURE WORK

7.1. Conclusion

This project set out to design and implement a comprehensive statewide, real-time work zone intelligence system capable of transforming raw CV telemetry, camera feeds, and cloud-based analytics into actionable operational insights. The resulting platform, ReactorIQ, integrates data ingestion pipelines, performance metric computation, machine learning models, and intuitive dashboards to provide transportation agencies with a continuous understanding of traffic conditions across active work zones. Through this work, the project showed that CV data can serve as a strong foundation for modern traffic monitoring, achieving a level of spatial coverage and temporal resolution that traditional fixed sensors alone cannot provide.

A major accomplishment of the project was the creation of a cloud-native analytics architecture capable of processing millions of CV records per hour with near-real-time latency. Built on AWS Glue, Apache Spark, and Apache Sedona, the system filters and maps telemetry onto roadway geometries and generates performance indicators such as queue length, delay, speed degradation, and slowdown frequency. The integration of machine learning models further enhanced the platform's capabilities, particularly in detecting lane closures from subtle shifts in lateral driving patterns. These innovations, combined with real-time alerting and an interactive dashboard, provide operators with improved visibility into merging turbulence, bottleneck formation, shockwave propagation, and behavioral responses to work zone configurations.

The platform supports both operational and planning needs. In real-time, transportation agencies can identify emerging congestion, respond to safety anomalies, and verify incidents through live camera feeds. Over longer time horizons, the system provides a rich archive of performance data that can inform work zone design, taper placement, contractor compliance, and seasonal behavior trends. Overall, this project demonstrates that integrating CV data, cloud analytics, and AI-driven interpretation can significantly enhance work zone operations and create a scalable blueprint for the next generation of ITS.

7.2. Limitations

Despite substantial progress, several limitations remain in the current system, many of which stem from inherent constraints in CV data and the complexities of real-world deployment. One of the most significant challenges arises from variability in CV penetration. Urban areas and major Interstates often provide dense streams of telemetry, but rural corridors and low-volume regions suffer from sparse data, which can delay or weaken the detection of queues or lane closures. Although smoothing and clustering techniques mitigate these limitations, sensitivity in low-volume environments remains constrained by data availability.

GPS noise and spatial uncertainty represent another important limitation. CVs occasionally report inaccurate or jittered positional data, particularly in urban canyons, multi-level interchanges, or heavily vegetated areas. Sedona-based map matching reduces misalignment, yet lane-level interpretations may still be affected by GPS drift. Variations in CV telemetry data delivery patterns also create inconsistent latency, with some vehicles reporting at near real-time while others upload data in batches. These inconsistencies introduce challenges for strict real-time interpretation and may briefly delay detection during rapidly evolving congestion episodes.

The integration of camera streams also introduces operational constraints. Although hundreds of cameras are available statewide, the quality and reliability of RTSP streams vary considerably. Cameras with outdated hardware, limited bandwidth, or unstable network connections produce buffering and inconsistent performance. Because the platform relies on these feeds for incident verification, such limitations can temporarily reduce situational awareness. The machine learning models that detect lane closures and unusual driving behavior may also struggle to generalize under atypical conditions such as holiday surges, snowstorms, or multi-event disruptions, where driving patterns deviate significantly from training data and may require periodic recalibration.

Scalability and cloud cost considerations further shape system design. While cloud-based architectures offer flexibility, running high-frequency analytics pipelines on AWS Glue, S3, and Lambda incurs nontrivial expenses, particularly at statewide scale. Although performance optimizations were implemented, further expansion may require hybrid compute deployments or more efficient job scheduling strategies. Collectively, these limitations highlight areas where continued refinement is needed and where natural improvements in CV penetration and infrastructure reliability will strengthen system performance over time.

7.3. Future Work

Future work can extend the platform in several meaningful directions. One of the most promising opportunities involves predictive analytics. By leveraging historical traffic patterns, CV trajectories, and real-time conditions, the system could anticipate queue formation minutes in advance and support proactive traffic management strategies. Such predictive capabilities would enhance safety and congestion mitigation by enabling earlier interventions.

Another important direction is the enhancement of camera-based analytics. Integrating advanced video inference through modern deep learning models would allow the platform to detect stalled vehicles, wrong-way movements, near-miss interactions, and crash precursors directly from live footage. Combined with CV telemetry, this multimodal understanding of roadway conditions could improve event detection accuracy and reduce incident response times.

Refining the lane closure detection algorithm remains another priority. More complex closure scenarios, such as partial shoulder restrictions, dynamically shifting tapers, multi-lane closures, and short-term mobile operations, require more adaptable models capable of interpreting diverse lateral positioning patterns. Integrating camera confirmation and additional sensor streams would further strengthen classification accuracy and robustness.

The platform also offers opportunities to expand its decision-support capabilities. By combining performance metrics, safety indicators, and predictive analytics, the dashboard could automatically recommend message sets for DMS, identify candidate taper lengths, flag potential contractor deviations, or propose operational adjustments during severe congestion. Additional enhancements such as historical playback tools and multi-site comparison analytics would support planning groups by enabling them to benchmark performance across work zones, seasons, and contractor teams.

Finally, scaling the platform to regional or multi-state deployments represents a long-term goal. Achieving this expansion would require standardized data schemas, enhanced cloud infrastructure, improved redundancy, and advanced load balancing. As CV adoption increases and cloud-based computing becomes more cost-effective, the platform is well-positioned to support broader geographic deployments. Ultimately, future work should continue to refine, expand, and operationalize the system to ensure that it evolves alongside advancements in transportation technology, data availability, and agency operational needs.

REFERENCES

[1] A. Dumka, R. Kandiboina, A. Joshi, S. Knickerbocker, N. Hawkins, and A. Sharma, "An automated framework for lane closure detection on highway using connected vehicle data and machine learning models," *Journal of Advanced Transportation*, Vol. 2025, Article ID 9579025, 2025.

[2] Federal Highway Administration (FHWA), Work Zone Facts and Statistics, Office of Operations, U.S. Department of Transportation, Washington, DC, 2024. https://ops.fhwa.dot.gov/wz/resources/facts_stats.htm.

[3] O. Ozturk, K. Ozbay, H. Yang, "Estimating the impact of work zones on highway safety," in *Transportation Research Board 100th Annual Meeting Compendium of Papers*, 2014.

[4] C. Lu, J. Dong, A. Sharma, T. Huang, and S. Knickerbocker, "Predicting freeway work zone capacity distribution based on logistic speed–density models," *Journal of Advanced Transportation*, Vol. 2018, p. 9614501, 2018.

[5] C. Lu, J. Dong, A. Sharma, T. Huang, and S. Knickerbocker, "Predicting freeway work zone capacity distribution based on logistic speed–density models," in *Proceedings of the Transportation Research Board 97th Annual Meeting*, (Washington, DC), 2018.

[6] Y. Yuan, M. Treiber, and S. Knorr, "The impact of work zone taper length on traffic stability," *Transportation Research Part C*, Vol. 131, pp. 240–253, 2023.

[7] A. Sharma, D. M. Bullock, and J. A. Bonneson, "Input–output and hybrid techniques for real-time prediction of delay and maximum queue length at signalized intersections," *Transportation Research Record*, Vol. 2035, No. 1, pp. 69–80, 2007.

[8] S. P. Anusha, A. Sharma, L. Vanajakshi, S. C. Subramanian, and L. R. Rilett, "Model-based approach for queue and delay estimation at signalized intersections with erroneous automated data," *Journal of Transportation Engineering*, Vol. 142, No. 5, p. 04016013, 2016.

[9] S. P. Anusha, L. Vanajakshi, and A. Sharma, "A simple method for estimation of queue length," in *Proceedings of the Transportation Research Board 92nd Annual Meeting*, (Washington, DC), 2013.

[10] V. Ahsani, M. Amin Naseri, S. Knickerbocker, and A. Sharma, "Quantitative analysis of probe data characteristics: Coverage, speed bias, and congestion detection precision," *Journal of Intelligent Transportation Systems*, Vol. 23, No. 2, pp. 103–119, 2018.

[11] V. Ahsani, A. Sharma, C. Hegde, S. Knickerbocker, and N. Hawkins, "Improving probe-based congestion performance metrics accuracy by using change point detection," *Journal of Big Data Analytics in Transportation*, Vol. 2, No. 1, pp. 61–74, 2020.

[12] Y. O. Adu-Gyamfi, A. Sharma, S. Knickerbocker, N. Hawkins, and M. Jackson, "Framework for evaluating the reliability of wide-area probe data," *Transportation Research Record*, Vol. 2643, No. 1, pp. 93–104, 2017.

[13] T. Huang, P. Chakraborty, A. Sharma, and C. Hegde, "Large-scale data-driven traffic sensor health monitoring," *Journal of Big Data Analytics in Transportation*, Vol. 3, No. 3, pp. 229–245, 2021.

[14] R. Mudge, H. S. Mahmassani, R. Haas, A. Talebpour, and L. Carroll, "Work zone performance measurement using probe data," Tech. Rep. FHWA-HOP-13-043, Federal Highway Administration, 2013.

[15] A. Sharma, V. Ahsani, and S. Rawat, "Evaluation of opportunities and challenges of using INRIX data for real-time performance monitoring and historical trend assessment," Tech. Rep., Nebraska Department of Roads, 2017.
[16] J. S. Bourne, C. Eng, G. L. Ullman, D. Gomez, B. Zimmerman, T. A. Scriba, R. D. Lipps, D. L. Markow, K.C. Matthews, D. L. Holstein, R. Stargell, "Scan 08-04: Best practices in work zone assessment, data collection, and performance evaluation," Tech. Rep., National Cooperative Highway Research Program, 2010.
[17] J. Lee and F. Mannering, "Sensor coverage limitations in freeway operations," in *Proceedings of the Transportation Research Board 100th Annual Meeting*, (Washington, DC), 2021.
[18] R. Ogle, T. Lomax, and M. Levin, "Real time freeway performance monitoring using probe data," *Transportation Research Record*, Vol. 2675, No. 6, pp. 520–532, 2021.
[19] K. F. Sadabadi, M. D. Fontaine, and A. L. Goodin, "Maryland state highway administration work zone performance monitoring," Tech. Rep. MD-16-SHA-UM-3-16, Maryland State Highway Administration, 2016.
[20] M. Fontaine, "Detector health and reliability study," Tech. Rep., Virginia Transportation Research Council, 2022.
[21] A. Bochkovskiy *et al.*, "YOLOv5 and modern detectors in adverse environments," in *CVPR Workshops*, 2021.
[22] S. Sun, M. Levin, and R. Ogle, "Integrating camera analytics into real time traffic management systems," *Journal of Transportation Engineering*, Vol. 146, No. 10, 2020.
[23] H. Edara and S. Sun, "Analyzing work zone bottleneck dynamics using high resolution speed and probe data," *Journal of Transportation Engineering, Part A: Systems*, Vol. 149, No. 5, 2023.
[24] M. Beecroft and J. Laval, "Using trajectory data to improve work zone safety: Applications to queue formation and shockwave detection," *Transportation Research Part B: Methodological*, Vol. 124, pp. 78–90, 2021.
[25] FHWA, "Temporary traffic control monitoring guide," Tech. Rep., 2021.
[26] FHWA, "Automated queue warning system pilot results," Tech. Rep., 2023.
[27] FHWA, "Work zone operations best practices guidebook," Tech. Rep. FHWA-HOP-20-008, 2020.
[28] S. Knickerbocker, V. Ravichandra Mouli, and A. Venkatachalapathy, "Deploying and integrating smart devices to improve work zone data for work zone data exchange," *Transportation Research Record*, Vol. 2675, No. 10, pp. 317–327, 2021.
[29] G. Vorhes and D. Noyce, "Low cost connected work zone devices with simplified architecture," *Transportation Research Record*, Vol. 2676, No. 7, pp. 685–696, 2022.
[30] FHWA, "Traffic management data dictionary (TMDD) and message sets for external interfaces, version 3.03c," Tech. Rep., 2023.
[31] TRB, *Estimating the Safety Effects of Work Zone Characteristics and Countermeasures: A Guidebook*, NCHRP Report 869, National Academies of Sciences, Engineering, and Medicine, Washington, DC, 2018.
[32] J. K. Mathew, H. Li, J. Desai, R. S. Sakhare, E. Saldivar Carranza, M. Hunter, Benjamin Scholer, and D. Bullock, "Integration of probe data tools into traffic management center operations," Tech. Rep., Joint Transportation Research Program, Purdue University, 2021.
[33] S. Hallmark, A. Sharma, B. Lawton, G. Basulto-Elias, N. Oneyear, and T. Litteral, "Assessing driver behavior at back of queues: Implications for queue warning system in

work zones," in *Proceedings of the Transportation Research Board 99th Annual Meeting*, (Washington, DC), 2020.
[34] R. Kandiboina, S. Knickerbocker, S. Bhagat, N. Hawkins, and A. Sharma, "Exploring the efficacy of large scale connected vehicle data in real time traffic applications," *Transportation Research Record*, Vol. 2678, No. 5, pp. 651–665, 2023.
[35] R. S. Sakhare, J. Desai, H. Li, M. A. Kachler, and D. M. Bullock, "Methodology for monitoring work zones traffic operations using connected vehicle data," *Safety*, Vol. 8, No. 2, p. 41, 2022.
[36] GM OnStar, "Trajectory data for DOTs," Tech. Brief, 2023.
[37] Ford Safety Insights, "Connected vehicle penetration trends," Tech. Rep., 2022.
[38] Audi, "TLI penetration overview," Tech. Rep., 2024.
[39] R. Kandiboina, S. Knickerbocker, and A. Sharma, "Investigation of heavy acceleration events from connected vehicle data and their associated causes," in *Proceedings of the Transportation Research Board 104th Annual Meeting*, (Washington, DC), 2025.
[40] S. R. Turochy, J. Liu, and T. Smith, "Multimodal transportation management using integrated dashboards," *IEEE Transactions on Intelligent Transportation Systems*, Vol. 25, No. 7, pp. 1394–1408, 2024.
[41] H. Li, L. Peters, C. Banuelos, J. Zaugg, A. Sharma, and D. M. Bullock, "Leveraging snow plow dashboard cams and connected vehicle speed data to improve winter operations performance measures," in *Proceedings of the Transportation Research Board 98th Annual Meeting*, (Washington, DC), 2019.
[42] S. Wang, S. Knickerbocker, and A. Sharma, "Big-data-driven traffic surveillance system for work zone monitoring and decision supporting," in *Proceedings of the Transportation Research Board 96th Annual Meeting*, (Washington, DC), 2017.
[43] P. Chakraborty, C. Hegde, and A. Sharma, "Data-driven parallelizable traffic incident detection using spatiotemporally denoised robust thresholds," *Transportation Research Part C: Emerging Technologies*, Vol. 105, pp. 81–99, 2019.
[44] K. Aghabayk, M. Sarvi, and W. Young, "Density estimation using connected vehicle data," *Transportation Research Part C: Emerging Technologies*, Vol. 121, Art. No. 102844, 2020.
[45] Google Cloud ITS Team, "Cloud native traffic operations," 2022.
[46] M. J. Islam, A. Sharma, and H. Rajan, "A cyberinfrastructure for big data transportation engineering," *Journal of Big Data Analytics in Transportation*, Vol. 1, No. 1, pp. 83–94, 2019.
[47] S. Poddar, P. Chakraborty, A. Sharma, S. Knickerbocker, and N. Hawkins, "Massively parallelizable approach for evaluating signalized arterial performance using probe-based data," *Journal of Intelligent Transportation Systems*, Vol. 27, No. 4, pp. 488–502, 2023.
[48] S. Knickerbocker, V. Ravichandra-Mouli, N. Hawkins, and A. Sharma, "Developing a research-grade Iowa work zone database," Tech. Rep., Institute for Transportation, Iowa State University, 2020.
[49] M. Fontaine and C. Dennison, "Turbulence analysis in work zones," Tech. Rep., Virginia Tech Transportation Institute, 2022.
[50] S. Wang, A. Sharma, and S. Knickerbocker, "Analyzing and improving the performance of dynamic message sign reporting work zone–related congestion," *Transportation Research Record*, Vol. 2617, No. 1, pp. 71–77, 2017.
[51] IDOT, "Connected vehicle operations annual report," Tech. Rep., 2023.

[52] M. Amin-Naseri, P. Chakraborty, A. Sharma, S. B. Gilbert, and M. Hong, "Evaluating the reliability, coverage, and added value of crowdsourced traffic incident reports from Waze," *Transportation Research Record*, Vol. 2672, No. 43, pp. 34–43, 2018.
[53] P. Kachroo and A. Sharma, "Theory of safety surrogates using vehicle trajectories in macroscopic and microscopic settings: Application to dynamic message signs controlled traffic at work zones," *Transportation Research Part C: Emerging Technologies*, Vol. 91, pp. 62–76, 2018.
[54] X. Wei, S. Knickerbocker, N. Hawkins, Z. Hans, and A. Sharma, "Leveraging commercial connected vehicle data for high-speed signalized intersection safety: A quantitative analysis of hard-braking and crashes," in *2024 IEEE World Forum on Public Safety Technology (WFPST)*, 2024.
[55] R. Kandiboina, V. Ravichandra-Mouli, G. Basulto-Elias, S. Knickerbocker, N. Hawkins, and A. Sharma, "Real-time freeway crash detection framework using connected vehicle waypoint data," *Transportation Research Record*, Vol. 2679, No. 3, pp. 565–578, 2025.
[56] B. Yu, H. Yin, and Z. Zhu, "Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting," *Transportation Research Part C: Emerging Technologies*, Vol. 102, pp. 1–13, 2019.
[57] M. Behl, J. Wang, and Z. Huang, "Apache Sedona: Scalable geospatial processing for big spatial data," in *Proceedings of the ACM SIGSPATIAL International Workshop on Big Spatial Data (BigSpatial)*, 2023.
[58] G. Davis, "Validation of queue detection using probe vehicle data," Minnesota Department of Transportation (MnDOT), St. Paul, MN, 2021.
[59] J. Hourdos, B. Jacobson, and R. Carlson, "Characterization of microscale turbulence at freeway lane drops using probe vehicle data," *Transportation Research Record*, Vol. 2676, No. 2, pp. 117–128, 2022.
[60] G. Astarita, V. P. Giofr´e, and M. Vitetta, "Safety impacts of speed variability and acceleration noise in freeway bottlenecks," *Accident Analysis & Prevention*, Vol. 127, pp. 180–192, 2019.
[61] N. Bharadwaj, P. Edara, and C. Sun, "Risk factors in work zone safety events: A naturalistic driving study analysis," *Transportation Research Record*, Vol. 2673, No. 1, pp. 379–387, 2019.
[62] S. Turochy and J. Smith, "Evaluation of real time queue detection algorithms using high resolution speed data," *Journal of Transportation Engineering*, Vol. 146, No. 10, 2020.
[63] K. Vu, H. Rakha, and M. L. Van Aerde, "Comparison of vehicle trajectory and loop detector data for identifying congestion patterns," *Transportation Research Record*, Vol. 2673, No. 3, pp. 452–463, 2019.
[64] D. Banerjee, H. Rakha, and T. Chase, "Shockwave identification and classification using connected vehicle trajectories," *IEEE Transactions on Intelligent Transportation Systems*, Vol. 24, No. 2, pp. 1331–1343, 2023.
[65] C. Sun, R. Nazem, and K. Heaslip, "Queue detection performance of probe vehicle speeds under varying market penetration rates," *Transportation Research Record*, Vol. 2675, No. 9, pp. 341–352, 2021.
[66] S. Anderson and M. Wang, "Queue formation and dissipation dynamics in recurrent freeway bottlenecks: A probe data study," *Transportation Research Part C*, Vol. 148, 2023.

[67] A. Khattak, J. Liu, and K. Wang, “Real time traffic information systems and standardized data formats,” *IEEE Transactions on Intelligent Transportation Systems*, Vol. 23, No. 1, pp. 321–335, 2022.
[68] X. He, L. Song, and Z. Li, “Distributed processing of probe vehicle data for real time congestion detection,” *IEEE Intelligent Transportation Systems Magazine*, Vol. 15, No. 1, pp. 88–100, 2023.
[69] G. B. Dadi, R. E. Sturgill, Z. Al-Shabbani, “NCHRP Synthesis 591: Use of safety management systems in managing highway maintenance worker safety,” National Cooperative Highway Research Program, 2022.
[70] T. Acosta, L. Quintero, and R. Sanchez, “Standards-Based Data Exchange for Multimodal Traffic Management Systems,” in *Proceedings of the Transportation Research Board Annual Meeting*, Washington, DC, 2023.
[71] P. Singhal and G. Gardner, “Dashboard design for real time transportation data platforms,” *Journal of Transportation Technologies*, Vol. 13, No. 2, 2023.
[72] U.S. Department of Transportation, Volpe National Transportation Systems Center, “Cloud Cost Optimization Strategies for State Departments of Transportation,” Cambridge, MA, 2024.
[73] G. Leduc, L. Bouchard, and P. Maciejewski, “Interoperability challenges in traffic management systems,” *IEEE Systems Journal*, Vol. 17, No. 3, pp. 5662–5674, 2023.
[74] J. Mensah and C. Qiao, “Impact of stale data on real time traffic estimation accuracy,” *Transportation Research Part C*, Vol. 143, 2022.
[75] A. Bermbach and S. Tai, “Cloud storage metadata performance under high frequency access patterns,” in *ACM Symposium on Cloud Computing*, pp. 110–118, 2021.
[76] Y. Yuan, M. Treiber, and S. Knorr, “Heterogeneous traffic flow observed from trajectory data: Empirical evidence and implications,” *Physica A*, Vol. 533, 2019.
[77] L. Leclercq, “Data driven traffic flow analysis using vehicle trajectories,” *Transportation Research Part C*, Vol. 118, 2020.
[78] J. Laval and C. F. Daganzo, “Lane changing in traffic streams: Kinematic waves model and impacts on capacity,” *Transportation Research Part B*, Vol. 39, No. 7, pp. 709–728, 2005.
[79] H. Rakha, R. K. Kamalanathsharma, and M. F. Alluri, “Identification of traffic shockwaves using vehicle trajectory data,” *IEEE Transactions on Intelligent Transportation Systems*, Vol. 15, No. 2, pp. 618–632, 2014.
[80] R. Herman and I. Prigogine, “A two fluid approach to town traffic,” *Science*, Vol. 204, No. 4389, pp. 148–151, 1979.
[81] FHWA, “A primer on work zone safety and mobility performance measurement,” Tech. Rep. FHWA-HOP-11-033, 2011.
[82] C. M. Day, D. Okaidjah, and S.Knickerbocker, “Usefulness and reliability of probe data when alerting work zone drivers using portable message signs,” Tech. Rep., Smart Work Zone Deployment Initiative, Institute for Transportation, Iowa State University, 2024.
[83] S. Fulari, *Leveraging Connected Vehicles to Enhance Traffic Signal Performance Measures*. PhD dissertation, Purdue University, West Lafayette, Indiana, 2019.
[84] E. Saldivar Carranza, H. Li, J. Mathew, M. Hunter, J. Sturdevant, and D. Bullock, “Deriving operational traffic signal performance measures from high resolution data,” *Transportation Research Record*, 2021.

[85] W. Abewickrema, K. Somarathna, D. N. M. Tan, and S. Buddhika, “An ensemble deep learning framework for real time maximum queue length prediction,” *Journal of Big Data Analytics in Transportation*, 2025.
[86] M. A. Gorji, M. Akbarzadeh, and S. N. Shetab Boushehri, “Evaluation and improvement of the urban transportation networks resilience in short term nonrecurring traffic congestion: A novel graph connectivity based criteria,” *Transportation Engineering*, Vol. 10, p. 100152, 2022.
[87] Y. Zhou and R. Bridgelall, “Review of usage of real world connected vehicle data,” *Transportation Research Record*, Vol. 2674, No. 10, pp. 939–950, 2020.
[88] H. H. Zhou, R. E. Turochy, and D. Xu, “Evaluation of work zone mobility by utilizing naturalistic driving study data,” Tech. Rep., STRIDE Southeastern Transportation Research Center, 2019.
[89] A. J. Huang and S. Agarwal, “Physics informed deep learning for traffic state estimation: Illustrations with LWR and CTM models,” *IEEE Open Journal of Intelligent Transportation Systems*, Vol. 3, pp. 503–518, 2022.
[90] P. Brimos, A. Karamanou, E. Kalampokis, and K. Tarabanis, “Graph neural networks and open government data to forecast traffic flow,” *Information*, Vol. 14, No. 4, p. 228, 2023.
[91] C. Zhao and H. Yu, “Observer informed deep learning for traffic state estimation with boundary sensing,” *IEEE Transactions on Intelligent Transportation Systems*, Vol. 25, No. 2, pp. 1602–1611, 2024.
[92] P. W. Holland, “Two measures of change in the gaps between the CDFs of test score distributions,” *Journal of Educational and Behavioral Statistics*, Vol. 27, No. 1, pp. 3–17, 2002.
[93] M. Awad and R. Khanna, “Support vector machines for classification,” in *Efficient Learning Machines*, pp. 39–66, Springer, 2015.
[94] S. Yue, P. Li, and P. Hao, “SVM classification: Its contents and challenges,” *Applied Mathematics – A Journal of Chinese Universities*, Vol. 18, No. 3, pp. 332–342, 2003.
[95] P. Cunningham and S. J. Delany, “k nearest neighbour classifiers – a tutorial,” *ACM Computing Surveys*, Vol. 54, No. 6, pp. 1–25, 2021.
[96] P. Oza, N. Abhyankar, A. Seth, and G. Sharma, “Road incidents and closures classification in a connected vehicular environment,” Tech. Rep., 2021.
[97] N. Noorhalim, A. Ali, and S. M. Shamsuddin, “Handling imbalanced ratio for class imbalance problem using SMOTE,” in *Proceedings of the Third International Conference on Computing, Mathematics and Statistics (iCMS2017)*, pp. 19–30, 2019.
[98] Wejo, “Connected vehicle data insights report,” Tech. Rep., 2023.
[99] A. Jacobson, D. Swiggett, and B. Anderson, “Evaluating the effectiveness of traffic management dashboards in operational decision making,” *Transportation Research Record*, Vol. 2678, No. 3, pp. 34–48, 2023.
[100] S. R. Turochy, A. B. Anderson, and M. S. Levin, “Leveraging data visualization tools for real time decision support in work zone monitoring,” *Transportation Research Record*, Vol. 2676, pp. 120–132, 2021.
[101] T. H. Chase, A. B. Jacobson, and R. M. Levens, “Designing effective interactive maps for traffic management centers,” *Transportation Research Part C: Emerging Technologies*, Vol. 90, pp. 122–133, 2020.

[102] G. Gardner, M. L. Van Aerde, and M. D. Leman, “Using connected vehicle data for traffic flow diagnostics,” *IEEE Intelligent Transportation Systems Magazine*, Vol. 17, No. 1, pp. 32–45, 2023.
[103] K. B. Rakha, D. M. Alluri, and R. K. Kamalanathsharma, “Real time queue detection using connected vehicle probe data,” *Transportation Research Part C*, Vol. 148, 2023.
[104] Amazon Web Services, “Amazon S3: Strong read after write consistency,” Tech. Rep., AWS Documentation, 2023.
[105] Y. Chen and Y. Zhang, “Low latency processing of real time traffic streams with spark structured streaming,” in *IEEE International Conference on Big Data*, pp. 1451–1460, 2022.