



**A USDOT NATIONAL
UNIVERSITY TRANSPORTATION CENTER**

Carnegie Mellon University



THE OHIO STATE UNIVERSITY



Bus on the Edge: Applications

PI: Christoph Mertz

<https://orcid.org/0000-0001-7540-5211>

FINAL RESEARCH REPORT

Contract # 69A3551747111

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 U.S. Department of Transportation's University Transportation Centers Program, in the interest of information exchange. The U.S. Government assumes no liability for the contents or use thereof.

Overview

Before the start of this project we developed a bus-on-the-edge system that uses cameras mounted on a bus to monitor infrastructure and traffic. The hardware components are a standard security system for transit buses: five cameras observing the surrounding of the bus and a computer to record the videos (Figure 1).



Figure 1 The left four images show the commuter bus, an image of the camera, the cabinet that contains the computer and electronics, and one of the cameras. The right diagram shows the field of view of each installed camera, where one faces forward and four side cameras face opposite directions from each other.

We have installed our own software on that computer. The main functions of the software is to analyze and manage the data. The videos from the five cameras are far too much data to upload to the cloud. Instead, it needs to be pre-analyzed on the bus itself and then data of interest is sent to a central location where it is analyzed more thoroughly and used for various applications (Figure 2).

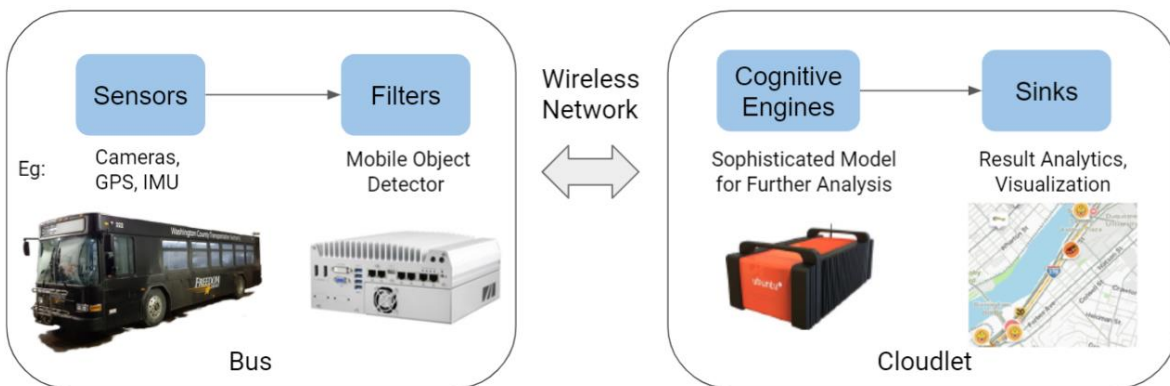


Figure 2 Schematic of data and analysis flow. The raw sensor data are pre-analyzed by a computer on the bus. Interesting data is wirelessly transmitted to a server ("cloudlet"), where it is analyzed in detail and finally sent to the end-user.

The computer has a cellular connection for time critical communication and a WiFi link to exchange large amounts of data. The system was installed on a FreedomTransit bus in February 2021 and we tested a first set of applications at that time. In the project of the past year we added several more applications to the system, two of them specifically

benefit the transit agency, one jointly done with a NSF project and the last one was done in collaboration with Argo. We will give a brief overview of these applications in this section; detailed descriptions are in the attached papers and Master's theses.

The first two applications concern issues at or around bus stops. Currently the transit agency finds out about issues through their drivers, complaints from users, or they have to send out staff. None of these is ideal. The bus driver has to operate the bus and deal with passenger and has little time to record issues. Users often complain only about major issues and are inconsistent in their reporting. Sending out staff is reliable, but it is expensive and time consuming. We want to use the bus-on-the-edge system to automate the reporting.

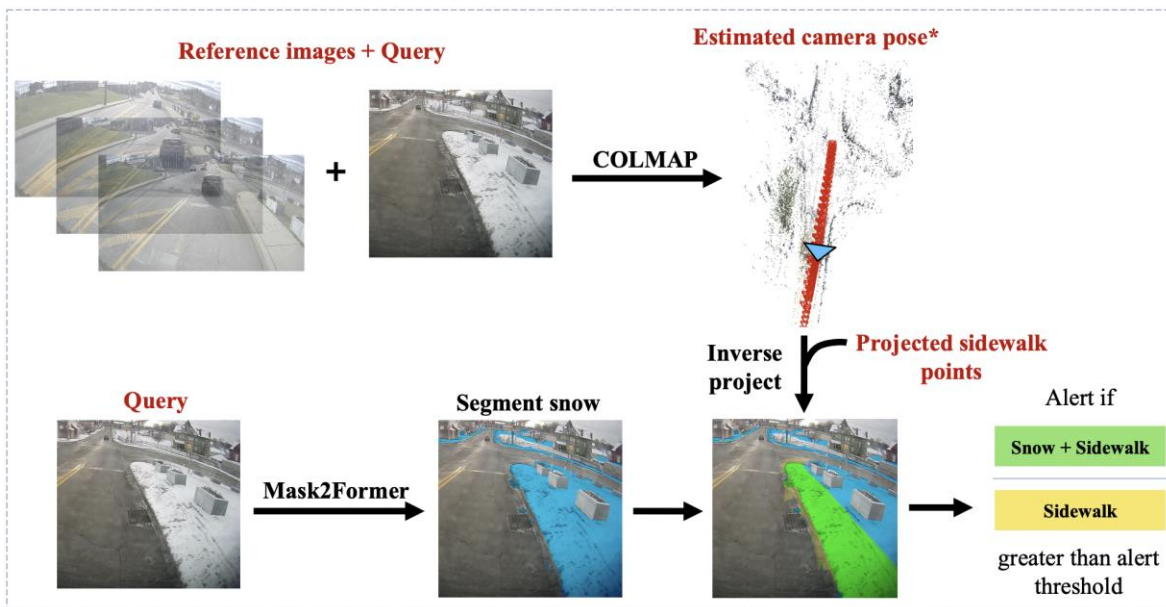
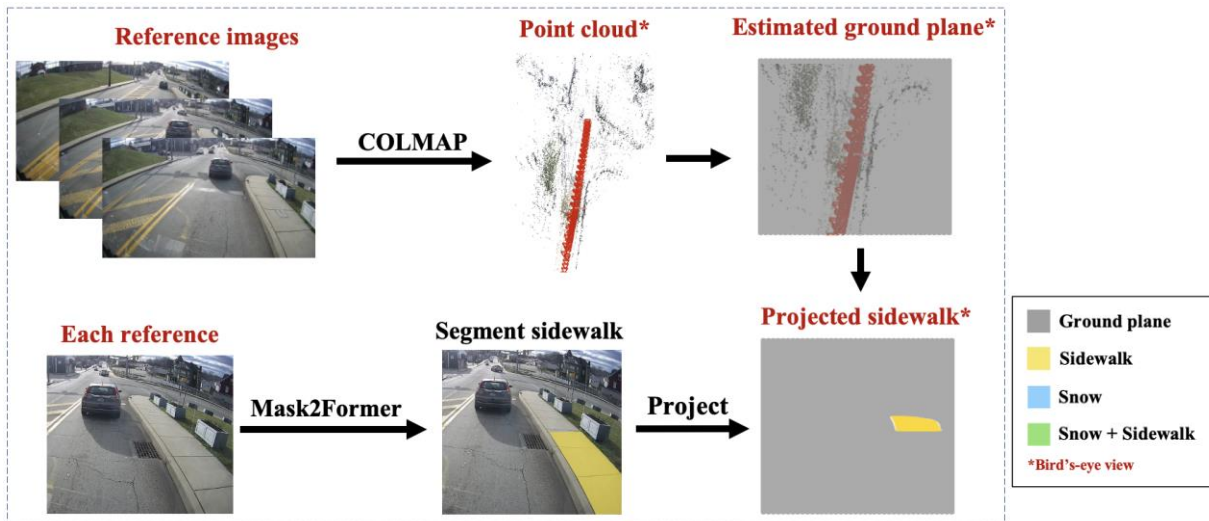


Figure 3 Detection of snow covered side walks: The method begins by using clear-weather images of the stretch to render a 3D model of the sidewalks as shown in (top). This involves running SfM on the reference images to render a point cloud of the scene, estimating the ground plane of the point cloud, and then projecting the portion of the sidewalk mask that is closest to the bus in each reference image onto the estimated ground plane. When the bus encounters snow coverage, the method compares the snow to the expected sidewalk area as shown in (bottom). This process involves estimating the camera pose of the query by re-running COLMAP with the query added to the reference images and using the estimated pose to project the saved sidewalk points into the query image (i.e. the inverse projection of (top)). Finally, if the proportion of the expected sidewalk area that is covered by snow exceeds a set alert threshold, it generates an alert.

The first application we developed is to detect snow on the sidewalk at or near bus stops.

It is important for all passengers to have an easy access to the bus, but it is especially important for people with mobility challenges to have a safe and unimpeded path to the bus. We developed algorithms that can find the sidewalks near a bus stop and then determine if they are covered with snow (Figure 3). These can then be used to trigger a notification to the transit agency or other users who then can act on this information. A detailed description of this application can be found in the attached paper “Leveraging Structure from Motion to Localize Inaccessible Bus Stops”¹

The second application we developed is detecting full trashcans. This application was suggested by the transit agency and will help them in their day-to-day operations. Managing trashcans is tedious and labor intensive. Emptying them when they are not full wastes time and money, emptying them too late makes a bus stop unwelcoming and makes people less likely to take public transportation. Knowing the status of the trashcans can therefore help optimize operations. Figure 4 illustrates the detection process.




Figure 4 Trashcan detection and analysis. In a first step on the bus side, the system finds the trash can. The image is sent to the server where a more powerful detector confirms the detection and a second detector determines if the trashcan is empty, full, or has a full trash bag next to it.


The attached paper “Detecting and classifying bus stop trash cans using camera-equipped public transit vehicles”² describes the application in detail.

In the third application, we focus on developing methods for extracting the spatial attribute and the temporal attribute from urban visual data. Specifically, we introduce a method of organizing large-scale urban visual data into a spatial-temporal structure by mining attributes inherent in the data. We demonstrate the effectiveness of our method by using videos from the bus to detect and analyze work zones within the captured videos. The results are illustrated as heat maps in Figure 5 and Figure 6. The details of the method and the results are attached as “Mining Spatio-Temporal Attributes of Anomalies through Large Ego-Vehicle Dataset”³

¹ To be published as: I. Panigrahi, T. Bu, and C. Mertz, "Leveraging Structure from Motion to Localize Inaccessible Bus Stops," Robotics Institute Summer Scholars Working Papers Journal 2022, vol. 10. 2022.

² To be published as: T. Storm and C. Mertz, “Detecting and classifying bus stop trash cans using camera-equipped public transit vehicles,” Robotics Institute Summer Scholars Working Papers Journal 2022, vol. 10. 2022.

³ T. Ma, “Mining Spatio-Temporal Attributes of Anomalies through Large Ego-Vehicle Dataset”, M.S. thesis, technical report CMU-CS-22-113, May 2022, <http://reports-archive.adm.cs.cmu.edu/anon/2022/abstracts/22-113.html>



(a) Heat map of spatial changes along run R_i



(b) Image at L_i



(c) Image at L_{i+2}




(d) Image at L_{i+14}



(e) Image from L_{i+15}

Figure 5 A heat map is indicating changes across the spatial axis and their corresponding images.

Changes in Anomaly Score Along the Temporal Dimension



(a) Heat map of temporal changes at location L_i



(b) Image at T_{i+1}



(c) Image at T_{i+2}



(d) Image at T_{i+3}



(e) Image from T_{i+4}

Figure 6 A heat map is indicating changes across the temporal axis and their corresponding images.

Our fourth application is High-Definition (HD) map updates. HD maps are very important for autonomous vehicles. They accurately describe all the traffic lanes to within a few centimeters and include information like pedestrian crossings. It is expensive to keep these maps up-to-date. Observing changes of the roads on a daily basis with our bus-edge system is an efficient way to update the maps. In this application, we concentrated on zebra crosswalks. Figure 7 shows six different examples of crosswalk changes that the system detected.

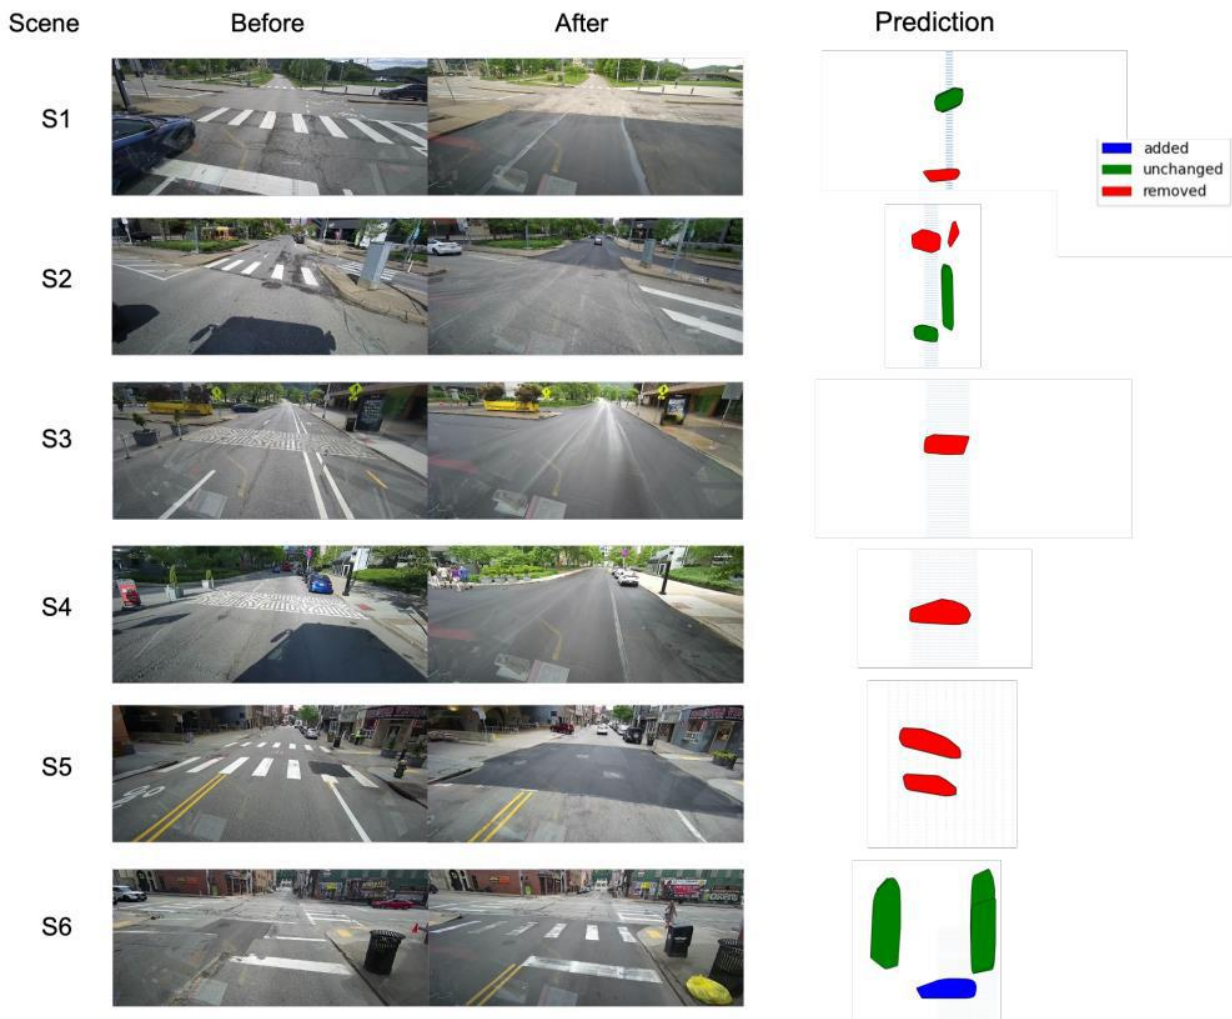


Figure 7 For six different scenes, one image each from the query and reference sequences is shown along with the change prediction. In each of the six scenes, one or more crosswalks have changed. Scenes 1-5 have crosswalks removed, while scene 6 has one plain crosswalk transformed into a zebra crosswalk.

The algorithms and methods we developed and the evaluation of the system is described in the attached “Towards HD Map Updates With Crosswalk Change Detection From Vehicle-mounted Cameras”⁴.

Figure 8 shows a concept how the results of these four applications could be displayed on an interactive map. The applications we developed will have several benefits to the transit agency. HD map update or detection and analysis of road construction will give them a new revenue stream, as the users of these applications will have to pay for the service.

⁴ T. Bu, “Towards HD Map Updates with Crosswalk Change Detection from Vehicle-Mounted Cameras”, Master's Thesis, Tech. Report, CMU-RI-TR-22-34, Robotics Institute, Carnegie Mellon University, August, 2022, <https://www.ri.cmu.edu/publications/towards-hd-map-updates-with-crosswalk-change-detection-from-vehicle-mounted-cameras/>

Detection of problems at bus stops will make their operations more efficient. Noticing when there is snow covering the sidewalks around the bus stop will help them with their core mission: making transportation accessible to all users.



Figure 8 Concept of an interactive map where the user can see the results of our applications.

Summary of publications, code, and datasets

Snow and sidewalk detection:

I. Panigrahi, T. Bu, and C. Mertz, "Leveraging Structure from Motion to Localize Inaccessible Bus Stops," Robotics Institute Summer Scholars Working Papers Journal 2022, vol. 10. 2022. (to be published at https://riss.ri.cmu.edu/research_showcase/working-papers-journals/)

Video "Leveraging Structure from Motion to Localize Inaccessible Bus Stops" by Indu Panigrahi: <https://youtu.be/sUctgSnSARM>

Code: https://github.com/ind1010/SfM_for_BusEdge

Dataset: <https://www.kaggle.com/datasets/indupanigrahi/busedge-sidewalks-and-more>

Poster: I. Panigrahi, T. Bu, and C. Mertz, "Leveraging Structure from Motion to Localize Inaccessible Bus Stops", RISS 2022 poster session, to be published at https://riss.ri.cmu.edu/research_showcase/research-posters/

Trashcan detection:

T. Storm and C. Mertz, "Detecting and classifying bus stop trash cans using camera-equipped public transit vehicles," Robotics Institute Summer Scholars Working Papers Journal 2022, vol. 10. 2022. (to be published at https://riss.ri.cmu.edu/research_showcase/working-papers-journals/)

Video “Detecting and Classifying Bus Stop Trash Cans” by Tim Storm:

<https://www.youtube.com/watch?v=cgPlnbnJLw>

Poster: T. Storm and C. Mertz, “Detecting and classifying bus stop trash cans”, RISS 2022 poster session, to be published at https://riss.ri.cmu.edu/research_showcase/research-posters/

Anomaly detection:

T. Ma, “Mining Spatio-Temporal Attributes of Anomalies through Large Ego-Vehicle Dataset”, M.S. thesis, CMU-CS-22-113, May 2022, <http://reports-archive.adm.cs.cmu.edu/anon/2022/abstracts/22-113.html>

HD map update:

T. Bu, “Towards HD Map Updates with Crosswalk Change Detection from Vehicle-Mounted Cameras”, Master's Thesis, CMU-RI-TR-22-34, Robotics Institute, Carnegie Mellon University, August, 2022

<https://www.ri.cmu.edu/publications/towards-hd-map-updates-with-crosswalk-change-detection-from-vehicle-mounted-cameras/>

Code to be submitted to: https://github.com/tom-bu/crosswalk_change_detector

Dataset: <https://www.kaggle.com/datasets/buvision/crosswalkchange>