Using Security Cameras to Count Transit Passengers

Kai Monast, MRP
Christopher Vaughan, PE
Porter Jones
Christopher Carnes

Institute for Transportation Research and Education

NCDOT Project 2017-16
FHWA/NC/2017-16
May 2017

Technical Report Documentation Page

19. Report No. **FHWA/NC/2017-16**	20. Government Accession No.	21. Recipient's Catalog No.
4. Title and Subtitle **Using Security Cameras to Count Transit Passengers**		22. Report Date **May 31, 2017**
		Performing Organization Code
Author(s) **Kai Monast, MRP, AICP, Christopher Vaughan, PE, Porter Jones**		Performing Organization Report No.
Performing Organization Name and Address **Institute for Transportation Research and Education** **North Carolina State University** **Centennial Campus Box 8601** **Raleigh, NC**		Work Unit No. (TRAIS)
		Contract or Grant No.
Sponsoring Agency Name and Address **North Carolina Department of Transportation** **Research and Analysis Group** **104 Fayetteville Street** **Raleigh, North Carolina 27601**		Type of Report and Period Covered **Final Report** **August 2016 to May 2017**
		Sponsoring Agency Code **2017-16**
Supplementary Notes:		

Abstract

The Institute for Transportation Research and Education at North Carolina State University performed research to determine if transit vehicle security camera footage could be paired with Autoscope software to automatically count transit passenger boardings. The research was limited to community transportation systems vehicles. The Autoscope software allows analog footage to be analyzed using the strategic placement of detectors and functions of detectors to determine where the background image changes. After attempting multiple detector placements in multiple vehicle types with multiple camera configurations, the results show promise that algorithms can be set up in a cost effective manner for counting passenger trips repeatedly on the same vehicle, but less promise that a standard algorithm can be developed that will accurately count passenger trips even on the same vehicle type with similar camera configurations. Because this research stems from a specific need for NCDOT to be able to validate passenger trip reporting from transit systems, two sampling plans are suggested as well as expected costs. One involves manual counts using spot-checks of data and the other involves more robust analysis of a single vehicle from each transit system operating deviated fixed route service.

Key Words Security Camera, Transit, Public Transportation, Community Transportation, Deviated Fixed Route, Automatic Passenger Counters	Distribution Statement

Security Classif. (of this report) Unclassified	Security Classif. (of this page) Unclassified	No. of Pages 18	Price

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

DISCLAIMER

The contents of this report reflect the views of the authors and not necessarily the views of the Institute for Transportation Research and Education or North Carolina State University. The authors are responsible for the facts and the accuracy of the data presented herein. The contents do not necessarily reflect the official views or policies of the North Carolina Department of Transportation or the Federal Highway Administration at the time of publication. This report does not constitute a standard, specification, or regulation.

Table of Contents

List of Tables

List of Figures

1. INTRODUCTION

The North Carolina Department of Transportation/Public Transportation Division (NCDOT/PTD) provides federal and state funding to community transportation system subrecipients. A requirement of this sub-allocation of funds is that NCDOT/PTD must collect and report service and financial statistics to the federal government. Recent audit findings have indicated that the Division must take a more proactive role in verifying the service statistics, especially passenger trips. In addition, NCDOT/PTD and the federal government are interested in moving to performance-based budgeting where funding is allocated based on yet to be defined performance measures. It is very likely that these performance measures will include passenger trips.

For these two reasons, it is imperative that the Division be able to verify the accuracy of reported passenger trips. To address this need, this research explores the possibility of using existing security camera footage to develop an automated process for counting passenger trips.

1.1. Research Need

Currently, there are currently two primary methods for verifying passenger trips. The first method is to purchase automated passenger counter technology which is expensive and requires significant data management skills. Community transportation systems lack funding to purchase this technology and they lack the technical skills to manage the data to ensure its accuracy. The second method is to conduct ride-along surveys where an individual rides on the vehicle and counts passenger trips. This method is not only expensive, but also has the drawback of transit systems knowing when a ride-along is occurring and thus being able to adjust their data collection during these periods. In other words, the transit system will be fully aware of when a sample is being conducted, and the sample results may not be fully valid or applicable to non-sample periods.

In 2010, NCDOT/PTD used American Reinvestment and Recovery Act funding to purchase security cameras for almost all publicly-funded community transportation systems vehicles. Vehicles acquired after the initial rollout also include security cameras. These cameras record all activity on the vehicle from multiple angles, depending on the vehicle type. Thus, the data exist to verify reported trip counts in the security camera footage. It is possible to watch the footage and count passenger trips, but this is time-consuming.

As an alternative solution, this research explores the development of an automated tool that will count passenger trips from security camera footage. The video footage is analyzed using Autoscope software which enables the researchers to place detectors in strategic locations for determining when passengers board a transit vehicle. The purpose of this research is to determine if the detectors can be placed in locations that will result in accurate counts of passenger activity and be more cost effective than traditional methods of validating trip counts.

1.2. Scope and Objectives

The objective of this research is to develop algorithms for counting passenger trips using existing security camera footage and test the accuracy of this technology in counting transit passengers. This is a proof of concept research project that attempts to pair existing data with a methodology that has been used in other industries, but not in transit. Surveillance footage from the vehicles is run through AutoScope software using different detector scenarios to determine if the footage can be used to accurately and reliably count passenger trips. The type of exploratory scope involves multiple iterations of trial and error.

2. LITERATURE REVIEW

The majority of respondents to the Transit Cooperative Research Program Synthesis 77 (1) use a combination of automated and manual solutions for tracking ridership data. The most common technology is Automatic Passenger Counters (APCs) plus farebox data. The report states that transit systems with less than 250 peak buses were likely to continue to rely on manual processes. Each APC unit, on the median, cost over $6,000 to purchase and $600 per year in maintenance per year. APCs do not work automatically, but require procedures for processing and reporting data.

Video surveillance cameras have been used to check APC accuracy (2). There are differences between APC counts and manual counts, but these differences were not found to be statistically significant at the 0.05 level on large fixed route buses. Manual counts are presumed to have random errors due to human limitations, whereas APC errors are likely to be both systematic and random. Comparing the APC counts to the manual counts often results in it being necessary to apply a calibration factor to the APC data.

Kotz et. al. received a Transportation Research Board IDEA grant for further study of a concept using vehicles weights to estimate ridership. They found that air suspension systems in larger buses can be retrofitted to include a pressure-sensitivity scale (3). The weight of the vehicle could be subtracted from the scale reading and the results can be divided by an average weight per person to estimate ridership. The pressure measurement system was found to be 87% accurate using an average passenger mass of 168 pounds (4). However, small vehicles like Light Transit Vehicles and high top vans do not have air suspensions systems.

The investigators have been involved in related research that utilizes similar methodology to the one proposed but counting vehicles instead of passengers (5, 6). In addition, the researchers are familiar with the automated passenger counter technology as deployed in transit vehicles as well as how to use automatic passenger counter data for reporting (7, 8).

Searching the TRID database and other sources revealed no previous research conducted directly on this topic. A doctoral dissertation from Utah State University in 2009 explored the capabilities of using special cameras for counting passengers (9). The primary camera vendor in North Carolina, Seon, offers passenger counting with an add-on technology that uses radio frequency identifiers for school bus riders (10). This technology is not possible for general public transit riders.

3. METHODOLOGY

Video data was collected for processing from three North Carolina local transit agencies. These include the Cape Fear Public Transportation Authority (WAVE Transit), the Western Piedmont Regional Transit Authority (WPRTA), and the Johnston County Area Transport System (JCATS). JCATS was not on the initial project study list, but was included later to bolster the sample size and to test the transferability of the algorithms to new environments.

This video was prerecorded, which provided two benefits: 1) the inability for the transit system to inflate passenger counts, and 2) the repeatability of this process using other prerecorded video from other agencies.

Video gathering was a time-consuming process, which is something the NCDOT will need to consider if deadlines are an issue when trying to process video. The length of the collection process was due to a couple of different issues, which include staff turnover at the transit systems, delayed responses from transit agencies, transit staff being unfamiliar with the process for capturing and transferring the data, and unrecognizable video formats for the video detection software. It is believed that, as transit systems become more familiar with the process of obtaining and transferring the data, the process will speed up substantially.

Screenshots of some of the transit agency recorded video are below in Figures 1, 2, and 3.

Figure 1 - WPRTA Video Screenshot

Figure 2 - WAVE Video Screenshot

Figure 3 - JCATS Video Screenshot

The video detection hardware and software used for this proof of concept was Econolite's Autoscope Terra Card and related software. This equipment is intended to be used at traffic signals for virtual detection loops, but ITRE has used it previously in unique implementations like vehicle detection in roundabouts to provide better accessibility for visually impaired pedestrians and wrong-way vehicle detection at the unconventional interchange known as the Diverging Diamond Interchange, or DDI. Because the Terra Card is typically placed in the field for real-time detection, it requires an analog video input from traffic cameras.

The native format of the security footage is digital whereas Autoscope requires analog video. Thus, all video must be reformatted to an understandable format for the Terra Card, which includes converting from digital to analog video or converting between various analog video formats. This often has to be done at real-time playback speeds, which increases the pre-processing time required for the transit videos. For example, a two-hour video requires two hours to be converted to analog. Once video is in a digestible format for Autoscope, it is played at real-time speed through the Autoscope Terra Card and processed through their software. Although the conversion process is time-consuming, it can be done unattended.

3.1. Video Detection

Virtual detectors can be placed anywhere within the confines of the video screen and will observe the pixels directly underneath the detector. If a portion of the pixels underneath a detector change, the detector will activate. "Functions" can be associated with one or more detectors to indicate a certain group of pixel changes in the video. For this project, detectors were placed over the entrance to the vehicle and along the walking path of the vehicle with functions tying these detectors together to capture passengers entering the vehicle.

Ideally, the detection algorithm would only capture passengers as they entered the bus, not when exiting. However, issues that arose include the following:

- counting the bus driver as a passenger;
- counting passengers as they exited the bus, sometime resulting in double counting passengers;
- counting passengers shifting in their seats as boardings;
- counting bags and mobility devices such as walkers as boardings;
- detecting shadows as the bus moved and counting them as passengers;
- and missing passengers altogether

Unfortunately, these issues are difficult to avoid. For instance, in a moving bus, shadows from the light streaming through the windows are unpredictable as the vehicle turns, meaning the algorithm has to be designed to account for them. However, as this software and hardware is intended to be used in a stationary position (at a signalized intersection observing one approach), it cannot completely account for all shadows. Also, the algorithm can be designed more specifically to rule out shadows and exiting passengers, but this results in a less conservative approach for entering passengers, meaning more are missed. Therefore, a fine balance is desired that will reduce false positives (shadows and exiting passengers being counted) while also reducing false negatives (entering passengers being missed).

The algorithm will never be perfect, because as has been mentioned already, this equipment was not designed with this scenario in mind. This limitation was expected because existing Automatic Passenger Counter technologies are also not perfect and must be calibrated and have their data adjusted. Note that there are other traffic detection software and hardware packages that may work better, but they are much more expensive and also often do not allow for user manipulation, meaning the developer has to design the algorithm for you. This is not a cost effective or easily repeatable process which is why Autoscope was chosen.

Autoscope works by taking a background image and then comparing an area of interest, marked by detectors, to the background that it took initially. This background image updates about every minute, but the refresh rate has to be changed depending on the situation. The unpredicatable motions of a vehicle frequently change the shadows inside the vehicle, making it difficult to arrive upon a standard background refresh rate.

When the detector area changes from the original image, Autoscope turns on the detector. Functions are used to say what effect the detector being on or off has. The three main ones used in this are ON, OFF, and M of N. M of N is used for a grouping of detectors, saying that M number of the set of N detectors must be on before the function conditions are met. Thus, it is possible to set up an array of detectors and use programming logic to record movements. Below are the specifics of how these particular algorithms were developed, along with a figure displaying the detectors and functions implemented for one of the buses (Figure 4).

3.2. Algorithm Development

The research team had a general idea of how to set up the detection algorithm. Once video was received from the participating transit agencies, the team was able to begin designing and testing the algorithm which involves hundreds of iterations at real-time playback speeds. The basic idea behind the algorithm is that detectors are placed at the entrance to the vehicle to detect when passengers first enter and this operation triggers the rest of the algorithm ("detector group 1" in Figure 4). Once a passenger enters the vehicle, the algorithm checks for movement across the next set of detectors ("detector group 2" in Figure 4), as it is possible that the door opens and no one enters the vehicle. When the first detector group is associated with the second set of detectors, it indicates the direction the passenger is traveling. This is intended to prevent counting exiting passengers who were already counted when they entered the vehicle. Figure 4 provides a schematic of the general algorithm design. Note that each vehicle required a slightly different algorithm, so they are not all the same. Even vehicles that are in the same class, such as 20-foot Light Transit Vehicles, have slightly different camera placements and angles. Unfortunately, this means that vehicles with similar setups are not identical, meaning that each vehicle required at least some algorithm manipulation.

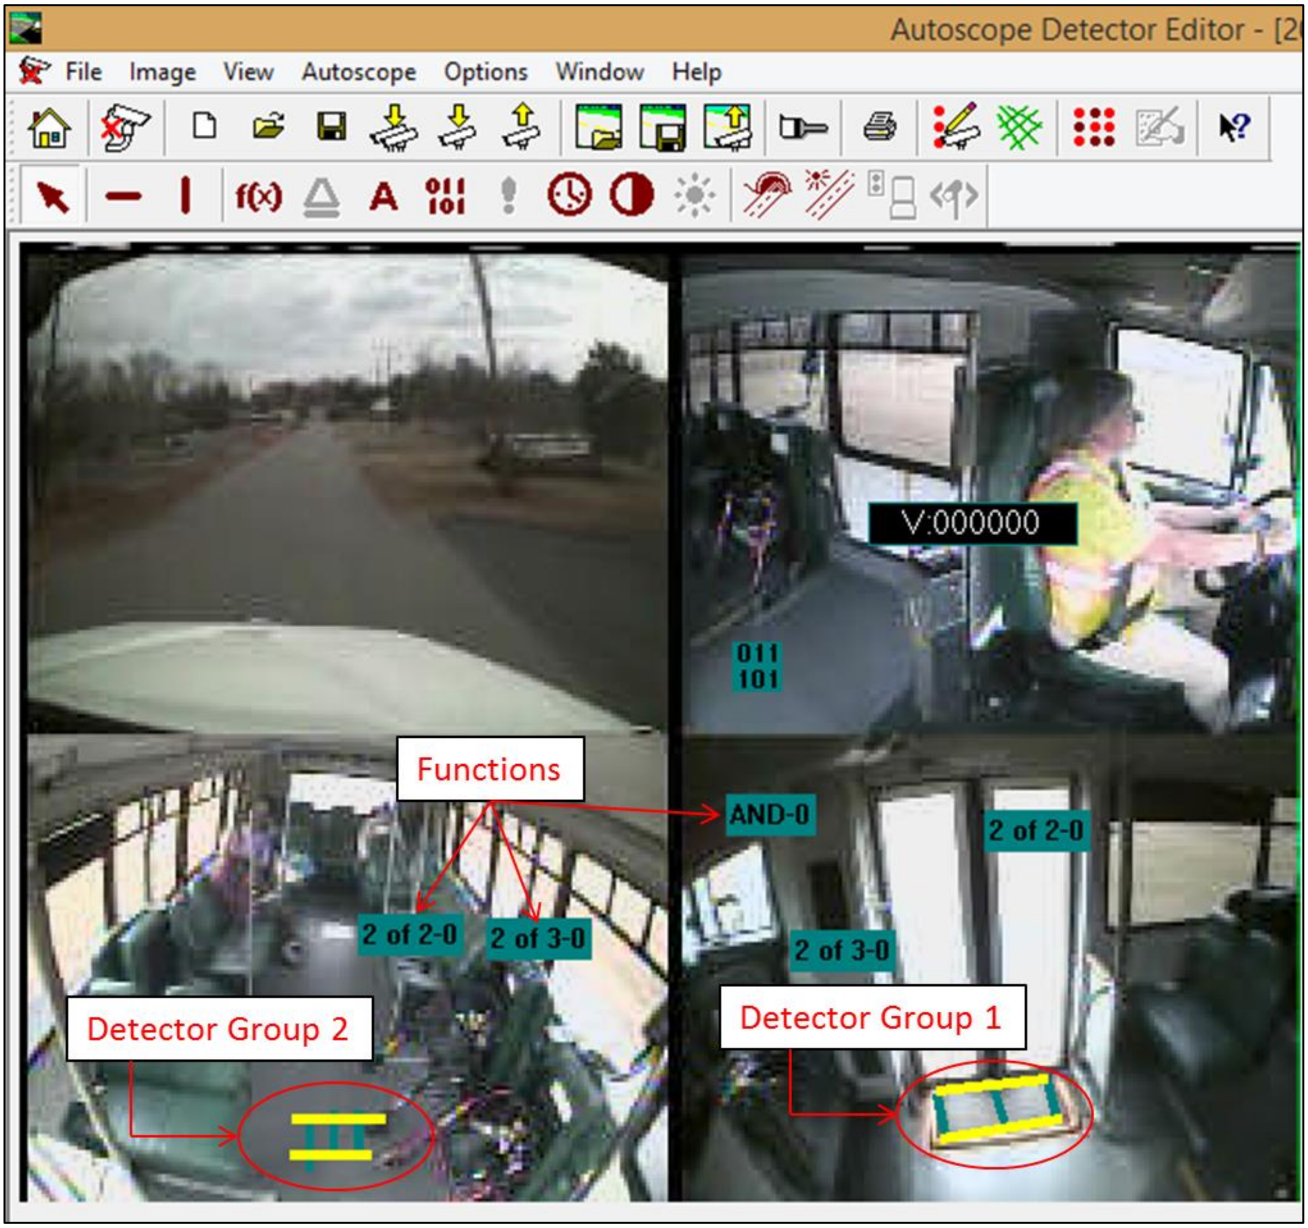

Figure 4 - Algorithm Development

3.3. Testing

Once the algorithm was developed for a vehicle type and video, it was tested by playing the video in real-time and processing it through the video detection package. The video data was automatically reduced by the software and polled to determine the number of passengers entering the bus. After processing a video through the software, manual counts determined the actual amount of entering passengers for comparison. These numbers were compared and the results are discussed below.

4. RESULTS

The algorithm consistently over-counted the number of passengers entering each vehicle, resulting in a high number of false positives. The table below displays the results of the comparison between the algorithm counts and the actual counts.

Table 1 - Actual vs. Algorithm Counts

	Actual	Algorithm	Difference (%)
WPRTA vehicles	5	14	280%
WAVE vehicles	11	55	500%
JCATS vehicles	17	99	582%

The percent differences between these counts are high and, more troubling, inconsistent. However, many of these high counts are due to a select few videos being processed very poorly through the detection algorithm. When omitting the outliers, the results are as follows:

Table 2 - Actual vs. Algorithm Counts (Outliers Removed)

	Actual	Algorithm	Difference (%)
WPRTA vehicles	5	14	280%
WAVE vehicles	1	3	300%
JCATS vehicles	8	16	200%

While these results are still not favorable, it appears that as the sample size of passengers increases, the accuracy improves, suggesting that videos with very low boarding rates may not be ideal for this method of collecting passenger boarding data, while videos that have frequent activity may be better suited for this type of automatic passenger counting. More importantly, though, the error in the algorithms is beginning to show a pattern of overestimating trips by factors of 2 to 3. If the error is consistent and the error range is acceptable, then adjustments can be made after the fact.

The outlier videos primarily consist of high-top vans, suggesting that this method will work better for larger transit vehicles, as described later in this section. Also, it's worth noting that the videos omitted from the table above had significant issues with shadows activating the detectors.

4.1. Issues

Many scenarios for detector placement were attempted. This kind of detection works best with a defined path and single riders getting on the bus at one time. The main issues are 1) when the vehicle layout has multiple paths a person could take, 2) the driver being detected multiple times while securing a mobility device, and 3) multiple riders getting on in quick succession. Some frequently observed situations where video detection appears to be effective or ineffective include:

1. *Light Transit Vehicles with a central aisle where the entry is directly across from the driver*

Video detection can be effective when used in a situation in which there is a central aisle with all passenger seats behind the entry door because the passengers have a defined path upon which to travel. In order to get on the vehicle, passengers enter through the door, pass by the driver, and then walk down the central aisle toward the seats. Likewise, to get off the vehicle they walk down the aisle, pass by the driver, and then exit the bus. This kind of defined path is the best case.

2. *Light Transit Vehicles with two front seats forward of the door*

This type of vehicle has side boarding, but passengers may turn left or right when boarding. There is no defined path in this scenario because there are seats both forward and aft of the door. With this seating arrangement, passengers seem to prefer to sit in the seats forward of the door. The forward seats do not have a long enough camera view that can be used to successfully place detectors. Placing the detectors on the floor directly adjacent to the seats causes false positive when the passenger sits in the seat because of arms waving through the detector or bags being placed on top of it.

3. *High Top Vans*

The camera angles on high top vans presented problems. These smaller vehicles have cameras mounted high in the ceiling with short sites lines and narrow aisles. When a passenger boards the van, the person's head takes up the entire camera angle. Passengers sitting at the front of the vehicle fill the camera screen and often inadvertently trip the sensors. In addition, passengers tend to utilize the aisle to spread out, which also trips the sensors.

4. *Slow Boarding*

If the door was left open for longer than the background image refresh rate, the detector would normalize to the open door. That meant when the door eventually closed, there would be about a minute gap where shadows would cause issues because the program thinks the open door is normal.

5. *Multiple Boardings*

Because of the cameras point down at an angle, multiple customers entering the vehicle in close proximity to each other appear to be one passenger.

5. FINDINGS & CONCLUSIONS

This proof of concept research does show promise in using surveillance camera video paired with Autoscope technology in certain situations. The ideal situation for using this concept are when:

1. The same vehicle is being sampled multiple times
2. The vehicle is a light transit vehicle with a center aisle and no seats forward of the boarding door
3. The video clip lengths are long (greater than 2 hours)
4. Passengers do not require driver assistance
5. The videos are from mid-day or cloudy days to eliminate shadows
6. Passenger boardings are distributed throughout the video clip instead of many passenger boardings occurring at only a few stops

The research team's familiarity with community transportation service in North Carolina means that we understand that these ideal conditions are unlikely to occur with regularity.

Depending on the amount of video required for validation of the transit system boarding rates, it may be more effective to simply manually count the passengers using the existing security video footage. While this process is tedious, it would simply mean gathering the video from the local agencies and observing the video for passengers. This process could be conducted at faster than real time speeds (ex. 2x or 4x), resulting in efficient passenger counting even though the process is being conducted manually. This would also prevent the NCDOT from needing to purchase any specialized software or hardware. Ultimately, the method of using video detection to count passengers is most effective when long periods of data need to be collected for count validation (greater than one week). See the sampling section below for more discussion on this topic.

5.1. Sampling Methodology

Discussions with NCDOT's Public Transportation Division reveal that the true need for automatic passenger counting is for fixed and deviated fixed routes in rural and small urban areas and that something must be proposed to fill this information gap. The pure demand response trips are easily validated because trip reservations are made, drivers fill out manifests, and most of the trips involve billing a contract agency. The community transportation systems that operate fixed and deviated fixed routes do not have the volume of trips to warrant purchasing infrared beam Automatic Passenger Counter nor do they have the technical capacity to manage the data.

Querying the annual Operating Statistics reveals that there are 28 community transportation systems that operate Motor Bus service, which is almost always deviated fixed route. We present two sampling plans, one based on manual counts and one based on using the technology explored in this research. We believe that both are viable and cost-effective options.

5.1.1. Manual Sampling

The manual sampling plan represents a spot-check of data integrity. Manual sampling will result in a low error rate, but loses its cost effectiveness at high volumes.

It is recommended that NCDOT staff, while conducting site visits, randomly select a deviated fixed route vehicle for analysis, as well as the supporting trip logs/manifests showing the manual counts. Although the preferred time period would coincide with the Vehicle Utilization Data (VUD) collection period, this is not feasible for two reasons. First, there is a lag between when the VUD data are collected, analyzed, and distributed and during this time, the video data may be overwritten. Second, to overcome the issue with this lag, the transit systems could pick the videos to be uploaded. However, having the transit systems select the routes will not instill confidence that the transit systems are being accurate in reporting data for the routes whose data have not been uploaded.

If NC State undergraduate students count the trips manually from the surveillance footage, they can watch the recordings at two to four times playback speed. Two hours of data from twenty-eight transit systems sampled twice per year will result in two hundred and twenty-four hours of data per year and could likely be viewed in approximately 100 hours. At a fully-loaded rate of $20/hour, the cost for manual counting will be $2,000 per year.

5.1.2. Automated Sampling

The automated sampling plan is more robust than the manual plan to better meet the ideal situations for using this technology but will result in higher error rates, where the tool will tend to estimate 2x or 3x the amount of actual boardings.

Again, the vehicle selected for analysis should be randomly determined by NCDOT staff among a list of vehicles that provide deviated fixed route service. However, instead of a few hours of data, NCDOT should collect a full week’s worth on an external hard drive. NCDOT staff can then deliver the hard drive to ITRE staff for analysis. This volume of data, perhaps amounting to 40 hours of video, will make adjusting the algorithms for the individual vehicle setup more cost effective. We recommend conducting this analysis once per year.

For the automated plan, each NCDOT staff member responsible for collecting data will need an external hard drive which costs less than $100. We estimate that the time required for splicing the video together and adjusting the algorithms for each vehicle to be around 6 hours per vehicle. The cost for having an NCSU student perform this type of analysis will be around $3,400 for all 28 transit systems.

5.2. Manual Versus Automated Counts

If spot-checking the data is desired by NCDOT, we recommend that manual counts be used because 1) setting the video detection algorithms involves slight, but time-consuming changes even if the vehicle types and camera configurations are the same, 2) data conversion into DVD format takes time, and 3) the video detection results are not as accurate as manual counts. However, if NCDOT expects to sample the video more frequently and is willing to accept an increased tolerance for error, then the video detection method could become increasingly more cost effective.

6. REFERENCES

1. Boyle, D. 2008. TCRP Synthesis 77 Passenger Counting Systems. Transportation Cooperative Research Program. Washington, DC.
2. Thomas Kimpel, James Strathman, David Griffin, Steve Callas, and Richard Gerhart. Transportation Research Record: Journal of the Transportation Research Board 2003 1835:, 93-100.
3. Kotz, Andrew J. David B Kittelson, and William F Northrop. 2015. Transportation Research Record. Jounrla of the Transportation Research Board 2536. 37-43.
4. Northrop, William F., David Kittelson, and Aditya Salapaka. 2017. Development of a Mass-Based Automated Passenger Counter. Research Poster presented at the 2017 Transportation Research Board Annual Meeting. Transit IDEA J-04/IDEA 84.
5. Vaughan, C., Jagadish, C., Bharadwaj, S., Cunningham, C.M., Schroeder, B.J., Hummer, J.E., Findley, D., and N. Rouphail (in press). Long-Term Monitoring of Wrong-Way Maneuvers at Diverging Diamond Interchanges. Transportation Research Board.
6. Cunningham, C.M., Chae, K, Schroeder, B., Salamati, K., and C. Vaughan (in progress). Evolution of Yield Detection Algorithm to Aid Roundabout Crossings for Pedestrians Who are Blind, Journal of Visual Impairment and Blindness (JVIB)
7. A Guidebook for Using Automatic Passenger Counter Data for National Transit Database (NTD) Reporting. National Center for Transit Research at CUTR, University of South Florida, Tampa, FL. December 2010. http://www.nctr.usf.edu/pdf/77803.pdf
8. Kimpel, Thomas J. and James G. Stathman. Automatic Passenger County Evaluation: Implications for National Transit Database Reporting. Center for Urban Studies, Portland State University. April 2002. http://www.ltrc.lsu.edu/TRB_82/TRB2003-000089.pdf
9. Sallay, John A.; Automatic People Counting and Matching. Doctoral Disseration. Utah State University. 2009.
10. Seon website. http://www.seon.com/school-bus-safety/student-tracking. Accessed 10/12/15.