Validation of Freight Volume Modeling on Major Highway Links

June 2024

A Research Report from the Pacific Southwest Region University Transportation Center

Chrysovalantis Anastasiou, Ph.D. Student
Seon Ho Kim, Researcher, Associate Director, IMSC
John Krumm, Researcher, Associate Director, IMSC

Cyrus Shahabi, Principal Investigator
Luciano Nocera, Co-Principal Investigator
Genevieve Giuliano, Co-Principal Investigator

University of Southern California

Pacific Southwest Region UTC

USC University of Southern California

TECHNICAL REPORT DOCUMENTATION PAGE

1. Report No. PSR-21-02	**2. Government Accession No.** N/A	**3. Recipient's Catalog No.** N/A
4. Title and Subtitle Validation of Freight Volume Modeling on Major Highway Links		**5. Report Date** 06/30/2024
		6. Performing Organization Code N/A
7. Author(s) Anastasiou: https://orcid.org/0000-0003-4966-5129 Kim: https://orcid.org/0000-0002-8410-0839 Krumm: https://orcid.org/0000-0003-4394-6704 Nocera: https://orcid.org/0000-0002-8100-4286 Giuliano: https://orcid.org/0000-0002-9257-8269 Shahabi: https://orcid.org/0000-0001-9118-0681		**8. Performing Organization Report No.** TBD
9. Performing Organization Name and Address METRANS Transportation Center University of Southern California University Park Campus, RGL 216 Los Angeles, CA 90089-0626		**10. Work Unit No.** N/A
		11. Contract or Grant No. [USDOT Grant, Caltrans 65A0674]
12. Sponsoring Agency Name and Address U.S. Department of Transportation Office of the Assistant Secretary for Research and Technology 1200 New Jersey Avenue, SE, Washington, DC 20590		**13. Type of Report and Period Covered** Final report (07/01/2023 – 06/30/2024)
		14. Sponsoring Agency Code USDOT OST-R
15. Supplementary Notes https://imsc.usc.edu:8103/datasets https://metrans.org/research/validation-of-freight-volume-modeling-on-major-highway-links https://doi.org/10.25554/7nts-we12		
16. Abstract The project aims at validating the feasibility of leveraging the existing infrastructure and existing and emerging sensors, including repurposed sensors such as CCTV cameras and WIM stations to estimate truck flow over a region of study situated north and east of the Ports of Los Angeles and Long Beach.		
17. Key Words CCTV, WIM, OD Matrix, Truck, Modeling, Highway		**18. Distribution Statement** No restrictions.

19. Security Classif. (of this report) Unclassified	**20. Security Classif. (of this page)** Unclassified	**21. No. of Pages** 29	**22. Price** N/A

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

Contents

About the Pacific Southwest Region University Transportation Center

The Pacific Southwest Region University Transportation Center (UTC) is the Region 9 University Transportation Center funded under the US Department of Transportation's University Transportation Centers Program. Established in 2016, the Pacific Southwest Region UTC (PSR) is led by the University of Southern California and includes seven partners: Long Beach State University; University of California, Davis; University of California, Irvine; University of California, Los Angeles; University of Hawaii; Northern Arizona University; Pima Community College.

The Pacific Southwest Region UTC conducts an integrated, multidisciplinary program of research, education and technology transfer aimed at *improving the mobility of people and goods throughout the region.* Our program is organized around four themes: 1) technology to address transportation problems and improve mobility; 2) improving mobility for vulnerable populations; 3) Improving resilience and protecting the environment; and 4) managing mobility in high growth areas.

U.S. Department of Transportation (USDOT) Disclaimer

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. This document is disseminated in the interest of information exchange. The report is funded, partially or entirely, by a grant from the U.S. Department of Transportation's University Transportation Centers Program. However, the U.S. Government assumes no liability for the contents or use thereof.

California Department of Transportation (CALTRANS) Disclaimer

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. This document is disseminated under the sponsorship of the United States Department of Transportation's University Transportation Centers program, in the interest of information exchange. The U.S. Government and the State of California assumes no liability for the contents or use thereof. Nor does the content necessarily reflect the official views or policies of the U.S. Government and the State of California. This report does not constitute a standard, specification, or regulation. This report does not constitute an endorsement by the California Department of Transportation (Caltrans) of any product described herein.

Disclosure

Principal Investigator Cyrus Shahabi, Co-Principal Investigators Luciano Nocera and Genevieve Giuliano, Chrysovalantis Anastasiou, Ph.D., student, Seon Ho Kim, and John Krumm researchers conducted this "Validation of Freight Volume Modeling on Major Highway Links" at IMSC, Viterbi School of Engineering, University of Southern California. The research took place from 07/01/2023 to 06/30/2024 and was funded by a grant from the USDOT in the amount of $100,000. The research was conducted as part of the Pacific Southwest Region University Transportation Center research program.

Acknowledgments

We wish to thank Caltrans for their support and input. Specifically, Deputy District Director Rafael Molina and Aung Than from Caltrans District 7, who helped with acquiring the CCTV data, Stan Norikane and Vicki Dickey from Traffic Operations Weigh-in-Motion who helped acquire WIM data, and Andre Tok from the Institute of Transportation Studies at UC Irvine who provided TAMS data. Our team would also like to thank Frank Law, Caltrans Task Order Manager, for his guidance and support.

Abstract

One of the most challenging problems in urban transportation planning is the lack of fine-grain data on freight movements. Cities and regions do not know how many trucks operate in the region and have only limited information on freight flows. Without a consistent and current source for freight volume and origin-destination data, it is difficult to manage or plan for freight in metropolitan areas. This project aims to develop methods to generate freight volume information, e.g., estimate hourly origin-destination counts (OD-matrices), from sensor observations. Available sensors include CCTV cameras to monitor the roadways, Weight-In-Motion Stations (WIM), and other available sensors such as Truck Activity Monitoring Systems (TAMS). Trucks must be detected and counted on the CCTV cameras, and truck observations with varying accuracy must be integrated in time and space. By continuously updating sensor data, we can generate fine-grained truck flow estimates on historical data and in a close to real-time fashion. With this work, we aim to validate the feasibility of detecting trucks in CCTV videos and estimating truck flow over the highway systems in a region of study most impacted by truck activity situated north and east of the Ports of Los Angeles and Long Beach.

Validation of Freight Volume Modeling on Major Highway Links

Executive Summary

One of the most challenging problems in urban transportation planning is the lack of fine-grain data on freight movements. Cities and regions do not know how many trucks operate in the region and have only limited information on freight flows. Without a consistent and current source for freight volume and origin-destination data, it is difficult to manage or plan for freight in metropolitan areas. This project aims to develop methods to generate freight volume information, e.g., estimate hourly origin-destination counts (OD-matrices), from sensor observations. Available sensors include CCTV cameras to monitor the roadways, Weight-In-Motion Stations (WIM), and other available sensors such as Truck Activity Monitoring Systems (TAMS). Trucks must be detected and counted on the CCTV cameras, and truck observations with varying accuracy must be integrated in time and space. By continuously updating sensor data, we can generate fine-grained truck flow estimates on historical data and in a close to real-time fashion. With this work, we aim to validate the feasibility of detecting trucks in CCTV videos and estimating truck flow over the highway systems in a region of interest (ROI) most impacted by truck activity situated north and east of the Ports of Los Angeles and Long Beach.

Our goal is to validate the feasibility of leveraging the existing infrastructure and existing and emerging sensors, including repurposed sensors such as CCTV cameras and WIM stations, to estimate truck flow. For this, we collected a dataset of CCTV, WIM, and TAMs sensor data over ROI and investigated the feasibility and usefulness of using traffic monitoring CCTV videos for truck detection and counting and created and studied novel truck flow estimation algorithms.

We have selected an ROI approximately 12 square miles in size north of the Ports of Los Angeles and Long Beach. This area was specifically chosen due to the significant impact of truck traffic on freight movement to and from major ports. The constant flow of trucks transporting goods from the ports to various distribution centers and end destinations contributes to substantial traffic congestion in the surrounding areas. This heavy truck traffic not only affects the efficiency of transportation logistics but also has significant environmental and health implications. Emissions from diesel trucks contribute to air pollution, which can lead to health issues for residents. The increased traffic volume also leads to greater wear and tear on infrastructure, necessitating more frequent repairs and maintenance.

To assemble the Truck Sensing Dataset, we worked with Caltrans District 7 to collect CCTV cameras, with the California Department of Transportation Traffic Operations Weigh-in-Motion to collect WIM station data, and with the UCI Institute of Transportation Studies to collect Truck Activity Monitoring System (TAMS) data.
Regarding sensor counts, the dataset contains 21 CCTV cameras, 20 WIM stations, and 6 TAMS stations. The ROI counts are 9 CCTV cameras, 14 WIM stations, and 2 TAMS stations. For all sensors in the ROI, data were available for Fri 11/3, 5 PM to Sun 11/5, 12 AM (19 hours), and Tue

11/7, 9 AM to Wed 11/8, 10 AM (13 hours). Therefore, we used this subset and time period in the experiments to evaluate the truck detection and counting and the truck flow estimation algorithms.

YOLO, a deep learning-based object detection algorithm, and StrongSORT, a robust tracking algorithm, were applied to Caltrans CCTV videos in the ROI to study the feasibility and usefulness of using traffic monitoring videos for truck detection and counting. For training the models, 9 one-minute clips are used from locations three locations containing recordings of roads with moving trucks, shot at 6 am, 10 am, and 2 pm. One frame per second (1fps) was used to sample images in labeling videos to reduce manual labor while not missing truck detection. We train the model to detect the four classes in our dataset. The model performance was assessed by comparing the results with the ground truth count. YOLOv5, a variant of YOLO, was used due to its high detection accuracy and very low inference times, allowing us to deploy it on video streams. Additionally, the StrongSORT provides robust object tracking even during extended periods of occlusion, a typical phenomenon on highways, especially during rush hours. We used 10 1-minute CCTV videos shot at location 255 between ~ 9:01-9:11 am to generate inference results. The 10 videos are stitched together to generate one 10-minute, 8-second-long video. Precision, Recall, and Mean average precision (mAP50) using an Intersection-over-Union (IoU) threshold of 0.5 was 0.762, 0.746, and 0.777, and the absolute difference between detected and actual normalized by an actual number of trucks was 0.21 in the overall sampled test cases.

Because CCTV and WIM sensors do not uniquely identify and track vehicles, extracting mobility patterns from their detections is challenging. We have proposed a framework named VPE, short for Visit Probability Estimation, that processes roadside sensor observations to estimate the probability with which a vehicle visits a road segment at a specific time. VPE is powered by LEM, short for Location Estimation Model, a novel mathematical model that calculates location transition probabilities while considering the sensors' reliability, and APD+, an algorithm that captures the uncertainty of movement between two endpoints. Our proposed algorithm APD+ to discover the most likely feasible paths the vehicle could have taken. We refer to the set of these paths as bridgelet. Subsequently, we weigh each constituent path of the bridgelet and aggregate them to produce a probability cloud, i.e., a mapping from road segments to the probability that the segment was visited during the trip. We detail the steps in our paper "*Estimating mobility distributions from uncertain roadside sensor datasets*" published at the 25th International Conference on Mobile Data Management (MDM'24). Our experiments on synthetic datasets show that the proposed methods achieve high accuracy while maintaining practical computation time. We applied our proposed framework, VPE, to several realistically synthesized datasets of roadside sensor observations. In our simulations, our methods were able to estimate visit probabilities accurately. Specifically, the mean Jensen-Shannon divergence between the estimated and actual (simulated) distribution was approximately 5%, while the mean F1 score was approximately 40%.

Introduction

In the following, we present a detailed report on the research. This includes (i) a curated Truck Sensing Dataset containing sensor data in the ROI that we have leveraged to produce experimental results and that other transportation researchers can use, (ii) state-of-the-art Truck Detection and Counting using a deep learning-based algorithm that we have specifically trained and tested on CCTVs in the ROI to investigate the feasibility and usefulness of using traffic monitoring CCTV videos for truck detection and counting, and (iii) a novel framework named VPE, short for Visit Probability Estimation, that processes roadside sensor observations to estimate the probability with which a vehicle visits a road segment at a specific time that we used to investigate the feasibility of truck flow estimation.

Truck Sensing Dataset

To assemble the Truck Sensing Dataset, we worked with Caltrans District 7 to collect CCTV cameras, with the California Department of Transportation Traffic Operations Weigh-in-Motion to collect WIM station data, and with the UCI Institute of Transportation Studies to collect Truck Activity Monitoring System (TAMS) data.

Other sensors in the area were also considered but not retained as they did not provide sufficiently accurate information. For example, single induction loop detectors (ILD) systems, while present in large numbers in the ROI, have been shown to provide accurate speed and occupancy measurements for passenger vehicles but not for trucks. In addition, we have examined the IMSC ADMS https://adms.usc.edu/app data consisting of IDL data as a source of traffic information that could be used by the flow estimation algorithms and found that in the ROI, the data was too sparse and unreliable due to many IDL sensors being defective. Figure 1 presents an overview of arterial and highway IDL sensor availability in the ROI from 2019 to 2023.

Finally, we have identified other sensors, such as the RFID sensors at the ports, that can be used to inform flow but did not collect as the focus was on flow estimation on the highway system. One advantage of the RFID system is that if available at multiple locations, it could be used as ground truth as it would uniquely identify trucks. Figure 1 and Table 1 provide an overview of the sensors available in the ROI spanning an area of approximately 12 square miles north of the Port of Long Beach and the Port of Los Angeles. The map of Figure 1 with all the sensor locations for the Truck Sensor Dataset is available online at:

https://www.google.com/maps/d/u/0/edit?mid=1fIdKIyqwpWPwW0RDsmf-pwg_3M0pfGo&ll=33.826202632011245%2C-118.37865048828125&z=11

The dataset contains 21 CCTV cameras, 20 WIM stations, and 6 TAMS stations regarding sensor counts. The ROI counts are 9 CCTV cameras, 14 WIM stations, and 2 TAMS stations.

For all sensors in the ROI, data were available for Fri 11/3, 5 PM to Sun 11/5, 12 AM (19 hours), and Tue 11/7, 9 AM to Wed 11/8, 10 AM (13 hours). Therefore, we used this subset in the experiments to evaluate the truck detection and counting and the truck flow estimation algorithms of sections CCTV Truck Detection and Counting and Truck Volume Modeling, respectively.

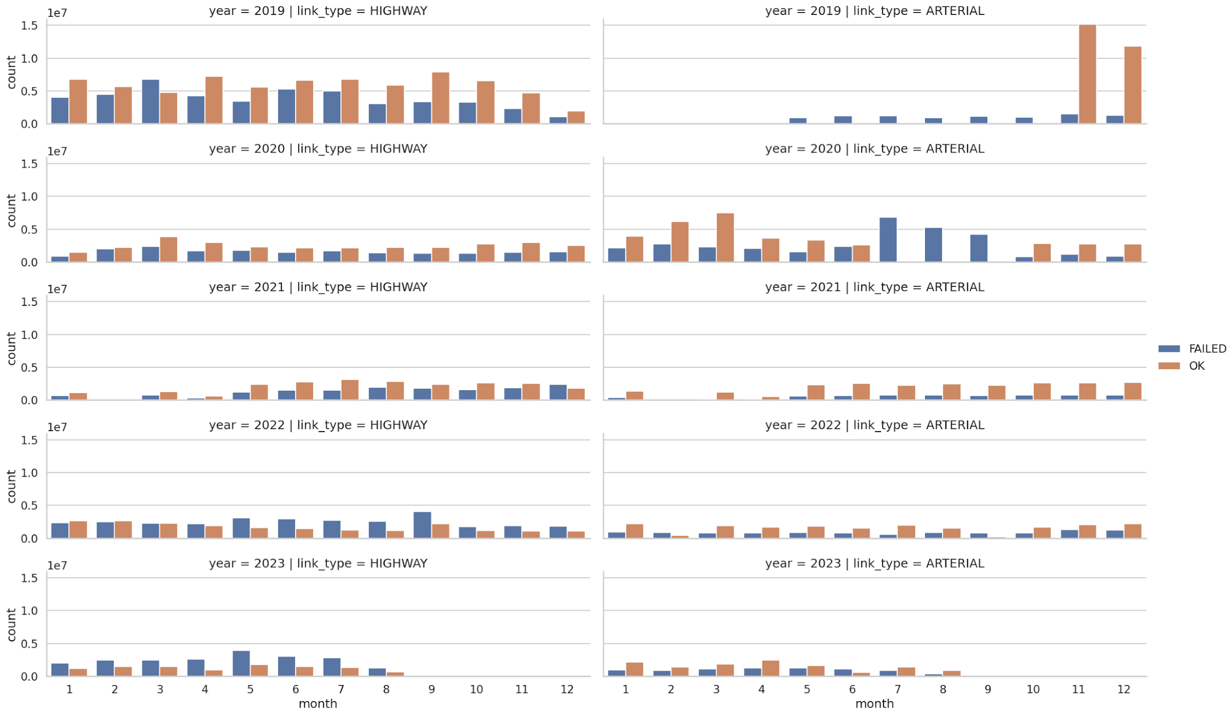

Figure 1. Arterial and highway IDL sensor availability in the ROI from 2019 to 2023.

Table 1. Periods with available sensors in ROI

Sensor	Periods Available	Count
CCTV	2023 11/02 5 PM to 11/05 12 AM 2023 11/07 9 AM to 11/08 10 AM	9
WIM	01/01/2019 – 08/31/2023 11/02/2023 – 11/08/2023	14
TAMS	11/01/2023 - 11/10/2023	2

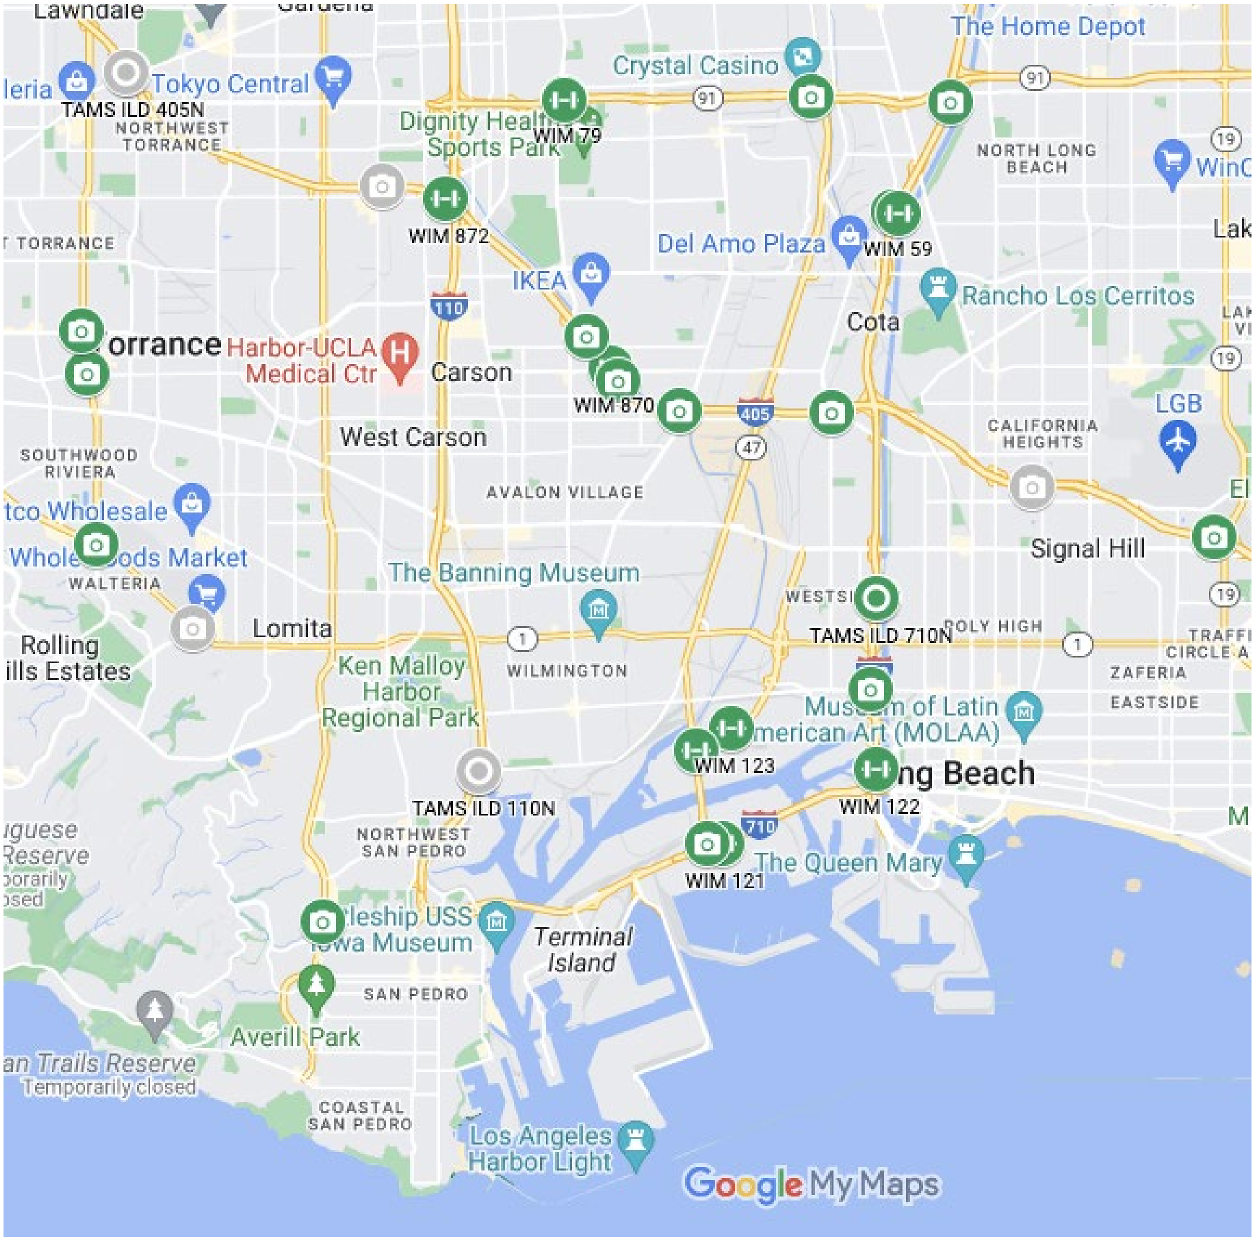

Figure 1. CCTV (cameras), WIM (weights), and TAMS (O) in ROI

Region of Interest

We have selected a region of interest (ROI) approximately 12 square miles in size, north of the Ports of Los Angeles and Long Beach. This area was specifically chosen due to the significant impact of truck traffic on freight movement to and from these major ports. The Ports of Los Angeles and Long Beach are among the busiest ports in the United States, serving as critical hubs for international trade. The constant flow of trucks transporting goods from the ports to various distribution centers and end destinations contributes to substantial traffic congestion in the surrounding areas. This heavy truck traffic not only affects the efficiency of transportation logistics but also has significant environmental and health implications. Emissions from diesel trucks contribute to air pollution, which can lead to respiratory problems and other health issues for residents. The increased traffic volume also leads to greater wear and tear on infrastructure, necessitating more frequent repairs and maintenance. By focusing on this ROI, we aim to better understand and address the challenges posed by freight transportation, develop strategies to mitigate negative impacts and improve the overall quality of life for communities in this region.

CCTV data

The CCTV video data consists of recordings from closed-circuit television cameras (CCTV) that are used by Caltrans to monitor the freeways. With this research, we aim to repurpose these sensors to count trucks as an input into the algorithms that estimate truck flow. Because Caltrans regularly updates its cameras, therefore, the resolution of the available cameras in the ROI can vary. However, when using the recordings to detect and count trucks, the footage resolution does not seem to be a limiting factor, so videos from older cameras can also be used. Table 2 provides an overview of the complete CCTV data collected with the help of Caltrans District 7 and shows which CCTV is available on highways. The table shows 12 CCTVs, 9 in the ROI and on a highway. Of the relevant 9 cameras in the ROI, there are two periods during which all 9 cameras have recordings available:

- Fri 11/3 5 PM to Sun 11/5 12 AM (19 hours)
- Tue 11/7 9 AM to Wed 11/8 10 AM (13 hours)

These are the specific CCTV recordings that we are leveraging in the experiments.

Table 2. Complete CCTV dataset availability for when WIM and TAMS data are available.

			11/2-11/5		**11/7-11/8**	
ID	**HWY**	**ROI**	**From**	**To**	**From**	**To**
221		X	11/3 1:00 PM	11/5 12:00 AM	11/7 8:00 AM	11/8 10:00 AM
222		X	11/3 12:00 PM	11/5 12:00 AM	11/7 9:00 AM	11/8 11:00 AM
223		X	11/3 11:00 AM	11/5 12:00 AM	11/7 9:00 AM	11/8 11:00 AM
224		X	11/3 11:00 AM	11/5 12:00 AM	11/7 9:00 AM	11/8 11:00 AM
255	X	X	11/3 1:00 PM	11/5 12:00 AM	11/7 9:00 AM	11/8 11:00 AM
270	X	X			11/7 9:00 AM	11/8 11:00 AM
276			11/3 9:00 AM	11/5 12:00 AM	11/7 9:00 AM	11/8 11:00 AM
277			11/3 8:00 AM	11/5 12:00 AM	11/7 9:00 AM	11/8 11:00 AM
282			11/3 7:00 AM	11/5 12:00 AM	11/7 9:00 AM	11/8 11:00 AM
283			11/2 11:00 PM	11/5 12:00 AM	11/7 9:00 AM	11/8 11:00 AM
284			11/2 11:00 PM	11/5 12:00 AM	11/7 9:00 AM	11/8 11:00 AM
285			11/2 10:00 PM	11/5 12:00 AM	11/7 9:00 AM	11/8 11:00 AM
286			11/2 10:00 PM	11/5 12:00 AM	11/7 9:00 AM	11/8 11:00 AM

323	X	X	11/2 10:00 PM	11/5 12:00 AM	11/7 9:00 AM	11/8 11:00 AM
324			11/2 7:00 PM	11/5 12:00 AM	11/7 9:00 AM	11/8 11:00 AM
327	X	X	11/2 7:00 PM	11/5 12:00 AM	11/7 9:00 AM	11/8 11:00 AM
335	X	X	11/3 5:00 PM	11/5 12:00 AM	11/7 6:00 AM	11/8 11:00 AM
336	X	X	11/2 3:00 PM	11/5 12:00 AM	11/7 5:00 AM	11/8 11:00 AM
337	X	X	11/2 6:00 PM	11/5 12:00 AM	11/7 6:00 AM	11/8 11:00 AM
954	X	X	11/2 3:00 PM	11/5 12:00 AM	11/7 8:00 AM	11/8 11:00 AM
956	X	X	11/2 1:00 PM	11/5 12:00 AM	11/7 7:00 AM	11/8 11:00 AM

WIM data

Thanks to Caltrans, we have collected WIM station data in Table 3, which includes data from recently commissioned WIM stations that are not yet mapped on the Caltrans website WIM map at https://dot.ca.gov/programs/traffic-operations/wim/locations. Data temporal coverage is from Jan 1, 2019, to Aug 31, 2023, and Nov 02, 2023, to Nov 08, 2023. As shown in Table 3, spatial coverage is on and around the Port of Los Angeles and Port of Long Beach in Caltrans District 7. The curated dataset includes the geographical coordinates of all WIM stations, which were geocoded by the Postmile Services.

Tables 3 and 4 summarize the WIM data. Table 3 provides sensor availability, and Table 4 shows a breakdown by year. There are 7 WIM stations in the ROI that are in service and that we can leverage for the experiments. WIM data files consist of individual truck records, including time lane and truck class information according to Caltrans WIM ASCII TRUCK RECORD FILE FORMAT, which we document in the dataset metadata.

Table 3. WIM station information

ID	Route	ROI	Availability	Issues
008	101		Jan 1, 2019 - Nov 22, 2022	Server issues
009	101		Jan 1, 2019 - Nov 22, 2022	Server issues
012	405		Jul 27, 2020 - Aug 31, 2023	
013	405		Jan 1, 2019 - Aug 31, 2023	
047	5			Under construction since Oct 13, 2018
048	5			Under construction since Oct 13, 2018
059	710	X	Jan 1, 2019 - May 23, 2022	Phone line issues
060	710	X	Jan 1, 2019 - Oct 21, 2021	Phone line issues
066	126		Jan 1, 2019 - Nov 23, 2022	Server issues
079	91	X	Jan 1, 2019 - Nov 23, 2022	Server issues

080	91	X	Jan 1, 2019 - Nov 23, 2022	Server issues
082	210			Under construction since Aug 3, 2018
083	210			Under construction since Aug 3, 2018
116	47			Abandoned since Mar 2012
121	710	X	Sept 1, 2021 - Aug 31, 2023	
122	710	X	Sept 1, 2021 - Aug 31, 2023	
123	103	X	Dec 1, 2021 - Mar 14, 2023	Phone line issues
124	47	X	Dec 1, 2021 - Oct 18, 2022	Phone line issues
870	405	X	Jan 1, 2019 - Jan 3, 2021	Phone line issues
872	405	X	Jan 5, 2022 - Aug 31, 2023	

Table 4. Stations available by year. Rows in yellow denote

Station	**ROI**	**2019**	**2020**	**2021**	**2022**	**2023/01-08**	**2023/11/2-8**
008		X	X	X	X		X
009		X	X	X	X		X
012		X	X	X	X	X	X
013		X	X	X	X	X	X
059	X	X	X	X			
060	X	X	X	X			
066		X	X	X	X		X
079	X	X	X	X	X		X
080	X	X	X	X	X		X
121	X			X	X	X	X
122	X			X	X	X	X
123	X			X	X	X	X
124	X			X	X		X
870		X	X	X			
872	X				X	X	X

TAMS data

This dataset contains the Truck Activity Monitoring System (TAMS), originally conceived and developed by the University of California, Irvine. TAMS sensors consist of existing Inductive Loop Detectors (ILD) upgraded with inductive loop signature technology and implementing state-of-the-art machine-learning classification models. Each row corresponds to a single WIM record providing the following information:

- lane_dir: directions N/S as 1 or 2
- lane: lane number, 1-3
- epoch: Unix timestamp
- tier2_class: vehicle class according to the five-vehicle category scheme of Table 5.

For Wed Nov 01, 2023, 00:00:02 to Fri Nov 10, 2023, 23:59:58 GMT-0800 (Pacific Standard Time), only two TAMS are available: I-710 N/S.

Table 5. Five-vehicle category scheme:

tier2_class_id	tier2_class	description
1	PC	Passenger Vehicle
2	SU	Single Unit Truck
3	Single	Truck with Single Trailer
4	Semi	Tractor with Semi-Trailer
5	Multi	Tractor with Semi-Trailer

CCTV Truck Detection and Counting

The goal of this study is to investigate the feasibility and usefulness of using traffic monitoring CCTV videos for truck detection and counting. Using a deep learning-based algorithm, i.e., YOLO, as the detection model, we explore the possible benefits of visual images in monitoring truck movement. To assess the effectiveness of our approach, truck traffic data were collected and labeled, encompassing both daytime and nighttime, employing our custom dataset sourced from the Caltrans. Using the collected dataset, the performances of deep learning models trained for nighttime and daytime conditions are assessed and compared. Additionally, leveraging the trained object detection models and the StrongSORT object counting algorithm [6], the number of passing trucks was counted from the videos. Furthermore, through a comparison of the results with the ground truth count, we investigate the effectiveness of our approach. This study will contribute to the practical and reliable image learning-based solutions that support 24/7 truck movement understanding, and the studies in urban traffic management, environmental sustainability and public safety.

Modeling

Some number of sample images from our CalTrans video datasets were manually selected and labeled to categorize typical instances of trucks that showed up in the videos in order to perform supervised machine learning, i.e., truck detection. One frame per second (1fps) was used for sampling images in labeling videos. This sampling frequency was chosen to reduce the manual labor in labeling while not missing truck detection. Considering that a passing truck appears for at least several seconds (i.e., in several images) in a video so our model can have enough chances to detect and count it.

9 one-minute clips are used from locations 270, 327, and 956 to construct a training dataset. These contain recordings of roads with moving trucks, shot at 6 am, 10 am, and 2 pm. Frames are sampled from these videos at the rate of 1 frame per second. Some statistics pertaining to the training set are given in Table 6.

Table 6. CCTV Datasets used for modeling.

Split	Images	Labeled Images	Backgrounds	Instances
Training	368	267	101	788
Testing	185	127	58	232
Total	553	394	159	1020

Methods

We have devised two strategies for truck detection leveraging the cutting-edge object detection algorithm YOLOv5 [4]. Our implementation of the truck detection and classification component is based on the YOLOv5 network. YOLOv5 is a two-stage object detection and classification framework that improves upon previous-generation variants. The main benefit of YOLOv5 is that, besides its high detection accuracy, it incurs very low inference times, allowing us to deploy it on video streams. We train the model to detect the four classes in our dataset: [class description]. Additionally, to handle the small imbalance in our dataset, we perform upsampling and augmentation during training. After classification, the detected objects are forwarded to the tracking component. The tracker processes the detected objects in a sequence of frames and makes associations of bounding boxes that contain the same object across multiple frames. In this study, we use the StrongSORT [6] algorithm. This algorithm builds on top of the traditional SORT algorithm with a pre-trained association metric. The main benefit of this algorithm is its ability to track objects even during extended periods of occlusion, a phenomenon that is typical in highways, especially during rush hours.

Figure 2 demonstrates the workflow of our truck detection and counting method. In details:

- An instance of the YOLOv5-small model is trained on frames of training videos.
- The fine-tuned YOLOv5-small model is then loaded into the StrongSORT model.
- A given testing video is expanded into individual frames, where each frame is then fed to the detection module, which is the fine-tuned YOLOv5-small model.
- Results from the detection stage are passed on to the tracking stage.
- The tracking module then generates results for tracked objects.
- A CSV file is generated for all objects detected and tracked by the model, containing object ID and timestamps in the video when the object was first and last seen by the StrongSORT model.

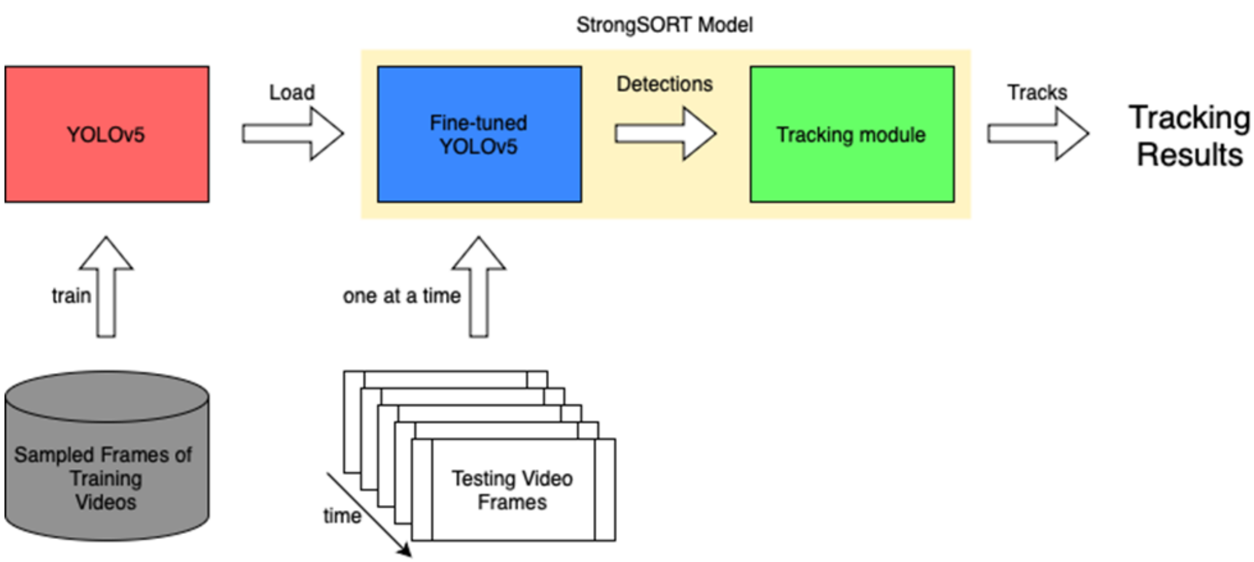

Figure 2. Workflow of Truck Detection and Tracking

Experiments and Results

In the first set of experiments, we evaluate the performance of overall truck detection and classification in terms of Precision (P), Recall (R), and Mean Average Precision (mAP@.5) using an Intersection-over-Union (IoU) threshold of 0.5. Table 7 summarizes the overall detection performance of the models.

In the experiments, we used the following setups:

1. Training the YOLOv5s model

- Epochs: 50
- Image Size: 640 x 640
- Batch Size: 32
- Optimizer: AdamW (lr = 0.01)

2. Inference using StrongSORT

```
STRONGSORT:
  ECC: True                # activate camera motion compensation
  MC_LAMBDA: 0.995         # matching with both appearance (1 - MC_LAMBDA) and motion cost
  EMA_ALPHA: 0.9           # updates  appearance  state in  an exponential moving average manner
  MAX_DIST: 0.2            # The matching threshold. Samples with larger distance are considered an invalid match
  MAX_IOU_DISTANCE: 0.7    # Gating threshold. Associations with cost larger than this value are disregarded.
  MAX_AGE: 90              # Maximum number of missed misses before a track is deleted
  N_INIT: 3                # Number of frames that a track remains in initialization phase
  NN_BUDGET: 100           # Maximum size of the appearance descriptors gallery
```

Figure 3: Hyperparameters for StrongSORT

Apart from the above, we use a confidence threshold of 0.6 and IoU threshold of 0.4 for detection.

To generate inference results, we used 10 1-minute CCTV videos shot at location 255 between ~ 9:01-9:11 am (07-17). The 10 videos are stitched together to generate one, 10-minute 8-second long video. The 10 videos used actually are a part of a larger dataset containing 2 hours worth of video content shot at 9 CCTV locations. From this test dataset, we generated the results in Table 7.

Table 7. Detection Accuracy

Conf	Precision	Recall	mAP50	mAP50-95
0.4	0.73	0.78	0.786	0.541
0.5	0.762	0.746	0.777	0.54
0.6	0.795	0.685	0.758	0.534

Table 8. Counting Results for Test Videos (IOU thresh = 0.4, Conf = 0.6)

Video	#Trucks Detected	Actual #Trucks	CER*
270_2pm_0000.122_0061.110.mp4	34	31	0.0967
327_2pm_0048.836_0112.353.mp4	9	5	0.8
956_2pm_0000.132_0060.233.mp4	9	7	0.2857
Overall	52	43	0.21

* CER = (absolute diff b/w detected and actual) / #actual trucks.

For truck counting using StrongSort algorithm, we used videos from three different CCTV locations (270, 327, 956). Table 8 shows the counting accuracy in CER. Depending on camera angles and other lighting conditions of CCTV mounting, detection and counting accuracy may vary significantly. For example, a video from the location 270 provided a very high accuracy while that from 327 showed a very low accuracy. Note that these are initial performance evaluation and we will update with more experimental results.

After developing the general truck detection and counting model, we applied it for multiple CCTV videos and generated results. Eight CCTV locations in the ROI were analyzed during 9-10 AM on Nov. 7, 2023 (see Table 9). For each CCTV location, the model detected and counted the number of passing trucks. For each detected truck, the result was recorded in the following tuple format: <TruckID, Time of First Appearance, Time of Last Appearance, Confidence, Direction> and the results were used in the truck volume modeling in the next Section.

Table 9. Counting Results at Eight Different CCTV Locations during 9-10 AM on Nov. 7, 2023

Location	South	North	West	East	Total
954	753	842			1595
956	1460	832			2292
336	524	178			702
335	1788	2527			4315
327	470	442			912
323	246	694			940
270	1401	1578			2979
255			1122	1079	2201

Discussions

Our experimental results demonstrated the feasibility that CalTrans CCTV analysis for truck detection and counting can be used in reality. The overall detection and counting accuracy were around 0.8 with basic models. The accuracy is expected to be higher when we handle more detailed real world considerations in image machine learning. Throughout the project, we identified the following issues in analyzing CalTrans CCTV video streams (Figure 4):

1) Road coverage issue: Most cameras are covering both sides of highways. But some cameras are not fully covering all lanes of the road so missing passing trucks (undercounting issue). On the contrary, a camera may cover not only highways but also nearby local roads. Then, an overcounting issue arises.
2) Environmental issue: Some cameras are heading East or West so facing a very strong Sunlight at sunrise or sunset time, which limits the identification of any objects in images. Another challenging case is foggy weather which blurs images.
3) Nighttime low light issue: Nighttime videos are fundamentally constrained by low light images so detection accuracy suffers.

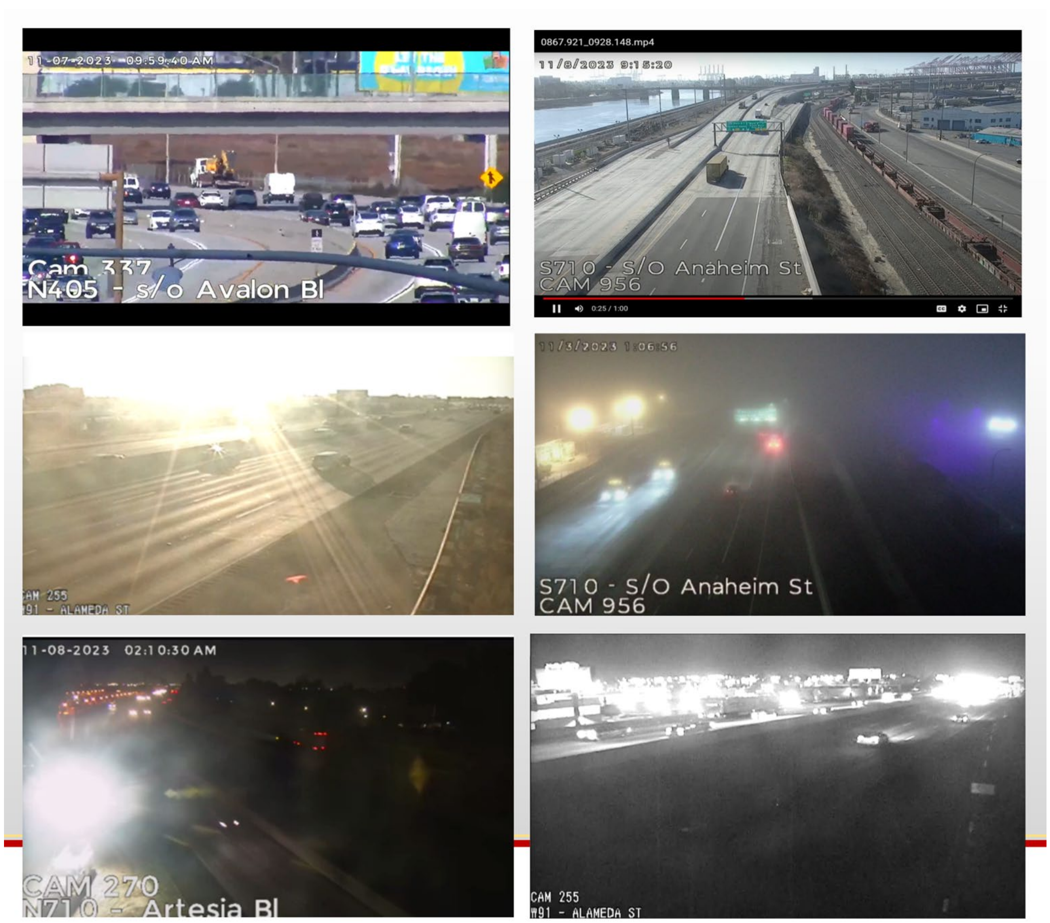

Figure 4. Real World Issues in CalTrans CCTV Video Analysis

Truck Volume Modeling

Because CCTV and WIM sensors do not uniquely identify and track vehicles, extracting mobility patterns from their detections is challenging. We have proposed a framework named VPE, short for Visit Probability Estimation, that processes roadside sensor observations to estimate the probability with which a vehicle visits a road segment at a specific time. VPE is powered by LEM, short for Location Estimation Model, a novel mathematical model that calculates location transition probabilities while considering the sensors' reliability, and APD+, an algorithm that captures the uncertainty of movement between two endpoints. Our experiments on synthetic datasets show that the proposed methods achieve high accuracy while maintaining practical computation time.

Modeling

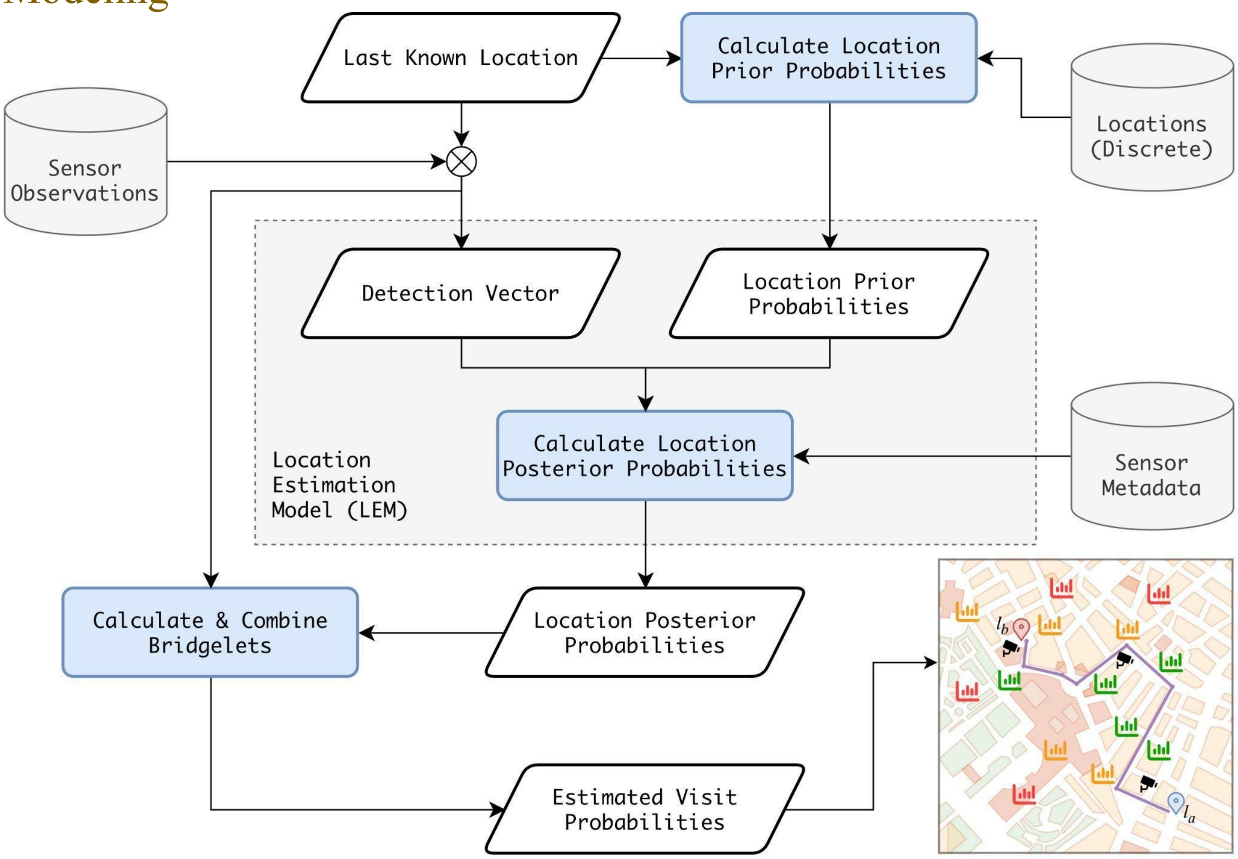

Figure 4. The architecture and data flow of the VPE framework.

Our proposed framework processes sensor observations to estimate the probability that a given vehicle visited, i.e., traversed, a road segment. Figure 4 shows the components of VPE and how they interconnect. VPE takes as input a set of sensors S, their observations D, and a set of locations L at which the sensors detect vehicles. The output of VPE is a visit probability distribution over the road segments at each time step and for each vehicle. At first, sensor observations are processed to generate the detection vector for a given vehicle for the current time step, where each sensor accounts for one element of the detection vector. Then, the location of the vehicle at the previous time step (referred to as the last known location) is used to make an initial estimation of the probability that the vehicle is at any of the locations in the next time step (prior probabilities). Subsequently, the sensor's characteristics and trustworthiness are used to refine these probabilities and account for errors in detection, i.e., false positives and negatives. A visit probability distribution is calculated between the last known location and each next location that has non-zero probability. VPE produces its final output by combining these individual visit probability distributions.

Methods

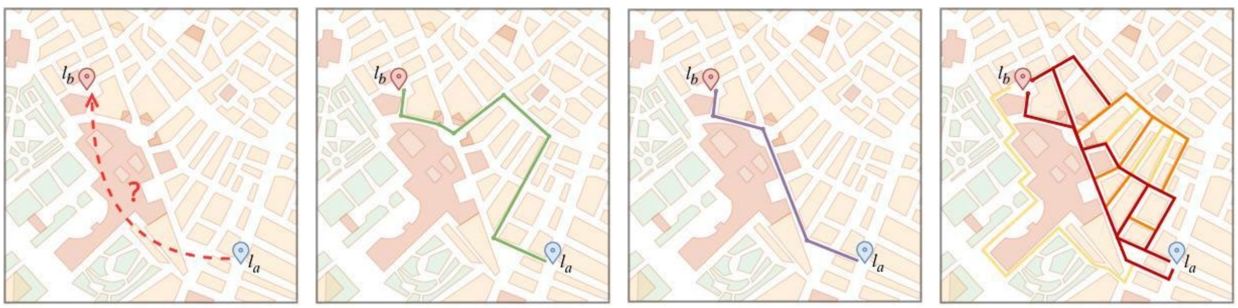

(a) The two endpoints (b) Actual trajectory (c) Shortest path (d) Probability distribution

Figure 5. Visit probability estimation between two endpoints. In (a), the two endpoints l_a and l_b are shown; in (b), the actual path (trajectory) is presented as a green line; in (c), the path is recovered using the fastest path method; and in (d) the probability distribution using road network-based bridgelets is plotted.

Consider the example shown in Figure 5. On the far left (Figure 5a), the two endpoint locations, l_a and l_b, are shown on a map. Next to it (Figure 5b), the actual path that the vehicle traveled on to get from the first location to the second is shown. The shortest path between the two locations is drawn in the middle-right map (Figure 5c). Evidently, the shortest path method misses several segments of the path, leading to inaccurate insights. However, on the far right (Figure 5d), a probability distribution over the road network segments is computed, capturing the mobility uncertainty more accurately. We observe that even though the edges that fall on the shortest path still carry a lot of weight (dark red), other possible edges retain some probability depending on how likely they are to have been used (with orange and yellow indicating a higher or lower probability, respectively.)

To estimate a probability distribution such as the one shown in Figure 5d, we first employ our proposed algorithm APD+ to discover the most likely feasible paths the vehicle could have taken. We refer to the set of these paths as bridgelet. Subsequently, we weigh each constituent path of the bridgelet and aggregate them to produce a probability cloud, i.e., a mapping from road segments to the probability that the segment was visited during the trip. We detail the steps in our paper "*Estimating mobility distributions from uncertain roadside sensor datasets*" published at the 25th International Conference on Mobile Data Management (MDM'24).

Experiments and Results

We conducted experiments to evaluate the performance of LEM in terms of Precision, Recall, and F1 scores. The following datasets are used in our experiments:

- **Road Network:** We obtain the road network for the metropolitan region of Los Angeles from OpenStreetMap. The road network graph contains 3,239,158 nodes and 4,190,761 edges.
- **Traffic Data:** We use the ADMS [1] system as the traffic data source and use the methods described in [3] to estimate the traffic at every road segment.
- **Trajectories:** We use a synthetic trajectory generator [2] to simulate 1000 realistic truck trajectories in a controlled environment where the traffic data is known.
- **Sensors:** We randomly distribute 300 sensors on the road network and replay the synthetic trajectories to generate sensor observations. We simulate various settings, from very reliable to very unreliable sensors.

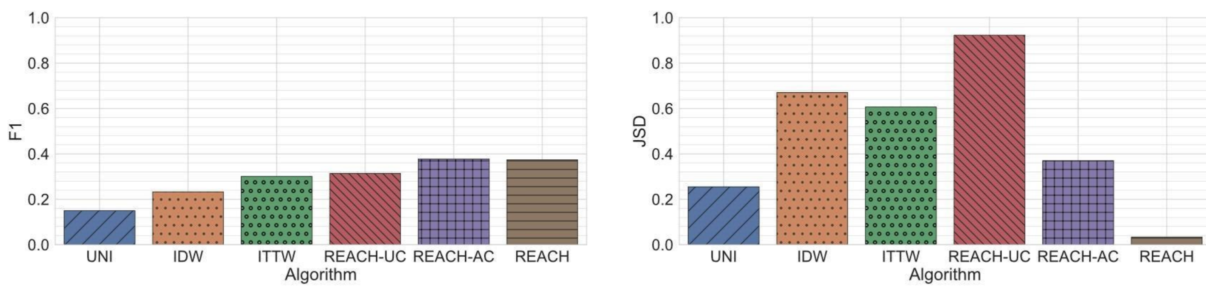

Figure 6. Comparison of LEM against baselines using F1 score and Jensen-Shannon divergence.

We varied the method used to generate the location prior probabilities and evaluate how the accuracy of LEM is affected. Specifically, we compare the following methods:

- **UNI:** Assigns the same probability to all locations in the detection vector (uniform).
- **IDW:** Uses inverse distance weighting to assign probabilities so that further away locations receive lower probability.
- **ITTW:** Similar to IDW but uses travel time instead of distance.
- **REACH-UC:** A variant of our approach that does not clip the probabilities.
- **REACH-AC:** A variant of our approach that only clips the final aggregated location probabilities using a threshold $\theta = 10^{-4}$.
- **REACH:** Our proposed reachability-based method with a clipping threshold $\theta = 10^{-4}$. The difference between this method and REACH-AC is that probabilities at intermediate time steps are also clipped.

Figure 6 shows the performance of all the algorithms. We observe that UNI exhibits the lowest F1 score as expected because it gives all detections equal probability, even if it is not feasible for the vehicle to travel from its last known location to that of the detection. Similar observations can be made for IDW and ITTW. However, these methods are more accurate than UNI because they implicitly assume that the farther away or the longer the vehicle travels to reach the detection, the less likely it is to be a true positive. They both fail to take into account the temporal dependency between the current and detection time steps, i.e., even if a detection is far, it may still be feasible to reach there in time.

On the other hand, the three reachability-based methods exhibit better performance, with REACH performing the best in terms of F1 score. Intuitively, this is attributed to the fact that detections that are not reachable are ignored. Finally, the distributions calculated by REACH exhibit the lowest JSD score. This means that the estimated distribution is very similar to the real distribution.

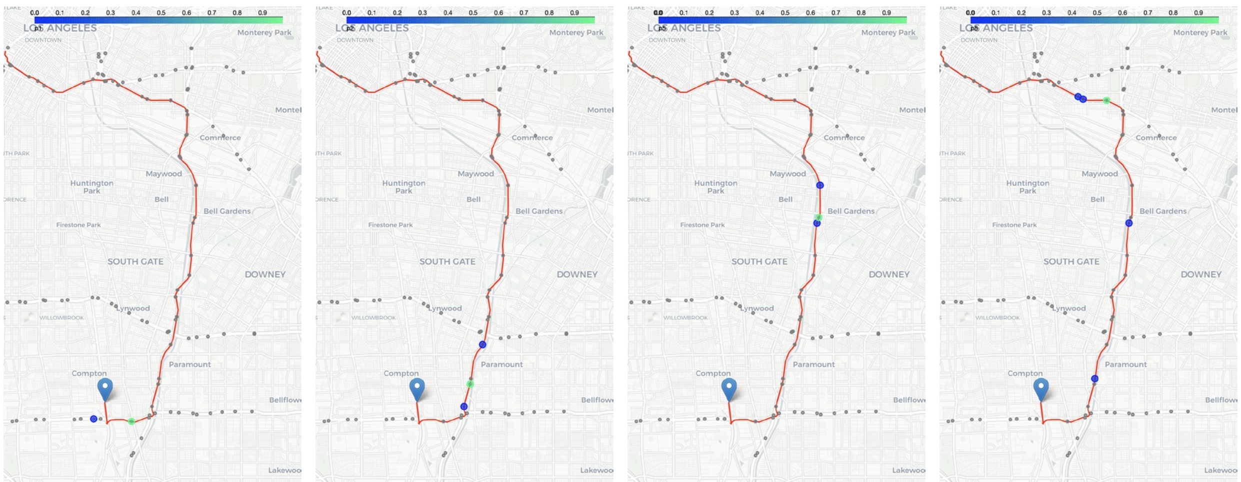

Figure 7. A visual example of a trajectory (red line), the set of sensors (grey dots), and the estimated location distributions (color-coded circles) at four different time steps.

In Figure 7, we provide a visual example of a trajectory and the vehicle's estimated locations at four different time steps. In the example, we plot the sensors using grey dots and the real trajectory of the vehicle as a red line for reference. Note that VPE is unaware of the trajectory but only of the sensors' observations. The only information provided to the algorithm for this example is the vehicle's initial location marked with a pin. We observe that at the first time step, VPE assigns a non-zero probability (color-coded circles) to a location that the vehicle did not visit. This is attributed to the fact that a sensor that observes that location is triggered at a time that is feasible to reach from the vehicle's current location. However, at subsequent time steps, the algorithm makes more accurate estimations. The reason behind this is that only a few detections are reachable by the vehicle.

Simulation with real-world sensors

Obtaining any real-world ground truth is nearly impossible. However, we still want to evaluate how well our proposed method does in a real-world setting. To that end, we simulate 5000 trajectories using real-world traffic data in the region of interest [1, 2] and we "replay" the simulated trajectories to generate the observations that 21 real sensors would have generated (10 WIM, 2 TAMS, 9 CCTV). In this setting, simulated trajectories are used to generate the ground truth truck flow for road segments in the region of interest. We evaluate the results on two families of metrics:

1. Set-based metrics: Precision, Recall, and F-1 scores

These metrics evaluate whether we estimate the precision on the "right" edges, i.e., the edges that incurred/observed truck volume in the simulation.

We compare three methods:

- IDW: Weights feasible pairs of consecutive observations using an inverse distance weighting function.
- ITTW: Similar to IDW but with an inverse travel time weighting function.
- APD: Our proposed reachability-based method that uses ITTW but with a pruning strategy to exclude pairs that are do not satisfy the reachability requirements (more details in our paper [])

The table below summarizes the results. We observe that APD achieves the highest F-1 score. Overall, APD significantly boosts recall only for a relatively small drop in precision. The reason behind this is because APD considers more than one paths between pairs of consecutive sensor observations. While this impacts precision (because it estimates flow on edges that belong to a feasible path but may have never been used by trucks), it provides a significant boost to recall for the same reason.

Method	**Precision**	**Recall**	**F1**
APD	0.753	0.864	**0.805**
ITTW	0.777	0.555	0.678
IDW	0.841	0.421	0.561

2. Estimation error metrics: MAE, RMSE

   These metrics evaluate how accurate the flow estimation is. We estimate the flow using the same three methods and compare their results.

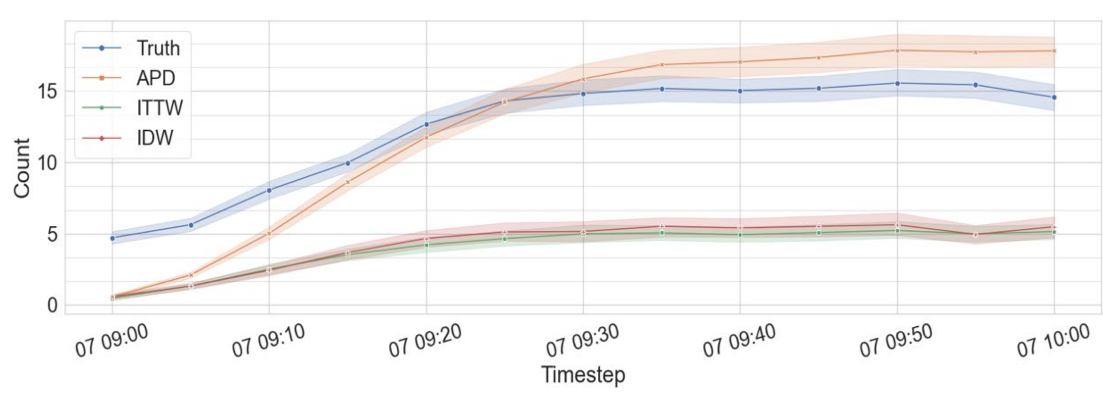

Figure 8. Mean expected truck flow of a road segment and the estimations generated by three methods: APD (proposed method), ITTW, and IDW.

We summarize the results in Figure 8. We observe that IDW and ITTW have similar performance with a MAE greater than 10. On the other hand, APD incurs a MAE of 6. In plain words, APD's estimations are on average off by 6 trucks whereas the other baselines are off by 10. The difference may seem insignificant but when multiplied by the number of road segments in a complex road network, it can add up to significantly large errors.

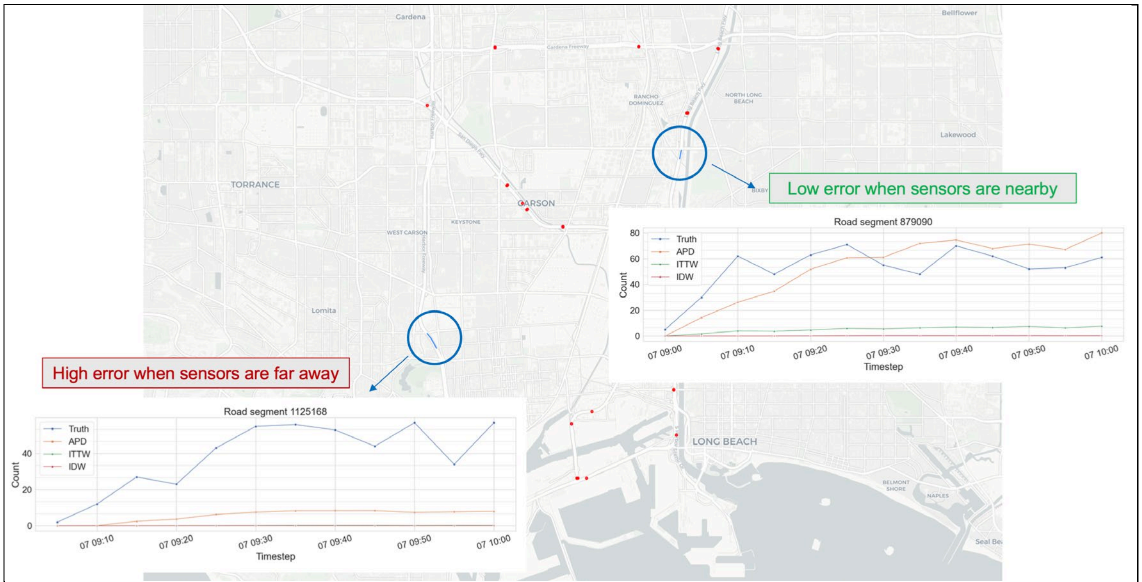

Figure 9. Flow estimations of two road segments shown on a map.

Another important observation, as shown in Figure 9, is that estimating the flow at road segments where no sensors exist nearby is quite challenging and typically incurs higher errors. This is expected because with no sensors nearby to provide us with insights, the uncertainty in those road segments becomes significantly higher.

References

[1] Anastasiou, Chrysovalantis, Jianfa Lin, Chaoyang He, Yao-Yi Chiang, and Cyrus Shahabi. "Admsv2: A modern architecture for transportation data management and analysis." In Proceedings of the 2nd ACM SIGSPATIAL International Workshop on Advances on Resilient and Intelligent Cities, pp. 25-28. 2019.

[2] Anastasiou, Chrysovalantis, Seon Ho Kim, and Cyrus Shahabi. "Generation of synthetic urban vehicle trajectories." In 2022 IEEE International Conference on Big Data (Big Data), pp. 359-366. IEEE, 2022.

[3] Anastasiou, Chrysovalantis, Juanhao Zhao, Seon Ho Kim, and Cyrus Shahabi. "Data-driven traffic index from sparse and incomplete data." In 2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC), pp. 2593-2598. IEEE, 2022.

[4] Zhu, Xingkui, Shuchang Lyu, Xu Wang, and Qi Zhao. "TPH-YOLOv5: Improved YOLOv5 based on transformer prediction head for object detection on drone-captured scenarios." In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 2778-2788. 2021.

[5] Jooyoung Yoo, Seon Ho Kim, Krish Sukhani, Min Sang Yoo, Cyrus Shahabi. Truck Detection and Counting in Low-Light Condition: Do We Need Infrared Camera? Accepted to appear in the IEEE International Conference on Big Data and Smart Computing (BigComp 2024), Feb. 2024.

[6] Du, Yunhao, Zhicheng Zhao, Yang Song, Yanyun Zhao, Fei Su, Tao Gong, and Hongying Meng. "Strongsort: Make deepsort great again." *IEEE Transactions on Multimedia* (2023).

Data Management Plan

This Data Management Plan was prepared to include the artifacts of this research and comply with recently enacted U.S. Department of Transportation Public Access Policies, which require that all data used or created under the University Transportation Centers program be accessible to the public.

Products of Research

The data that were collected and used for the study are listed as part of the Truck Sensing Dataset section and include:

- CCTV data: recordings for 22 cameras from Fri 11/03/2024 5 PM to Sun 11/5 12 AM (19 hours) and Tue 11/07/2024 9 AM to Wed 11/8 10 AM (13 hours) provided by Caltrans District 7.
- WIM data: contains the Weigh-in-Motion (WIM) data from 15 stations in California District 7 from Jan 1, 2019, to Aug 31, 2023, and from Nov 02, 2023, to Nov 08, 2023, as provided by Caltrans Traffic Operations.
- TAMS data: five-vehicle category scheme counts (passenger vehicle, single-unit truck, truck with single trailer, tractor with semi-trailer, and tractor with multiple trailers) for the TAMS ILD 710N/S from Wed Nov 01, 2023, 00:00:02 to Fri Nov 10 2023 23:59:58 PST as provided by the Institute of Transportation Studies of the University of California, Irvine.

In addition, we have produced as part of the research:

- Truck detection and counting models on working on image data and producing probabilistic truck observations
- Truck flow estimation algorithms with corresponding experimental results on synthetic and Truck Sensing Datasets data.

Data Format and Content

Describe the format, or file type, of the data and the contents of each file.

- CCTV data consists of approximately 2TB of recordings in mp4 format stored in folders according to the 3-digit camera identifier.
- WIM data consists of day-level ASCII files named with year, month, day, and station ID. Each file consists of individual WIM records in ASCII TRUCK RECORD FILE FORMAT: This file shall include every "truck record" in the daily data file. Each field shall be comma delimited and padded with blanks to complete the fixed logical record length. For axle weight-only weighing (in lieu of right and left wheel weighing), either the "AXLE n RT. WEIGHT" or the "AXLE n LT. WEIGHT" field may be used for the "AXLE n WEIGHT".
- TAMS data for TAMS I-710 N/S are stored in a single file named i710_data.csv, with headers lane_dir, lane, epoch, and tier2_class. lane_dir: encodes the N/S directions N/S as 1 or 2, lane: is the lane number, 1-3, epoch: is the Unix timestamp of the record, and tier2_class: is the vehicle class according to the five-vehicle category scheme (1:

Passenger Vehicle, 2: Single Unit Truck, 3: Truck with Single Trailer, 4: Tractor with Semi-Trailer and 5: Tractor with Semi-Trailer (Multi).

Data Access and Sharing

The general public can access the data by contacting the researchers and receiving access to a Google Drive that contains all the data used in the research with README files describing each dataset, its metadata, including collaborators' information, spatial and temporal coverage, provenance and license, and a data dictionary. Because of their size, only a sample of the CCTV cameras included in the ROI is in Google Drive at the following public link:

https://drive.google.com/drive/folders/1BCIQ4YoHNahkMJwT9OqTjEsFvofxQ6zf?usp=sharing

Reuse and Redistribution
Restrictions on how the data can be reused and redistributed by the general public. TBD according to Caltrans requirements.