Work Order No. 2021-03-UT-ATD

Video Analytics for Vision Zero

FINAL REPORT

Prepared by the Center for Transportation Research at

The University of Texas at Austin

for the

City of Austin Transportation Department

UTA19-000382, Work Order No. 2021-03-UT-ATD

DO 2400 21042307532

December 2021

THE UNIVERSITY OF TEXAS AT AUSTIN
CENTER FOR TRANSPORTATION RESEARCH
3925 W. Braker Lane • Suite 4.11080 • Austin, TX 78759 • (512) 232-3100 • ctr.utexas.edu

Contents

Executive Summary

Work Order 2021-03 refines the tools and work conducted in earlier Video Analytics for Vision Zero tasks by testing the use of automated object recognition/tracking in video data streams to support the assessment of pedestrian safety using existing traffic monitoring cameras. The specific tasks pursued in the scope of WO 2021-03 include: 1) refine the video recording workflow; 2) video recording and analysis at the Lamar/Rundberg and Lamar/Payton Gin locations; 3) analysis of bus stop activities and crossing event detection; 4) validation of crossing detection with PHB device activations; and 5) creating a web interface for interactive video reviewing. The two use cases examined in work task 2 illustrate how a scalable tool for automated analysis of data collected from monocular traffic cameras can allow agencies to leverage existing infrastructure in analysis and mitigation of pedestrian safety concerns.

The value of this work lies not with the result of this particular case study, but on a computational approach to conducting correlation analyses using a street view captured by a traffic camera. As the methodology for the computational approach is improved this work can be scaled to other locations or other use cases, such as in determining if mid-block crossings are a result of nearby bus stops. If a correlation is found and identified early, addressing pedestrian safety issues at mid-block crossings could prevent loss of life.

The web interface leverages a framework developed for creating web applications on cyberinfrastructure, allowing users to interactively explore video clips and summary statistics. After initial login, all traffic camera locations are shown in a map. Once recording results are selected, detailed results are shown and users may select to review results from different use cases. Summary statistics compiled using the video detection framework are shown as boxplot and stack column charts to give users an overview of the inference results. The user can also export the graph and data as csv or image files for sharing or additional studies.

1. Introduction

The rise in pedestrian fatalities in the US over the past few years has led many transportation agencies to refocus efforts on implementing proven countermeasures to improve pedestrian safety. As approximately 75 percent of pedestrian fatalities occur at mid-block locations, focus is primarily placed on strategies to facilitate safer pedestrian crossings. One solution, installation of a pedestrian hybrid beacon (PHB), has been shown to reduce serious injury and fatal crashes by 15 percent and pedestrian crashes by nearly 70 percent [1]. In the past decade the City of Austin (CoA) has installed more than seventy-five PHBs at locations characterized by documented pedestrian demand, crash history, and long distances between safe crossing opportunities, among other factors. When studying potential PHB locations, CoA has traditionally relied on manual observations to quantify pedestrian movements, such as counting the number of crossings during a given time period. However, manual observation provides information only for limited time periods, making it challenging to quantify the evolution of pedestrian behavior over time or in response to a treatment such as PHB.

The collection and analysis of video data at critical locations provides an opportunity to analyze pedestrian movements and to provide a verifiable account of road user behavior. The former reduces the need to rely on ad hoc decision making [2]. However, when analyses are conducted by human observers there is a limit to the number of locations and analysis periods that may be considered. Automated approaches to effectively recognize, analyze, and store pedestrian activities over time are needed. The technical challenges associated with pedestrian activity analysis using video data from traffic cameras are different from those faced when conducting traffic flow analyses. Regular roadside cameras have wide and deep fields of view. Pedestrian activities occupy only a small portion of the field and, at many locations, are only present sporadically. Further, pedestrians appear smaller in size than cars and are more frequently subject to visual obstructions from other objects within the scene.

Incorporating Internet of Things (IoT) and smart devices within an intelligent transportation system (ITS) usually comes with substantial up-front costs for installation and deployment. At the same time, advances in algorithm development and software design bring new opportunities to increase utilization of existing transportation infrastructure. To address such challenges, we have continued to develop a video processing pipeline [3] to improve pedestrian detection and tracking. We have implemented additional features, including select areas of interest from video frames, reporting additional measures of object detections, incorporating heuristics based on best practices, additional visualization summarization, and manual video curation utilities.

2. Tasks Performed

2.1 Task 1: Refine the Video Recording Workflow

The video recording and processing pipeline has been reworked to facilitate new security requirements and new use-case requirements. The new recording pipeline is shown in Figure 1.

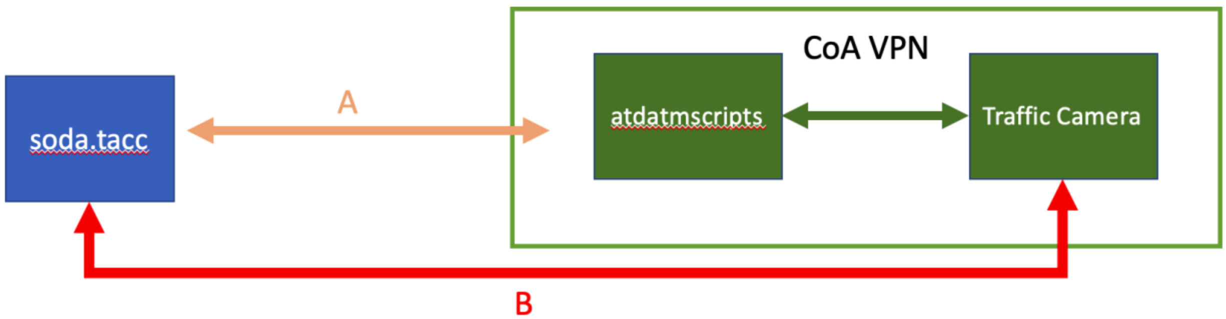

Figure 1. Video recording pipeline: A) Establish connection between video recording server and CoA network; B) Push stream from the selected traffic camera to recording server.

The video recording pipeline requires two separate steps. First, a secure connection must be established between the video recording server (soda.tacc.utexas.edu) and the CoA traffic camera management node (atdatmscripts) through the VPN (Step A). Once the secure connection is established, a command must be run on atdatmscripts to push the video stream from traffic cameras of interest to the video recording node to start recordings. Since Step A requires multi-factor authentication (MFA), the two separate steps are necessary. It also prevents a fully automated workflow for recording video on-demand due to the MFA requirement. Regardless, the new recording pipeline is much more simplified when compared to the previous workflow and it is more stable over longer periods of time.

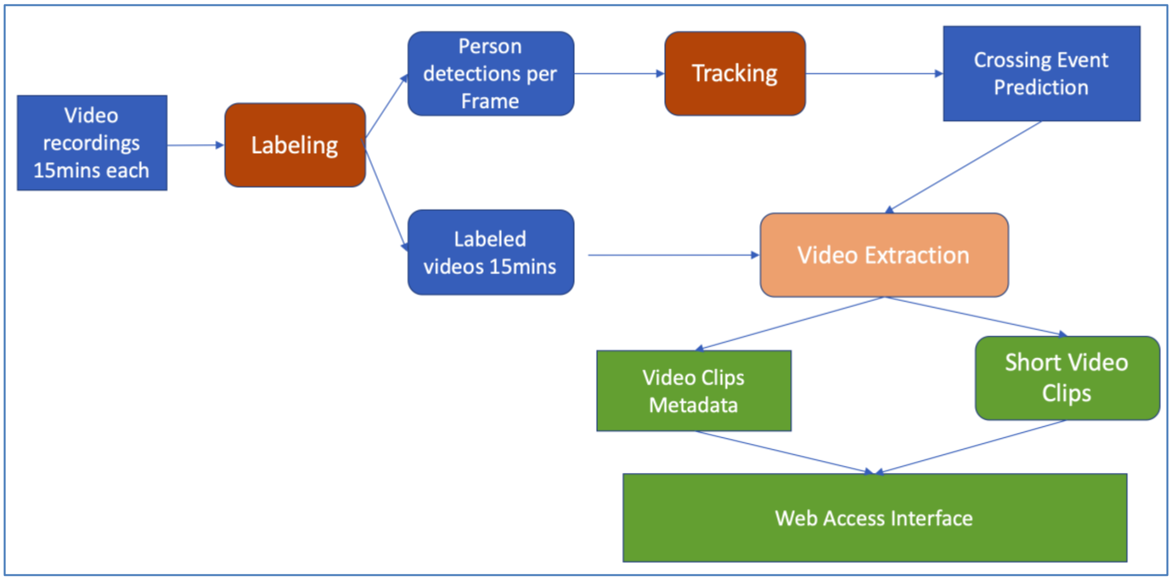

Figure 2. Video processing workflow overview.

Figure 2 shows the video processing workflow. In this redesigned workflow, we target the use case of pedestrian crossing event detection. In addition to an output for the prediction of pedestrian detections, we have added a component to automatically extract a twenty-second clip for each detection from the raw video footage. The extracted video clips are then organized and stored in a pre-defined hierarchical structure on disk so that the video clips can be procedurally accessed and served from the website for further review (see details in Task 5).

2.2 Task 2: Recording and Analysis of Pedestrian Crossing at Lamar and Payton Gin

Raw videos originate from IP cameras in the CoA private network, which has limited accessibility. To facilitate access, CoA set up a proxy server to forward selected video feeds from the IP cameras to a storage cluster hosted at the Texas Advanced Computing Center (TACC), where the recorded video can be processed by a high-performance computing cluster. Processed data is saved in a storage server, which is accessed by the Vision Zero project server for results dissemination purposes. The core algorithm utilizes a convolution-neural-network-based object detection system, YOLOv2, to analyze each frame of an input video (Redmon et al., 2016; Redmon and Farhadi, 2017). For each frame, the algorithm outputs a list of objects that includes their locations in the frame, class label, and confidence of recognition. We have defined recognition according to seven classification labels that are most relevant: person, car, bus, truck, bicycle, motorcycle, and traffic light. To improve algorithmic performance and maximize utilization of multi-node computing clusters, we have also adapted the YOLOv2 implementation for parallel execution (Huang et. al., 2017).

A location at Lamar Boulevard north of Payton Gin Road, hereafter referred to as the Payton location, was selected by the CoA Vision Zero team for a focused study on pedestrian crossing detections. There are two bus stops visible from the camera mounted at the Payton Gin Road and Lamar Boulevard intersection. The bus stops are located on Lamar Boulevard, north of the intersection. In this effort, our goal is to infer bus stop activity (i.e., how many people are waiting for buses throughout the day?) using video from the camera. A view of the camera and locations of two bus stops are shown in Figure 3. We have studied traffic camera videos recorded during period from March 7, 2020, to March 19, 2020, at this location. The recording for each day is from 10:00 to 20:00. The raw recording data requires approximately 40–50 GB of storage per day, with a total of approximately 700GB of storage needed for the month of recorded video.

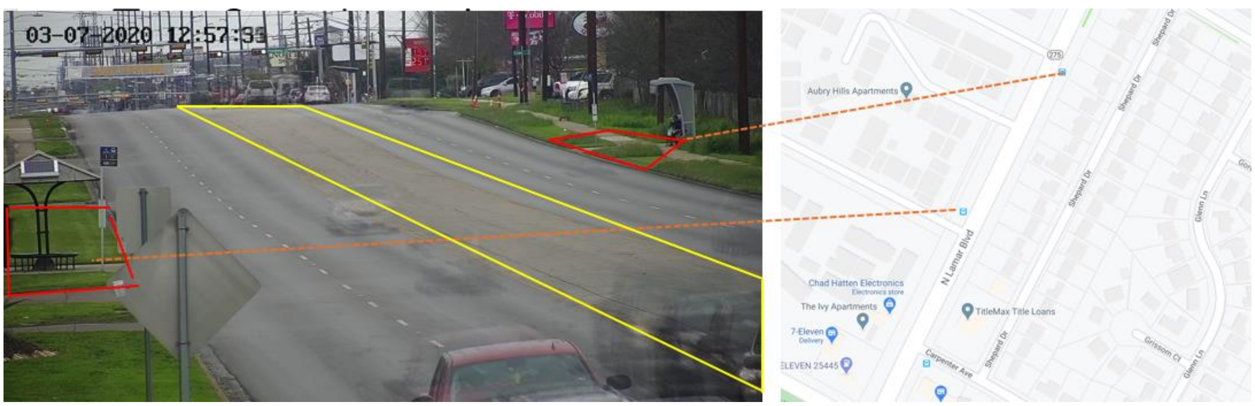

Figure 3. Illustration of bounding boxes and their locations on map.

To support different types of pedestrian activity detection, we enabled the user to self-define the region-of-interest (ROI) over the view of the video footage. The region-of-interest can be simply defined as a set of arbitrary convex shapes overlayed on the footage. Multiple groups of region-of-interests can be defined and analyzed separately. Figure 3 shows an example of an ROI along the center turn lane of Lamar Boulevard (yellow bounding box) and ROIs on two bus stops—one on each side of the road (red bounding box).

Visual summaries of person detections during the weekdays are illustrated in Figure 4. Figure 4 shows crossing detections in the Lamar Boulevard median using the video processing framework. Pedestrian crossing activities are colored by time of day, with yellow representing the AM peak period, green the midday period, blue the afternoon off-peak, and red the PM peak period. Based on a qualitative review of the crossings shown in Figure 4, most crossings occur during the PM peak period, with midday crossings being the second busiest time of day.

Figure 4. Visual summary results at Payton location during weekdays (03/09/2020–03/13/2020 and 03/16/2020—03/19/2020). Detections at bus stops (left) and at middle of roads (right) are colored for four time periods: 7:00–10:00, yellow; 10:00–13:00, green; 13:00–16:00, blue; 16:00–19:00, red.

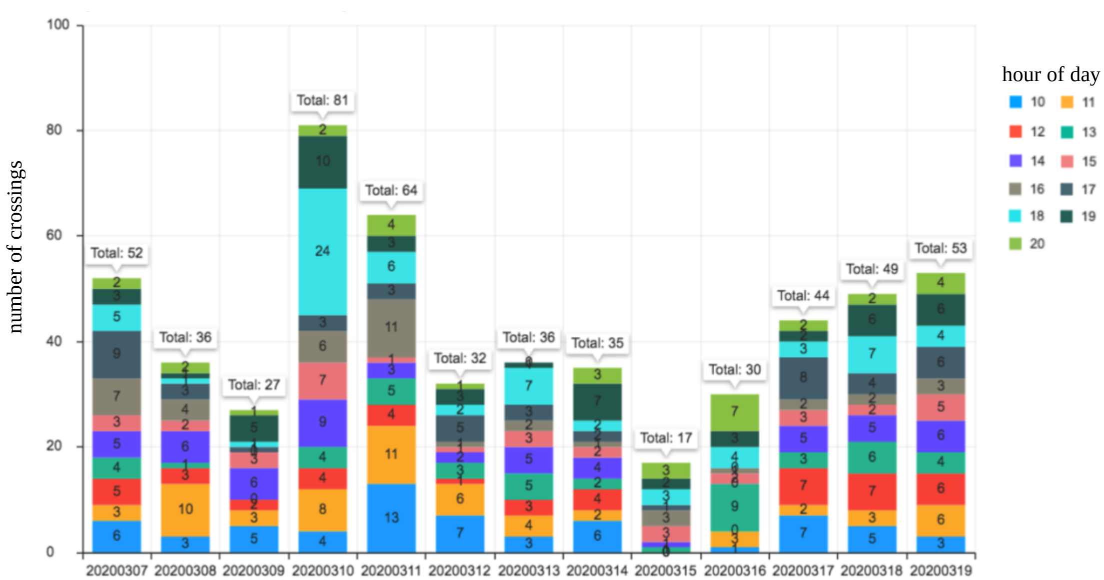

Figure 5. Number of crossing events inferred per hour per day.

We have inferred the number of mid-block crossing per hour per day for days during which data was collected, based on the algorithm described in (Xu et. al., 2019), as shown in Figure 5. Inferred crossing events for each day are illustrated as a stacked column. For each column, detection for each hourly block is illustrated with different colors, starting with 00:00 at the bottom and ending with 23:00 at the top. The number in each color block represents the number of inferred crossings in that particular hour. Daily crossing events inferred for this period range from 17 to 81 with a median of 36 crossing events per day. The results also show large variations among total crossings of different days as well as large variations in hourly inferences for some hours.

Figure 6 shows a box plot of average hourly detections for pedestrian crossings by time of day at the mid-block location studied. The lines extending from the boxes indicate variability outside the upper and lower quartiles. On average, there are 4.2 crossing events inferred per hour.

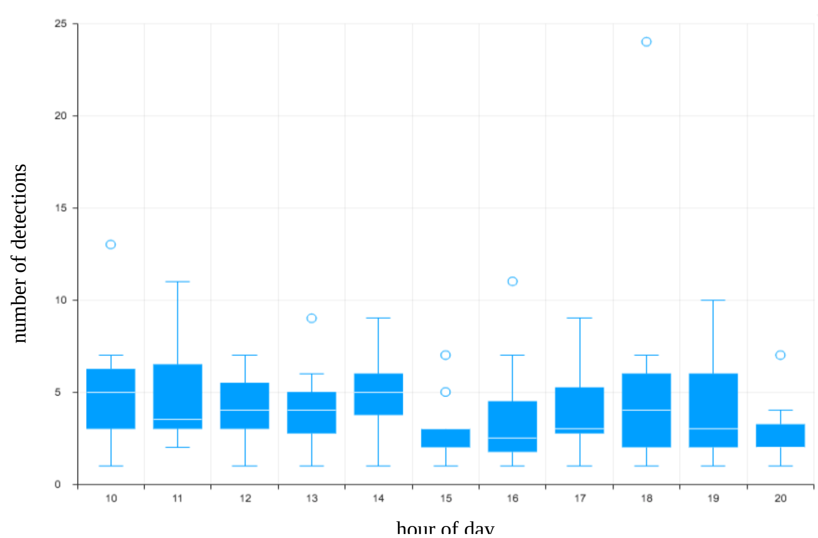

Figure 6. Summary of average hourly detections and outliers (circles on the plot). The box plot is computed by excluding outliers to show minimum, 25 percent quartile, median, 75 percent quartile, and maximum values.

2.3 Task 3: Analysis of Bus Stop Activities at Lamar and Payton Gin

The number of people waiting for buses in each hour over the recording periods are inferred according to the method described in Figure 7. Since the bus stops are located on the sidewalks where other foot traffic occurs frequently, we needed to develop a process to infer the number of people waiting for a bus. We first identified time intervals (referred as sessions) during which people are continually detected within the regions of interests. To reduce the impact of occasional false positive identifications, a minimum time threshold (min_session_time) is used so that each session is longer than the specified threshold. When waiting at bus stops, people often move very little and are easily blocked from the video camera by other objects. This can lead to false negative detection since those people will not be detected with the described methodology. A second threshold, min_no_detection, specifies a minimum time window that separates two consecutive sessions. If two sessions are within the min_no_detection threshold they will be merged into one, as it is assumed the two sessions are tracking the same single person. Once sessions are defined, the number of unique persons are further inferred within each session. Figure 7 summarizes the algorithms used to infer the number of people waiting for buses.

We chose five seconds as the threshold value for min_no_detection and fifteen seconds as the threshold value for the min_session_time. The bus stop shown on the right side of frame (Figure 3) is the bus stop for northbound traffic (NB). The bus stop on the left side of frame (Figure 3) is the bus stop for southbound traffic (SB). Figure 8 shows the hourly number of people inferred at the NB stop for each day. The inferred total number of people at the NB stop ranges from 5 to 22 with a median of 12 for each day. However, the inferred total number of people waiting at the SB stop has a much larger variant (between 4 to 84, with a median of 33).

Pseudo code for identify number of people waiting at bus stops.

Input: ***D***: sequence of detections {***f* , *b*** | ***f***: frame index, ***b***: bounding box of detected object in the frame}, ***ROI***: regions of interest, ***min_session_time***, ***min_no_detection***

Output: ***S***: sequence of sessions {(**f_b, f_e, *p***)| f_b: begging of a session, f_e end frame index of a session, p is the number of unique persons within the session}

1: D D_{ROI} filter out all detections in D within region of interest specified by ROI as D_{ROI}

2: D_{ROI} S_D: merge consecutive detections in D_{ROI} to form list of session which is a sequence of detections.

Figure 7. Pseudo code to infer number of people waiting for bus from recognition results.

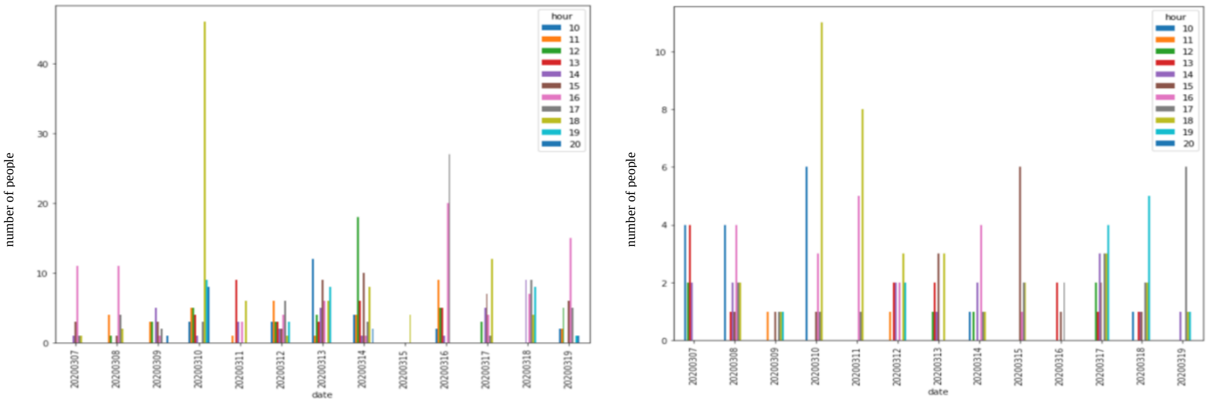

Figure 8. Inference of number of people waiting for bus during each hour of the day for recorded days at the SB stop (left) and NB stop (right).

In Figure 8, notice several significant spikes of the number of persons identified at the bus stop area. For example, the maximum number of persons 57, was inferred for 18:00–19:00 on March 10, 2020, at the SB stop. Further manual review of recordings shows the inferences was due to two persons staying at the bus stop area for an extended time, which caused repeated detections.

In subsequent hourly pattern analysis, we incorporated an outlier detection filter to remove inferences that are clearly out of normal expectations. Further breaking down by hour, the median number of people waiting for buses per hour at the NB stop ranges from 1 to 3 across different hours of the day during the time period. For the SB stop, the median number of people waiting for buses per hour ranges from 1 to 6 across different hours of the day during this time period. When the data is combined, the median of the inferred number of people waiting at both stops per hour (across all hours and all days) is about 4.

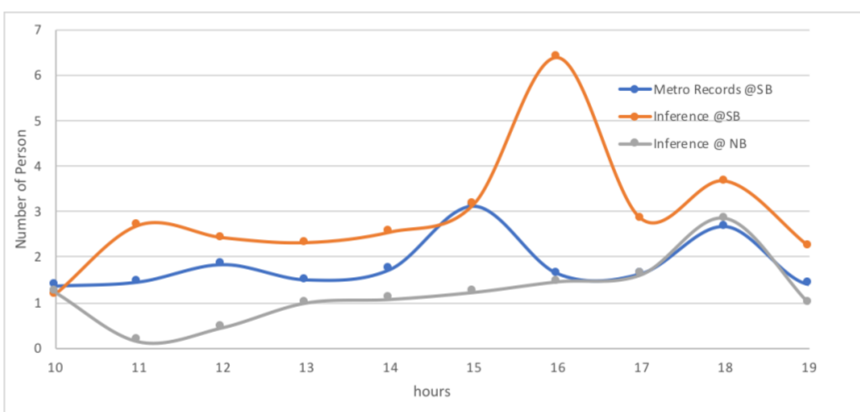

Figure 9. Average of inferred number of persons at bus stops for each hour during the recording periods at SB and NB stop. For comparisons, we also show the boarding records from CapMetro at the SB stop.

To check the accuracy of number of persons inferred, we obtained boarding records at SB stop during the same time periods from the bus operator, CapMetro (Figure 9). The results show our approach inferred higher values than the operator's records. The mean absolute error (MAE) between our estimate with CapMetro record is 1.14. A primary reason for the discrepancy is that our approach is based on persons detected in the bus stop area rather than a direct inference for the number of people boarding the bus, as recorded by CapMetro. There are instances where person is detected around the bus stops but never boarding a bus. Manual video review confirmed those observations.

Table 1. Correlation analysis between inferred crossing events with inferred number of people waiting at SB, NB, and both bus stops.

Time	**10:00**	**11:00**	**12:00**	**13:00**	**14:00**	**15:00**	**16:00**	**17:00**	**18:00**	**19:00**
#SB	-0.19	0.20	0.01	0.39	0.14	0.05	-0.07	-0.36	**0.91**	0.40
#NB	-0.08	-0.03	0.44	0.37	0.12	0.10	0.50	0.14	**0.79**	0.06
#BOTH	-0.22	0.19	0.07	0.54	0.18	0.09	0.09	-0.31	**0.92**	0.35

To study connections between people at bus stops and inferred crossing events, a correlation analysis has been conducted for each hour of the day. In the correlation analysis, inferred values for the same hour on different days are treated as a series. Therefore, for each hour, four time series are computed: number of crossing events (CRO), number of people waiting at the NB stop (#NB), number of people waiting at the SB stop (#SB), and the sum of people waiting at both stops (#BOTH). Each time series contains thirteen data points for each day. Table 1 shows correlation results for each hour between CRO with #SB, #NB, #BOTH. Typically, a strong correlation between two series will result a value between 0.7–1.0. Correlation values between 0.3–0.7 can be considered as moderate correlations. The results only show strong correlation (with correlation score greater than 0.7) for 18:00–19:00 at both SB and NB stop. However, this correlation may be due to the excessive false detections often observed during that hour.

2.4 Task 4: Validation of Pedestrian Crossing with PHB Activations

A PHB was installed at the Lamar and Cooper Drive intersection and is viewable from the Lamar/Rundberg Lane camera (Figure 10). We have recorded video from July 19, 2021, to July 30, 2021, at this location. Each day has about fifty-two videos at fifteen minutes recording length. The raw recording is about 20GB per day with about 500GB total.

Figure 10. Visual summary of detection from the Lamar/Rundberg camera from July 19, 2021 to July 30, 2021. The detections are separated into two regions of interest. On the left, a mid-block crossing at the bottom part of the road (between Cooper Drive and Rundberg Lane). On the right, the area is focused on the crosswalk area installed at the Cooper Drive intersection.

An overview of all detections, separated into two regions of interest, is shown in Figure 10. On the left, it shows a mid-block crossing at the bottom part of the road (between Cooper Drive and Rundberg Lane). On the right, detection is focused on the crosswalk area at the Cooper Drive intersection. Detection results are shown in Figure 11.

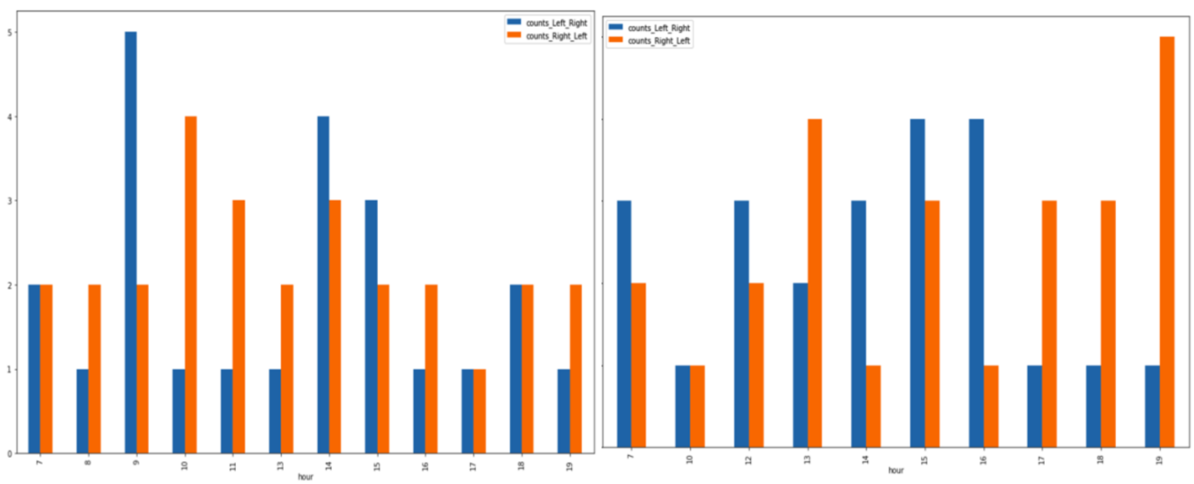

Figure 11. Average hourly detection for midblock crossing (left) and use of crosswalk (right).

We obtained the PHB activation records to compare with prediction results. For comparison purposes, we focused on PHB activation between 07:00 and 21:00 each day. The daily comparison is shown in Figure 12.

The accuracy for each day is computed as:

$$accuracy = 1 - \frac{|\,total\ predicted\ crossing\ at\ crosswalk - total\ PHB\ activations\ (7-21)|}{total\ PHB\ activations\ (7-21)}$$

The average accuracy over the eleven-day period is 0.89 with an error margin ranging from 0–19 percent.

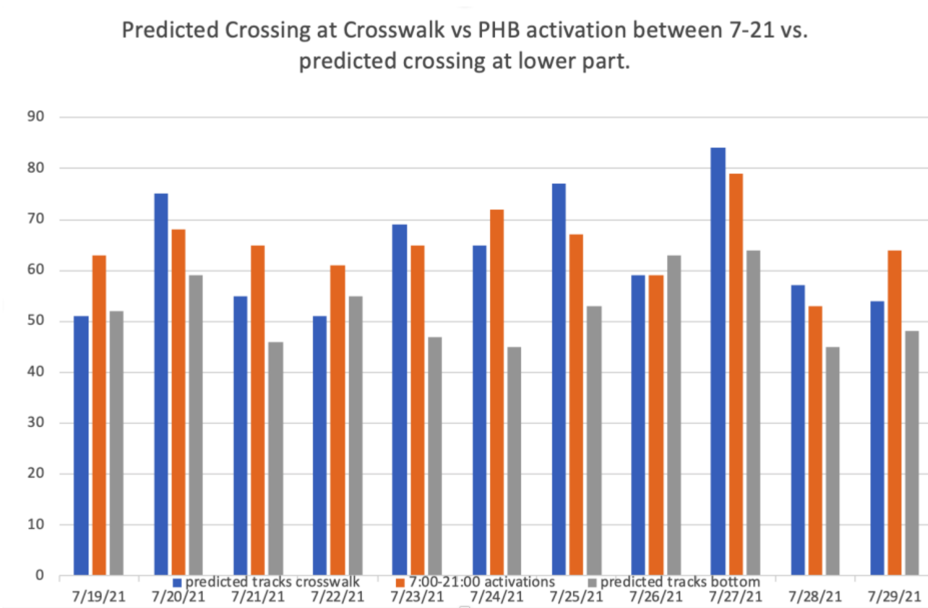

Figure 12. Comparison of daily predicted crossing using crosswalk (blue), PHB activations during 7-21 hours per day (red), and predictions of mid-block crossing inference (grey).

2.5 Task 5: Web Interface for Interactive Exploration

While our approach aims to provide robust and automated methods for detecting and inferring pedestrian activities, results indicate outliers occur frequently among the inference results. These outliers can be results from both model limitations and actual unexpected events on the road. To support human intervention and to gain more insights from the inference results, we have implemented a web interface for reviewing video clips (soda.tacc.utexas.edu). The web interface leverages our framework developed for creating web applications on cyberinfrastructure (Xu et.al., 2019). To further reduce data risks and protect personal privacy, the web user interface supports authentication so that only people with account credentials can access the raw video footage (Wang et.al., 2018).

The website is password protected and can only be accessed via a granted credential account. After initial login, all traffic camera locations are shown in a map. If past recordings are available, a solid camera icon will be used instead (Figure 13). A user can mouse over the solid camera icon for more details and a link to detailed results.

Once the recording results are selected, detailed results are shown on a new page (as illustrated in the screenshot of the web interface shown in Figure 14). At the top, users may select to review results from different use cases. Summary statistics compiled using the video detection framework are shown as boxplot and stack column charts to give users an overview of the inference results.

The boxplot on the left shows average hour pedestrian crossing events for the chosen date range; the bar diagram on the right displays the daily number of detections for each recorded day, with each hour represented by a different color band. The graphs are interactive: Users can click on a time segment to see a list of clips associated with inferences during that time segment. The user can also export the graph and data as csv or image files for sharing or further study. The clip list table can also be interactively navigated by users to bring up clips of detections, as shown on the left side.

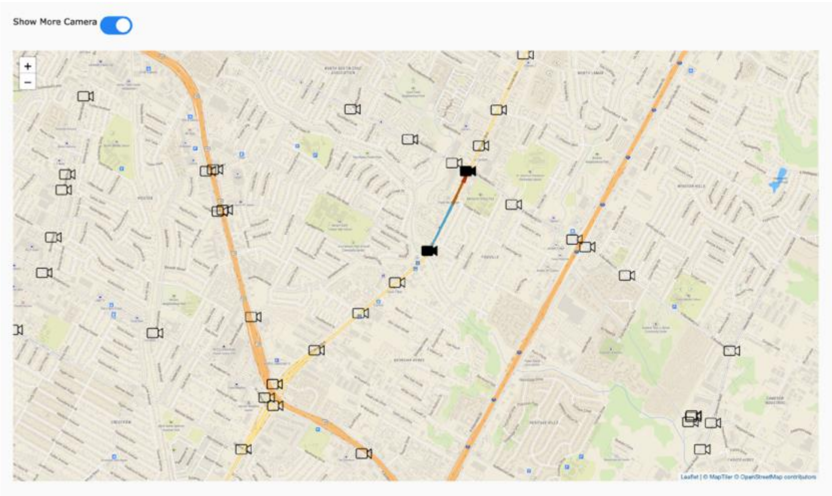

Figure 13. Map view of all camera locations. When recordings are available, a solid camera icon will be shown with indication on the direction and road segment of the recordings.

Figure 14. A web interface for interactive detection video reviewing purposes.

In Figure 14, the video still shown in the bottom left shows an example of the object labeling and tracking algorithm applied to Lamar Boulevard north of Payton Gin Road. Pedestrians are identified and crossings tracked, as shown by the pink bounding boxes in this video. Note that other objects are also identified and tracked (car, bus, truck, bicycle, motorcycle, and traffic light). The data for all objects is stored with no personal identifiable information, and it can be used for future research efforts that have yet to be identified. The user can click on individual video clips to the right of the video image, to observe each pedestrian crossing. The user can also download the individual video clips of reviewing offline.

3. Findings, Summary, and Discussion

The City of Austin has more than 400 CCTV cameras installed at intersections in the Austin area. These cameras are commonly used for manual traffic monitoring, with no long-term recording or archiving of videos. Artificial intelligence technologies can greatly reduce the effort involved in analyzing video data. The framework presented here can facilitate research traditionally based on manual field and video data analysis. The intent is that this work effort will promote further work on video data applications and integration. A unique advantage of our framework is that it converts video recordings into query-able information while not saving personally identifiable information, which can accommodate multiple subsequent use cases without re-processing (Huang et al., 2017) or risks related to privacy issues.

The use cases presented in this work illustrate the benefits and limitations of the proposed methodology. Our video aggregation pipeline has the potential to support long-term pedestrian activity monitoring. The flexibility of the data selection and filtering capabilities is expected to enable further applications. In addition to the visual summaries described in this study, quantitative outputs can be generated to facilitate the comparison of conditions across different locations or time ranges, and to evaluate the impact of infrastructure changes and construction scenarios, among others.

Bus stops present additional challenges in detecting pedestrian activities. Although our motivation is to identify how many people are waiting for buses and how long they have been waiting, we also identified cases where pedestrians dwell at bus stops extended time periods that appear unrelated to bus activities. These observations can add bias to bus waiting time predictions. Furthermore, there are high numbers of duplicative detections over time, as pedestrians tend to move less while waiting for a bus to arrive.

Our case study leverages the latest AI technology to infer both crossing events and number of pedestrians waiting at bus stops in lieu of a manual field count. This approach has the potential to scale to other locations and for longer periods of time than is possible using human data collection methodologies. In the use cases reported here, we focused on one selected location in Austin and analyzed traffic videos collected during a period in March 2020. We reported our inference results for pedestrian mid-block crossings and bus stop activities and conducted a correlation analysis between the two variables for each hour of the day. Our results show no overall strong correlations are observed between person detection in bus stops with mid-block crossing for this data sets. To assist traffic engineering practitioners in reviewing and to gain insights from the recording and the inference results, we also built a secure web interface to allow a user to interactively explore inference results.

The value of this work lies not with the result of this particular case study, but on a computationally aided approach to conduct correlation analyses using a street view captured by a nearby traffic camera. As the methodology for the computational approach is improved, this work can be scaled to other locations or other use cases, such as determining if mid-block crossings are a result of the nearby bus stops. If a correlation is found and identified early, addressing pedestrian safety issues around mid-block crossings could prevent loss of life.

The web interface leverages the framework developed for creating web applications on cyberinfrastructure. It allows users to interactively explore video clips and summary statistics. After initial login, all traffic camera locations are shown on a map. Once recording results are selected, detailed results are shown and users may select to review results from different use cases. Summary statistics compiled using the video detection framework are shown as boxplot and stack column charts to give users an overview of the inference results. The user can also export the graph and data as csv or image files for sharing or additional studies.

References

Sayed, T., M. H. Zaki, J. Autey (2013). Automated safety diagnosis of vehicle–bicycle interactions using computer vision analysis, *Safety Science*, vol. 59, pp. 163–172.

F. M. Puscar, T. Sayed, A. Y. Bigazzi, and M. H. Zaki, “Multimodal Safety Assessment of an Urban Intersection by Video Analysis of Bicycle, Pedestrian, and Motor Vehicle Traffic Conflicts and Violations,” 2018.

St. Aubin, P., N. Saunier, L. Miranda-Moreno (2015). Large-scale automated proactive road safety analysis using video data. *Transportation Research Part C: Emerging Technologies*, vol. 58, part B, pp. 363–379.

Kuciemba, S., K. Swindler (2016). Transportation Management Center Video Recording and Archiving Best General Practices. Federal Highway Administration report no. FHWA-HOP-16-033.

Zangenehpour, S., L. F. Miranda-Moreno, N. Saunier (2015). Automated classification based on video data at intersections with heavy pedestrian and bicycle traffic: methodology and application. *Transportation Research Part C: Emerging Technologies*, vol. 56, pp. 161–176.

Kastrinaki, V., M. Zervakis, K. Kalaitzakis (2003). A survey of video processing techniques for traffic applications, *Image and Vision Computing*, vol. 21, issue 4, pp. 359–381.

Weijia Xu, Natalia Ruiz-Juri, Ruizhu Huang, Jennifer Duthie, and John Clary. 2018. Automated pedestrian safety analysis using data from traffic monitoring cameras. In Proceedings of the 1st ACM/EIGSCC Symposium on Smart Cities and Communities (SCC '18). ACM, New York, NY, USA, Article 3, 8 pages. DOI: https://doi.org/10.1145/3236461.3241972

W. Xu, N. Ruiz, K. Pierce, R. Huang, J. Meyer and J. Duthie, "Detecting Pedestrian Crossing Events in Large Video Data from Traffic Monitoring Cameras," 2019 IEEE International Conference on Big Data (Big Data), 2019, pp. 3824-3831, doi: 10.1109/BigData47090.2019.9005655.

Weijia Xu, Ruiz, N., Huang, R., Duthie, J., Meyer, j., Clary, J., (2019) Deep learning methods to leverage traffic monitoring cameras for pedestrian data applications, 26th ITS World Congress, Oct. 21-25, 2019, Singapore (Best Technical Paper Awards)

Huang, L., W. Xu, S. Liu, V. Pandey, N. Ruiz Juri (2017). Enabling versatile analysis of large scale traffic video data with deep learning and HiveQL. In Proceedings *2017 IEEE International Conference on Big Data (Big Data)*, Boston, MA

Redmon, J, S. Divvala, R. Girshick, A. Farhadi (2016). You Only Look Once: Unified, Real-Time Object Detection. In Proceedings *IEEE Conference on Computer Vision and Pattern Recognition*, Las Vegas, NV.

Redmon, Joseph, and Ali Farhadi. "YOLO9000: better, faster, stronger." In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 7263-7271. 2017.

Weijia Xu, Ruizhu Huang and Yige Wang, "Enabling User Driven Big Data Application on Remote Computing Resources," *2018 IEEE International Conference on Big Data (Big Data)*, Seattle, WA, USA, 2018, pp. 4276-4284. doi: 10.1109/BigData.2018.8622006

Yige Wang, Ruizhu Huang, and Weijia Xu. 2018. Authentication with User Driven Web Application for Accessing Remote Resources. In Proceedings of *the Practice and Experience on Advanced Research Computing (PEARC '18)*. ACM, New York, NY, USA, Article 2, 7 pages. DOI: https://doi.org/10.1145/3219104.3229290