Development of a Sensor Platform for Roadway Mapping : Part A - Road Centerline and Asset Management

Final Report

Brian Davis
Max Donath

Intelligent Vehicles Laboratory
Department of Mechanical Engineering
University of Minnesota

CTS 14-09

1. Report No. CTS 14-09	2.	3. Recipients Accession No.
4. Title and Subtitle Development of a Sensor Platform for Roadway Mapping: Part A - Road Centerline and Asset Management		5. Report Date June 2014
		6.
7. Author(s) Brian Davis, Max Donath		8. Performing Organization Report No.
9. Performing Organization Name and Address Department of Mechanical Engineering University of Minnesota 111 Church Street SE Minneapolis, MN 55455		10. Project/Task/Work Unit No. CTS Project #2013031
		11. Contract (C) or Grant (G) No. (c) 99008 (wo) 74
12. Sponsoring Organization Name and Address Minnesota Department of Transportation 395 John Ireland Boulevard, MS 330 St. Paul, MN 55155		13. Type of Report and Period Covered Final Report
		14. Sponsoring Agency Code
15. Supplementary Notes http://www.cts.umn.edu/Publications/ResearchReports/		

16. Abstract (Limit: 250 words)

Collecting information about the roadway infrastructure is a task that DOTs at all governmental levels need to accomplish. One way to increase the operational efficiency of these efforts is to use a relatively inexpensive mobile data collection platform that acquires information that is general enough to serve multiple purposes. The design and evaluation of one such platform that costs roughly $40,000 is described. It primarily consists of a differential GPS receiver providing vehicle location, and a LIDAR scanner that generates geometric profiles of the area between the vehicle and just beyond the road's edge. The vehicle collects data along the road by driving it in both directions. The system post-processes the data to automate feature extraction. For roads with simple geometry such as two-lane, undivided highways, the road's centerline can be calculated by finding the midline between the vehicle's paths from each direction of travel. Algorithms process the LIDAR scans to automatically detect the presence of curbs and guardrails, which is then combined with location information to yield the position of these features in world coordinates. The centerline calculation was determined to be accurate to within 6 cm in areas where its use was applicable. Curbs and guardrails were generally detected with an accuracy of better than 10 cm. The results demonstrate that it is feasible to use a relatively inexpensive mobile data collection system to acquire road centerline and roadside features such as curbs and guardrails.

| 17. Document Analysis/Descriptors

Driver support systems, In vehicle sensors, Vehicle safety, Ran off road crashes, Crash analysis, Crash data, Digital maps, Asset management, Pavement management systems, Traffic surveillance, Monitoring, Laser radar, Mobile computing, Highway safety, Evaluation and assessment, Performance measurement, Needs assessment | 18. Availability Statement

No restrictions. Document available from:
National Technical Information Services,
Alexandria, Virginia 22312 |

19. Security Class (this report) Unclassified	20. Security Class (this page) Unclassified	21. No. of Pages 61	22. Price

Development of a Sensor Platform for Roadway Mapping:
Part A - Road Centerline and Asset Management

Prepared by:

Brian Davis
Max Donath

Intelligent Vehicles Laboratory
Department of Mechanical Engineering
University of Minnesota

June 2014

Acknowledgments

Table of Contents

List of Figures

List of Tables

Executive Summary

Collecting information about the roadway infrastructure is a task that DOTs at all governmental levels need to accomplish. This information can include maintaining base maps for network level analyses and crash reporting, tracking of road furniture for asset inventory, assessing pavement quality to inform maintenance decisions, highway performance monitoring, road safety assessment needed to generate highway safety improvement plans and more. Generally, organizational groups within a DOT only collect the information they need to meet their functional goals, thus data collection techniques are usually specialized and have limited scope. This can lead to multiple groups collecting different types of data along the same road.

One way to increase the operational efficiency of these efforts is to use mobile data collection techniques that acquire information that is general enough to serve multiple purposes. These systems can be designed for installation on vehicles that are already traveling the roads for other purposes in order to reduce operational costs.

This report details the design and evaluation of a mobile data collection platform that costs roughly $40,000, which consists primarily of a differential GPS receiver, a LIDAR scanner and a computer. The GPS receiver is augmented with a real-time differential correction network to provide vehicle position information. The LIDAR generates a geometric profile from the outside edge of the data collection vehicle to roughly 2 lane widths away (24 feet). The vehicle collects data along the road by driving it in both directions.

In addition to methods for collecting the data, the report also describes the algorithms used to post-process the collected data to reduce the need for costly manual extraction of road features. For roads with simple or well-understood geometry (i.e. two lane roads symmetric about their centerline), the road's centerline can readily be calculated. This is accomplished by considering the vehicle path data along each direction of travel and then calculating a midline between these paths.

The LIDAR data is processed by algorithms that automatically detect the presence of pre-defined shapes in the scans. The system currently can detect curbs and guardrails, but can be extended in the future to detect any shape that is visible and unique. The algorithms detect these shapes within the scan and compute their lateral position with respect to the vehicle. This is then combined with vehicle location data to determine the position of these features in world coordinates. The algorithms then further process this information to filter out outliers and cluster nearby points into single, continuous roadside features.

The system was tested in two areas for which high accuracy (<10 cm) maps were available. One was an urban, divided highway and the other was a relatively rural two lane undivided highway. The centerline calculation was determined to be accurate to within 6 cm in areas where its use was applicable. Curbs and guardrails were generally detected with an accuracy of better than 10 cm.

Future work seeks to augment this system with video image processing to detect lane markings and to further increase the capabilities and accuracy of the system. The results of this work demonstrate that it is feasible to use a relatively inexpensive mobile data collection system capable of determining the location of a road centerline for roads with straightforward geometries and of roadside features such as curbs and guardrails.

1   Introduction

Roadway departure crashes represent a major road safety issue on both the national and state level. Infrastructure-based countermeasures including pavement treatments (e.g., rumble strips) have had success, but not enough to significantly reduce Run-Off-Road (ROR) fatalities.

The original motivation for this project was to examine how pavement monitoring vans within MnDOT (and other state DOT's), which already log all the roads in Minnesota on a multi-year cycle, could with additional sensors, be used to also collect lane boundary map data. Such high accuracy lane boundary maps could be used to develop new types of lane departure warning systems that could be effective in northern tier states where lane markings are often worn or covered by snow.

A recent study by CH2M Hill [1] indicated that there are more severe crashes on the county road system than on the state's road system. This study specifically identified lane-departure crashes (run off road and head-on) and curve-related crashes along rural roadway segments as those crashes that need the most attention.

As we worked towards developing a system for mapping the roads specifically on the types of roads where ROR crashes are significant (rural roads), it became clear that there were many other benefits of such a system. As such, we stepped back and explored other applications of a system capable of mapping roadway infrastructure at accuracies of 0.1 meters. These will be described here. In this document, we lay out the technologies that are available for mobile mapping, which provide a context for the work described here. We then identified test roads that were mapped by outside contractors at a higher accuracy than our goal. We followed up by evaluating the proposed system against the higher accuracy data to demonstrate that it actually met the aforementioned specifications.

Collecting information about the roadway infrastructure is a task that DOTs at all governmental levels need to accomplish. This information can include maintaining base maps for network level analyses, tracking of road furniture for asset inventory, assessing pavement quality to inform maintenance decisions, and more. Generally, DOT groups collect only the information they need to perform their function. This means that the techniques used are specialized so that they are the most efficient way to collect a particular piece of roadway data, but usually are not readily applied to the collection of a different piece of data. For example, a team that surveys the location of a guardrail would not be well suited to also collect pavement quality information. This leads to multiple, separate data collection efforts on roads.

One way to increase the operational efficiency of these efforts is to use mobile data collection techniques that acquire information that is general enough to serve multiple purposes. Here, the term mobile data collection is being used as an aggregate of all techniques that use sensors to collect data from a vehicle traveling a road.

In addition to collecting data for multiple groups at once, mobile data collection systems may be installed on vehicles that are already traveling the roads for other purposes. Furthermore, existing mobile data collection vehicles may be augmented with additional sensors to acquire new

roadway information without requiring the added operational costs to fund additional drivers and operators. For example, a pavement monitoring/video logging van may be fitted with equipment that allows for the collection of guardrail locations for inventory purposes.

It is important to note that there is a tradeoff between the capabilities of a system and the resources required to fully utilize them. Systems capable of collecting massive amounts of highly accurate data can be very expensive to purchase and operate. Additionally, the cost associated with processing the data to extract the needed information in a useful format can also be very costly. In evaluating mobile data collection systems, it is important to weigh the benefits of combining data collection efforts against the costs associated with doing so.

The work described in this report will discuss the design, implementation, and evaluation of a mobile data collection system that seeks to balance this tradeoff. The proposed system is much less expensive than high-end, data collection systems but still sophisticated enough to acquire multiple types of roadway information.

1.1   Prototypical Mobile Data Collection System

The term "mobile data collection" as used here, refers to any data collection technique that uses sensors to collect data from a vehicle traveling a road. This term is used as a generalization of the more common term "mobile LIDAR" or "mobile LIDAR scanning". The reason for this deviation is to include systems that do not use LIDAR. This section discusses the technologies generally used in a mobile LIDAR scanning system noting that a mobile data collection system may include only a subset of these technologies.

LIDAR, or light detection and ranging, is a method for determining the range between an object and the scanner. LIDAR scanners in general contain a spinning internal element that allows for multiple ranges to be measured in quick succession. The result is a scan, or a collection of distances reported with respect to the scanner. If the position of the scanner is known through the use of a positioning system, then the individual points it has measured can then be represented as an x,y,z position in world coordinates. When multiple scans are combined, this creates a point cloud, which is a collection of individual points representing the physical geometry that the scanner has detected.

Mobile LIDAR scanning refers to LIDAR data collected from a moving terrestrial vehicle as opposed to collection methods using an airplane or a stationary surveying station. Mobile LIDAR scanning systems consist of at minimum, a positioning system and a LIDAR scanner. The LIDAR provides information about the physical geometry of the roadway and the immediately surrounding area. The positioning system, generally consisting of a GPS[1] receiver and an inertial measurement unit, provides the vehicle's position. This information is then combined to generate a point cloud, a 3D model, of the geometry local to the roadway where the data was collected. In addition to this hardware, these systems can also be augmented by cameras, to provide for video logging needs. This raw data is then processed to extract the

---

[1] It would be more correct to say GNSS (Global Navigation Satellite System) which is the general term for satellite positioning systems, an example of which would be the American GPS (Global Positioning System). Most high end GNSS receivers including the one used in this project also receives signals from the Russian GLONASS.

required roadway information in a useful format. Further information is included in the overview of a prototypical mobile LIDAR system available in an NCHRP report discussing considerations for LIDAR data collection [2].

1.2   Uses for Mobile LIDAR Data

As described above, the raw data that a mobile LIDAR scanner outputs is a point cloud requiring further processing to yield usable information. Choosing what and how to extract the required data from the point cloud is determined by the application for the resulting information as well as the resources available to perform the road feature extraction.

This section describes a number of different applications for mobile LIDAR data. Each of these require different information be extracted from the raw data and depending on the requirements of the application, may necessitate different types or configurations of mobile LIDAR scanning equipment.

1.2.1   GPS-based driver assist systems

GPS-based driver assist systems seek to assist a driver by providing feedback about whether their vehicle is properly positioned within its lane. Therefore, these systems' efficacy depends on their ability to accurately and quickly determine the vehicle's side to side (lateral) lane position within the lane.

The most ubiquitous method for doing so is vision-based where one or more forward facing cameras are used to see the lane markings and determine the vehicle lane position. One existing product currently available is the Mobileye [3]. In addition to offering an aftermarket product, they also work with many auto manufactures to create OEM systems for integration with the vehicles' other safety sensors.

Vision based systems' efficacy is dependent on being able to detect visible lane markings. This can be an issue when lane markings have degraded or conditions like rain, fog, or snow block the view of the road. This is particularly problematic in climates like Minnesota's where snowfall is common and persistent, creating large portions of the year where such a system might be ineffective.

This motivates GPS-based driver assist systems that instead of sensing the lane markings in real time instead rely on a priori information about the road geometry and then correlate that with precise position information from a GPS receiver. The a priori information about the road geometry is stored in the form of a map that provides lane and road boundaries. This data is generally collected using mobile data collection methods. The resulting data can then be processed to yield the required road and lane features.

1.2.2   GPS-based lane departure warning systems (LDWS)

Lane departure crashes, also referred to as run-off-road (ROR) crashes represent a major road safety issue on our nation's roadway system. This has motivated infrastructure-based warnings such as signs to warn drivers of difficult road geometries and rumble strips to warn drivers when they enter the shoulder. However, even with these technologies, lane departure crashes are a common cause of roadway fatalities and injuries. In data for 2012 published by the Minnesota

Department of Transportation (MnDOT), 51% of fatal crashes were due to vehicles leaving their lane [4]. Additionally, a 2014 report prepared by CH2MHill lists run-of-road crashes among crash types representing the greatest opportunity for reduction [1].

Lane departure warning systems seek to alert the driver if or before their vehicle leaves the lane so that they can regain control and avoid a lane departure crash. These systems typically do not include any active vehicle control but rather provide some aural, haptic, or visual feedback to a disengaged driver to alert them to focus on the driving task because they are about to leave the lane in an unsafe way.

1.2.3   GPS-based Lane keeping

Lane keeping systems are also concerned with keeping the vehicle within the lane but for these systems, the goal is to assist trained, engaged drivers with a difficult driving task. Because the drivers are trained to use such a system, they can be more sophisticated, providing more information and even limited active control of the vehicle.

This type of system has been installed on some Alaska DOT and Minnesota DOT (MnDOT) snowplows [5]. In these applications, plowing is made more difficult by snow covering the lane markings or whiteout conditions obscuring the driver's vision. Here, the driving task can be very difficult even for an operator who is fully engaged. In this environment, the system provides assistance to ensure the driver's safety. The system consists of a heads up display (HUD) that superimposes virtual lane markings in the driver's field of view. The projection of the virtual lane markings in the HUD is based on a coordinate system positioned at the driver's eyes so that they line up with the true lane markings. The driver can thus continue to operate the vehicle even under total whiteout conditions. Additionally, if the vehicle is about to leave the lane, vibrations in the seat further warn the driver on the left or right as needed. Figure 1.1 shows the snowplow outfitted with the system as well as the driver looking through the HUD.

Figure 1.1: Instrumented snowplow and driver looking through HUD

Another implementation of this type of lane keeping system has been deployed in ten MVTA busses [6]. When transit busses travel on the interstate during rush hour, they are permitted to travel on the shoulder to bypass stopped or slowly moving traffic. However, this shoulder lane is much narrower than a general purpose lane and provides only a few inches of clearance on either side of the bus' mirrors. Maintaining precise lane position in these situations is not only very

difficult, but also critical to avoid lane departure that could cause a collision. This system expands on the snowplow system by including a motor that gently applies torque to the steering column and helps steer the bus. This allows the driver to be more confident in their ability to use the shoulders during conditions when they may not otherwise do so. Figure 1.2 shows the exterior of the bus fitted with the equipment and Figure 1.3 shows the driver looking through the HUD.

Figure 1.2: Exterior of instrumented MVTA bus

Figure 1.3: Driver looking through HUD

1.2.4 Inventory

In addition to the safety systems described above, mobile data collection can also assist the DOT by making existing operational tasks more efficient and cost effective. One application that can be improved is the process by which road assets are inventoried. It is necessary to keep track of the existence and location of road assets so that the groups responsible for their maintenance have a framework to keep track of their condition and estimate the resources needed for their upkeep.

Some methods for asset inventory can be time consuming, requiring workers to physically inspect and use inefficient surveying methods to locate road features. By utilizing mobile LIDAR scanning, it is possible to detect the locations of certain features by processing the resulting point cloud. From a single data collection effort, a single point cloud may be generated that could provide locations of guardrails, bridge abutments, bridge heights, ADA curb cuts, among other

features. Using a mobile data collection system that has a camera also may allow for extraction of road sign information. It is important to note that the smaller a feature of interest is, the more difficult it is to locate it. Highly accurate scanning systems can be very costly. Similarly, each additional feature to be extracted from the point cloud increases the resources required to do so.

1.2.5   Roadway Assessment
There are also mobile data collection systems that are used to determine the condition of the roadway itself. A mobile LIDAR scanner with extremely high accuracy or a specialized pavement assessment system may be used to generate information about the road surface's profile which can then be used to evaluate pavement quality and identify areas requiring maintenance. Road edges can also be evaluated to determine areas where unsafe edge conditions exist. Specialized sensors can also augment these systems to determine ride quality, provide greater detail about the road surface, or evaluate paint marking retro-reflectivity.

1.2.6   Construction Planning
Point clouds that are sufficiently accurate (~1 cm) can be used for planning road construction projects as a replacement for surveying. Here, the information provided by the LIDAR characterizes the existing roadway and its surrounding geometry so well, that planners can use the data to determine where the new road surface will be laid. Additionally, this enables new construction techniques such as stringless paving where computer controlled equipment lays the new road guided by information generated from the point cloud. However, the quality of systems capable of generating point clouds this accurately is reflected in the cost and resources required to purchase and operate the equipment.

1.3   Existing Mobile Data Collection Systems
Mobile data collection systems as defined here represent a very broad collection of systems. This section will focus on mobile LIDAR scanning systems as well as less expensive systems containing only a subset of the sensors found on mobile LIDAR scanning systems. Specialized sensors such as those used to generate very fine road profile or lane marking retro-reflectivity data are outside the scope of this analysis. Although they are useful for their particular task, they lack the flexibility and wide applicability that enables operational efficiencies when collecting many different road features from a single data collection effort. It is important to note that for applications where this data is required, systems can be augmented with these sensors. Alternatively, vehicles that already have these specialized sensors may be augmented with LIDAR and GPS to increase the capability of the platform.

1.3.1   Commercially available mobile LIDAR scanners
This report will highlight just two commercially available mobile LIDAR scanners representing a range of available commercial products. One is a less expensive, entry level system and the other is a more expensive but better performing system. The specifications and pricing information given for these products are based on a combination of published specifications and interviews with company representatives. The first is the Trimble MX2 [7], which comes in a single head or dual head configuration. The dual head configuration is shown in Figure 1.4.

Figure 1.4: Trimble MX2 System (from: [7])

This corresponds to whether the unit has a single or two LIDAR scanners, respectively. When it's in its less expensive, single head configuration, it is capable of collecting 36,000 points per second. When processed, this data can be used to create a point cloud that is accurate to 2 cm in the horizontal and 5 cm in the vertical. Including software to process and manage the data, the system starts at around $200,000. If the needs of the project are more demanding, another option would be to use the Riegl VMX-250 [8] shown in Figure 1.5.

Figure 1.5: Riegl VMX-250 System

It is a dual head scanner, which allows it to collect 600,000 points per second. The added accuracy of this hardware allows for a point cloud that is even more accurate at 1 cm horizontal and 2 cm vertical. However, to achieve this performance, this system's price can exceed one million dollars.

It is important to note that these are only two examples of commercial mobile LIDAR collection systems and there are many available systems with varying prices and performances. Research conducted at UC Davis in 2011 describes a number of these systems and how they fit into a data collection and usage plan. The UC Davis survey published by the Washington DOT discusses a handful of common, commercially available mobile LIDAR scanning systems including the

Riegl VMX-250, Trimble MX8, Topcon IP-S2, Ambercore TITAN, Optech Lynx, and StreetMaper 360 (Appendix A of [9]). A separate report by UC Davis published by the California DOT also discusses some commercially available systems but also analyzes point cloud accuracy for the StreetMapper, TITAN, and Lynx (Chapter 3 of [10]).

In addition to the scanning hardware itself, there are other costs associated with the operation of these systems. For example, there is the analysis that must be done before and after data collection to generate useful data at the accuracies discussed above. First, there must be one or more local GPS corrections stations to provide data for post-processing the raw observables collected by the data collection vehicle. Additionally, surveyed control points must be used to ensure accuracy, particularly in the vertical direction (perpendicular to the ground), which is more challenging for GPS to determine. These control points are used in post processing as ground truths as they are known positions that are also observable in the collected point cloud data. The generated point cloud has high relative accuracy, which is to say that the relative distance between features within the data are highly accurate; the control points allow for a better correlation between the relative distances between features and their absolute position within the world's coordinate frame.

Another consideration is the manpower required to operate and maintain these systems. Most of these systems require both a driver as well as an operator to run the collection software and monitor the process in real time. After the data is collected, the data must be processed to turn the raw point cloud data in to map data that can be used by a lane departure warning or other driver assist systems. This involves examining the point cloud data and locating the position of features to be represented in the resulting map. Although some tools exist to help automate some of the process, there is still considerable work that must be performed by a human to ensure the accuracy of the resulting map. These additional considerations require extra steps beyond simply driving the roads which in turn add to the time required and costs for these projects.

The advantage of this tier of mobile LIDAR scanning is the data produced has multiple applications. By generating data accurate on the order of 1 cm, it is useful not only for mapping, but also for pavement assessment, asset tracking, bridge heights, and construction planning. As a result of this high fidelity, it is also possible to further break down roadway information into multiple types of features for a single lane. For example such a system might be able to separately identify paint markings, the road-shoulder interface, curbs and guardrails as well as the edge of the shoulder. There are many advantages in using these types of systems however, they come with high costs. For this reason, it is critical to examine the requirements of the project to see if the extra accuracy afforded by such a system is necessary.

1.3.2 *Mapping systems using only GPS*
Another approach to creating maps to enable lane departure warning contrasts with the solution provided by commercially available scanning hardware. This technique focuses on the utilization of a lean system capable of delivering only the minimum final product, but in return is relatively inexpensive. This method again involves physically driving the roads to be mapped, but uses only a differential GPS receiver utilizing real-time corrections data such as that provided by Trimble's Virtual Reference Station (VRS) Continuously Operating Reference Station (CORS) network software [11, 12]. This mapping system logs the path of the vehicle as it travels a road

in both directions. The midline is then calculated between the two vehicle paths which are assumed to be the center of their respective lanes. This midline then represents the road centerline. Then, using a known or assumed lane width, the edges of the lane are then be calculated as a standard offset from the road centerline.

The obvious advantage of this method is that it is relatively easy to collect many miles of this data because there is no need for prior surveying to create control points. The hardware is also relatively inexpensive requiring only a GPS receiver with access to a corrections network and a computer to log the data. For example, data has been collected using this method on prior projects requiring only about $20,000 in equipment mounted on an existing vehicle. However, by collecting data using this approach, it is not possible to observe the position of, or even detect the presence of different road edge conditions. Additionally, the accuracy of the calculated lane boundaries is limited by the assumption that the lane width is a constant distance and that the vehicle is always centered within the lane.

This is the mapping technique for deploying driver assist systems used by the Intelligent Vehicles Lab of the University of Minnesota. It was used most recently in two snowplow projects - in the Donner Pass in California and at MnDOT's district 8 maintenance station at Windom, MN [2]. It is also used for mapping by the Public Works Department of Polk County, MN. In these projects, the accuracy of the maps created were not evaluated rigorously, but were found to be suitable for the application by ensuring that the systems operated as expected.

---

[2] This driver assist system technology has now been licensed to MTS Systems Corporation in Eden Prairie, MN.

2 Proposed System

The two approaches discussed above highlight the tradeoff between the cost of collection and the accuracy and road feature set that is generated. There is however, room in the middle for a mobile data collection technique that is less expensive than using commercial mobile LIDAR scanning products but can still provide useful feature data. Such a system would be particularly useful for situations where the extra information and data afforded by expensive collection methods was unneeded or would otherwise go unused. By using a suite of relatively inexpensive sensors, such a system can provide accurate road feature identification and positioning.

In addition to better matching the hardware capabilities to the data requirements, additional cost savings can be derived from automating the feature extraction process. That is to say, if the raw sensor data can be processed with little or no human input, the resources required to extract the necessary features from the raw data can be drastically reduced.

The focus of the system developed in this work is to generate high accuracy road centerlines for incorporating into state and county base maps. The system also seeks to detect selected road features such as the position of guardrails and curbs to facilitate inventory and GPS-based driver assist systems. The goal is to capture these road features accurate to 10 cm using a system whose components cost roughly $40,000. This accuracy level is considered to be a good tradeoff between being accurate enough to be used in base map, driver assist and inventory applications while still being cost effective to collect on a wide scale.

Such a system would collect a relatively sparse point cloud of data. Instead of those collected by commercially available mobile scanning systems, this system would collect data only in the area between the side of the vehicle to roughly 2 lane widths away. By restricting the data to only the areas of interest, this reduces system needs and processing power. This section describes the hardware used to accomplish this as well as the algorithms needed to generate useful information.

2.1 Hardware
2.1.1 GPS
The positioning system used in this project is a single GPS receiver that is augmented with real-time differential corrections provided by the statewide CORS network operated by MnDOT. The corrections data is sent over the internet using a wireless modem with a cellular data plan. Many other states operate a publically available network and in addition to these, there are a number of private VRS networks to which access can be purchased. A map of the Minnesota system is shown in Figure 2.1.
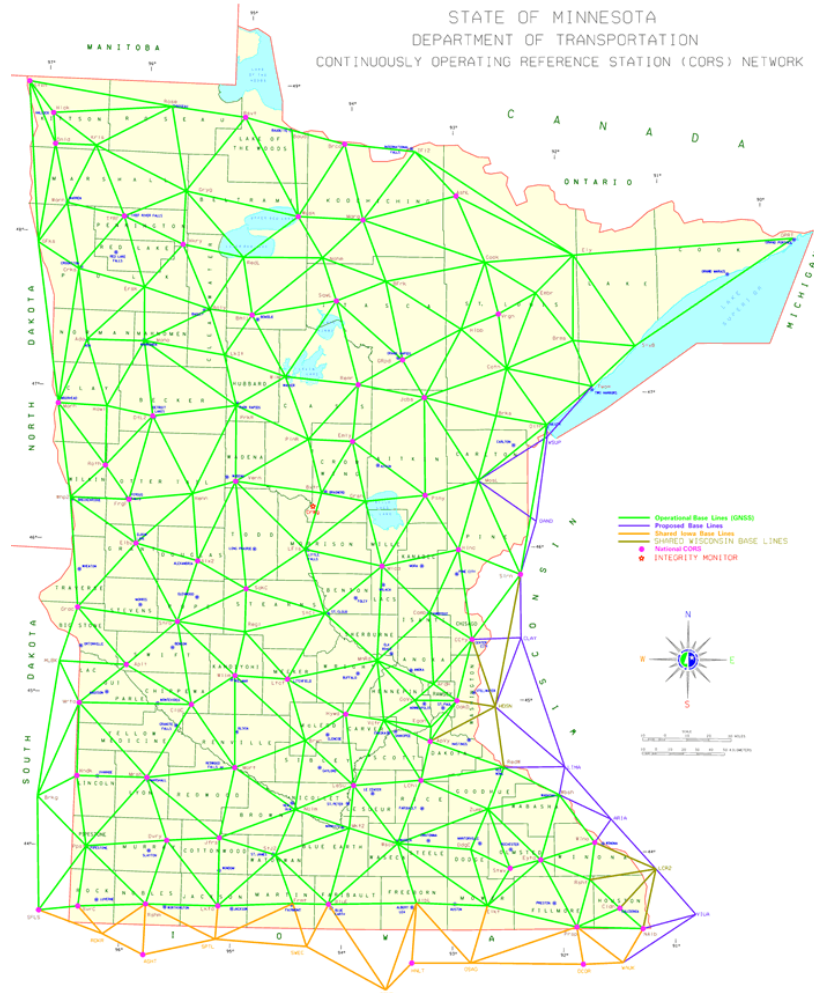
Figure 2.1: Minnesota CORS network as of March 11, 2014 (from: [12])

The positioning system used here differs from the prototypical commercially available mobile LIDAR scanning systems in a few ways. First, the GPS receiver is not augmented with either an inertial measurement unit (IMU) or a distance measurement indicator (DMI). Position information is only obtained from the GPS receiver and interpolated for sensor readings collected between position fixes. The data collected from the receiver is only the real-time position fixes as opposed to the raw observables, which could be used for post processing. This reduces the complexity of the post-processing which in turn reduces the resources required. Lastly, this system does not use control points, which are surveyed markers on the ground that show up in the point cloud and are used to link the point cloud to the real-world coordinates. Creating and surveying these control points presents a massive resource need for projects. By choosing to not utilize this technique, this drastically reduces the amount of work that must be completed before data collection. This in turn increases the efficiency of collecting data on a wide scale.

The GPS receiver used in this project is the Trimble R7 which is paired with the Trimble Zephyr Antenna. These are shown (not to scale) in Figure 2.2.

Figure 2.2: Trimble R7 receiver and Zephyr-2 antenna (from: [13,14])

It is a receiver capable of calculating position fixes with a stated horizontal accuracy of 8 mm + 0.5 ppm (ppm calculated with respect to the distance to the nearest correction station) RMS. The receiver is set to provide these fixes 10 times per second. For a data collection vehicle moving at 55 mph, this corresponds to a position update every 2.5 meters. This receiver is commercially available and costs about $18,000. In order to communicate with the VRS network, the receiver must be connected to a computer that has internet access. This is accomplished by using a mobile modem that provides internet access through the use of a cellular network.

2.1.2 LIDAR

Using a LIDAR scanner allows for the detection of road boundaries identifiable by geometric features. The scanner used in this system is mounted to the vehicle such that it scans in a plane perpendicular to both the road surface and the side of the vehicle. In this configuration, the scanner provides the geometric profile of the edge of the road and what is beyond it each time a scan is completed.

The scanner used in this project is the SICK LMS 511 shown in Figure 2.3.

Figure 2.3: SICK LMS 511 LIDAR (from: [15])

This unit costs about $10,000 with is significantly less than other, higher end scanners. It is capable of operating with an overall accuracy of about 1 cm, which is similar to the LIDAR scanners discussed above. However, beyond that, there are some major differences. Table 2.1 shows a summary of the key specifications for the scanner used in this project as well as two higher end units [15,16,17]. The first is the Velodyne HDL-64E. The second is the Riegl VQ-250 which is the LIDAR scanner used in the Riegl VMX-250 system described above. Both these LIDAR scanners are shown in Figure 2.4.
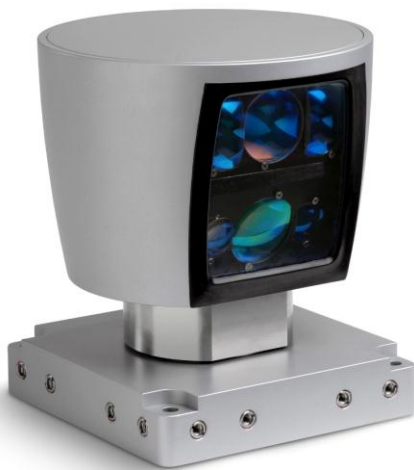
Velodyne HDL-64E (from: [16])   Riegl VQ-250 (from: [17])
Figure 2.4: LIDAR scanners researched for comparison

Table 2.1: Summary of select LIDAR scanners (specifications from manufacturers)

	SICK LMS 511	Velodyne HDL-64E S2	Riegl VQ-250
Scan frequency	25 - 100 Hz	5 - 15 Hz	100 Hz (max)
Simultaneous scans	1	64	1
Systematic error	25 mm		
Statistical error	7 mm		
Angular resolution	0.1667° - 1°	0.09° - 0.35°	.001° - 0.72°
Field of view	190°	360°(h), 26.8°(v)	360°
Weight	3.7 kg (8.2 lbs)	13 kg (29 lbs)	11 kg (24.3 lbs)
Accuracy	12 mm	20 mm	10 mm
Precision			5 mm
Beam divergence	4.7 mrad		0.35 mrad
Price	$10,000	$80,000	$200,000
Data Rate	28,525 pts/sec	1.333 mil. pts/sec	300,000 pts/sec

In examining the table, it is clear that these three scanners represent a range of both cost and raw data output. The LMS 511 is capable of taking 1141 range measurements over 190 degrees at 1/6 degree intervals and does so 25 times per second. This results in 28,525 individual points per second. On the other hand, the Velodyne unit takes 64 simultaneous scans at 0.09 degree intervals while it spins a complete 360 degrees five times per second. This allows it to collect 1.333 million points per second. The LMS 511 and HDL-64E S2 have adjustable scan frequencies which allow for more scans per second however by increasing this, they have to reduce their angular resolution. Configuring these specifications are a tradeoff between how much data is collected per scan and how frequently scans occur.

Increasing the data collection rate increases the density of the data collected. This means that the individual points collected by the scanner are closer together providing for a higher precision in the detection of smaller roadside features. However, these increases in sensor capability are also reflected in their price. The LMS 511 LIDAR used in this project costs much less than the higher end scanners.

The density of the collected points is not only affected by the update frequency of the scanner, but also by their location with respect to the scanner. For example, the Sick LMS 511 LIDAR's spinning element sends a pulse at a constant 1/6 degree interval. As the lateral distance from the car increases, so does the distance between consecutive scanner pulses. This is illustrated in Figure 2.5.
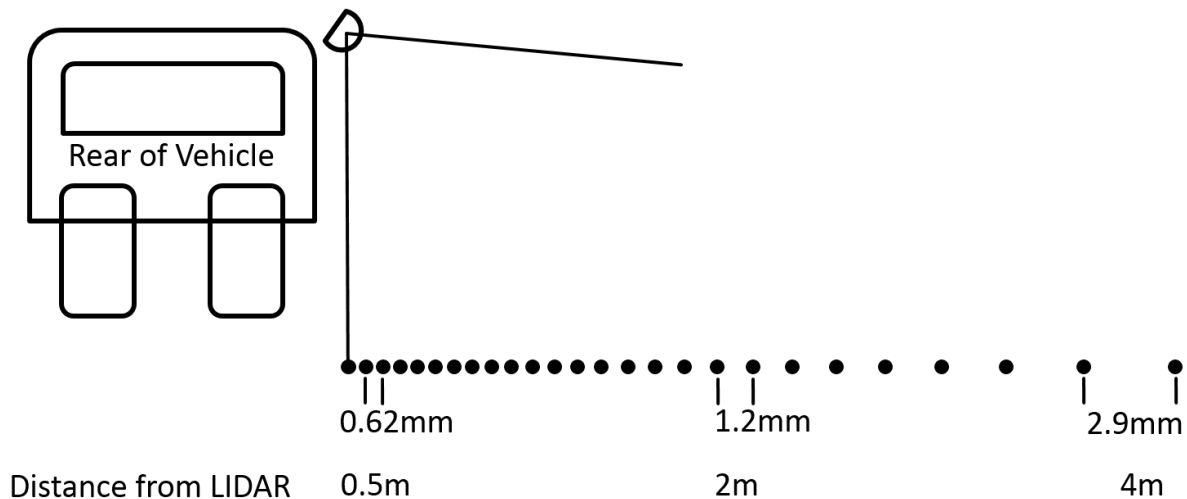
Figure 2.5: Illustration of pulse spacing vs. lateral distance from vehicle

Another factor that affects the point density is the collection vehicle's speed. For example, for a collection vehicle moving at 55 mph, each LIDAR scan is taken 1 meter down the road from the last one. This doesn't affect the accuracy of each individual scan, but it creates gaps in the data, which can cause ambiguity when trying to find the beginning or end of a roadside feature or otherwise trying to detect any feature that doesn't run alongside of the road.

Another consequence of the vehicle's speed is that the LMS 511 takes a finite amount of time to complete a scan. When operating at 25 Hz, it follows that a single scan takes 1/25 of a second. This means that for a vehicle traveling at 55 mph, in the time the LIDAR takes to scan 3.6 meters (approximately 1 lane width), the vehicle has moved 16 cm. Generally, there isn't enough lateral variation in the distance between the curb and the vehicle over this distance to have a large effect, but it can occasionally cause anomalies if in that 16 cm, a feature end or begins. For the purposes of this project, these issues are not formally addressed noting that there are other methods for determining the validity of and filtering the results.

2.1.3   Computer
A computer is required to communicate with the sensors in order to initialize and then acquire the data. In our case, the computer used is a robust industrial computer from VersaLogic that costs around $3000. It must be fast enough to run the sensor drivers and collect the data, but because no post-processing or analysis is performed on the vehicle, it can have relatively moderate specifications. Although the computer does not affect the accuracy of the system in any way, it is included in this discussion because it performs a critical function without which the system could not function.

2.1.4   Mounting Equipment
It is important that the equipment be attached securely to the data collection vehicle. The method utilized here is a two-part solution consisting of the vehicle mounts and the sensor pod. The vehicle mounts connect to the vehicle and provide a connection point for the sensor pod. The bars are mounted so that they are parallel to each other and the ground and perpendicular to the car's direction of travel.

Once the vehicle mounts are installed, the sensor pod attaches easily by sliding on the ends of the bars and clamping in place with set screws. The advantage of using a system like this is that the vehicle mounts can be left on the car even when the sensor pod is removed. The sensor pod itself is adjustable so that the orientation of the LIDAR can be changed allowing for its field of view to be focused in different directions. For this project, the focus was the space from the side of the vehicle to roughly two lane widths away.

The LIDAR's ability to be repositioned means that each time it is moved, it must be recalibrated. This is performed by moving the vehicle so that the corner between the ground and a wall is visible in the scan data. Then knowing that the ground is horizontal and the wall is vertical, the LIDAR's orientation and the height off the ground can be calculated.

For the primary research vehicle, a Chevrolet Impala, the mounts are a commercially available off the shelf roof rack that consist of two aluminum bars attached to the top of the vehicle with clamps that attach to the door frames. The vehicle mounts and sensor pod are shown attached to the Chevorlet Impala in Figure 2.6.

Figure 2.6: Vehicle mounts and sensor pod on Chevorlet Impala

To mount to the MnDOT Metro District's video logging van, custom vehicle mounts were fabricated because of the unique geometry of the top of the van. Here, vehicle mounts were constructed that clamped on to the rain gutter of the van which then support similar aluminum bars. This provides the same connection point for the sensor pod allowing for the same sensor pod to be used on either vehicle. The vehicle mounts for the van along with the sensor pod in this configuration are shown in Figure 2.7.

Figure 2.7: Vehicle mounts and sensor pod on Metro District's video logging van

2.2   Post-Processing Software

The methods by which raw data is processed into useful roadside features can be resource intensive. Some commercial tools exist to assist in extracting the desired features from relatively dense 3D point clouds. Trimble's Trident software and Certainty 3D's Topodot software are two examples of such software. Although they can aid with feature extraction, they do not fully automate the process. This has motivated our research that seeks to partially or fully automate the process of extracting features from point cloud data. Doing so could reduce the analysis time as well as enable the use of real time point cloud data in driver assist systems [18, 19]. The work done here to automate the process seeks to reduce the human time required as well as take advantage of the comparatively sparse data collected by the Sick LMS 511 LIDAR.

2.2.1   Centerline Calculation

The main purpose of the GPS receiver is to determine the location of the vehicle when individual LIDAR scans are taken. However, the vehicle's position data is useful on its own as well. Knowing the vehicle's path as it travels a road provides general information about its geometry. Depending on the a priori information known about a road or the assumptions made, vehicle path data can be used to calculate a road centerline.

17

For example, consider a two lane, undivided rural road where vehicle path data has been collected for one pass in each direction. The centerline of the road can then be calculated by determining the midpoints between the two vehicle paths. The validity of this calculation depends on two main factors. First, the vehicle must have been driven in the center of the lane for each pass. Second, the road must be symmetric about its centerline, which is to say that the lane width in each direction must be the same. It follows that deviations between the calculated centerline and the actual road centerline are due not only to sensor errors, but also are a function of variations in the vehicle's lateral position within the lane and the road's symmetry.

Because of these requirements, this method breaks down in areas where the road geometry is not symmetric such as urban areas with complex road geometries or divided highways characterized by a median or some other type of separation. For these types of roads, a centerline located in the median isn't always useful data.

The Highway Performance Monitoring System (HPMS) is a national level highway information system that includes data on the extent, condition, performance, use, and operating characteristics of the nation's highways [20]. The HPMS Field Guide [21] requires that centerline data for divided highways be reported such that there are two centerlines, one for the road surface running in each direction. This is to say that each direction of a divided highway is reported similar to separate one way roads where the centerline is the midline between the outermost through lanes.

To collect this data, multiple passes for a single road would then be required and for complex road geometries, it may not be obvious for a data collection driver which lanes are the outermost through lanes. For this reason, the centerline calculations described here are only considered valid for non-separated, symmetric roads.

The method used for calculating the road centerline is relatively simple. Vehicle path data is identified from two passes on a single road, each from opposing directions of travel. Then, for each individual vehicle position in one direction, the closest individual point from the other direction is identified. Then a centerline point is calculated as the midpoint on the line segment connecting the two points.

2.2.2   Shape finding

As the vehicle passes road features such as curbs and guardrails, these objects appear in the LIDAR data. Because of the relatively constant cross-section shape of these objects, it is possible to automatically find these shapes within the data. The algorithm implemented in this project extends the work done by the Intelligent Vehicles Laboratory related to operating a GPS-based bus driver assist systems in urban canyons. This previous effort [22] detected curbs and jersey barriers by using a LIDAR mounted in the same fashion as in this project. By detecting the distance between the known curb location and the bus, it was then possible to determine the lateral position of the bus within the lane when it could not otherwise be determined by the onboard GPS. Here, the algorithm works in reverse. The vehicle position is determined by GPS and the unknown feature location is then calculated.

A prototypical road feature model or template is defined for every shape to be located. The shape models that the algorithm currently can detect are shown in Figure 2.8. The black lines and curves represent the model shape which correspond to their respective feature's cross section. The vertical red line shows where on the feature their position is reported.
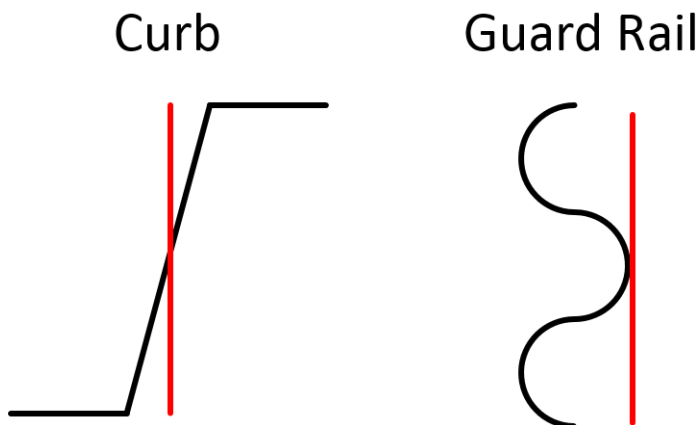
Figure 2.8: Illustration of prototypical curb and guardrail shapes

This is a relatively basic model that represents the profile of the feature to be detected. The model is then moved through the scan to determine how well the model fits the scan at each individual scan point. Then the location where the model matches the scan the best is determined to be the location of that particular feature. The similarity between the model and the scan is determined by calculating the mean square error between the two. This value is then thresholded to reduce occurrences of the algorithm reporting a found shape when it shouldn't.

Locating features within the scan only yields their lateral distance from the vehicle. This information must then be combined with the vehicle's position and heading at the time of the scan in order to place the feature within the world's coordinates. For each scan, the vehicle's position is determined by interpolating between the GPS fix immediately before and after the scan. It is assumed that in the 0.1 seconds between the fixes, the vehicle travels at a constant speed and along the segment connecting the two fixes. This is necessary because the GPS receiver does not provide heading or speed information.

2.2.3 Road feature clustering and filtering

The result of the shape finding algorithm is a collection of uncorrelated points each corresponding to a position where a particular shape was detected. However, it is necessary to further process these points in order to remove outliers and determine which points when combined represent a single continuous road feature.

The first step is to cluster the points into groups such that the generated clusters are all at least 2 meters from each other. To filter out outliers, any group containing less than 10 points is deleted. Next, nearby clusters are connected if the gap between them contains deleted points less than 4 meters apart.

3 System Evaluation Methods

3.1 Data Collection and Analysis

Our evaluation seeks to determine the accuracy with which the system can calculate a road centerline and determine the position of curbs and guardrails. This was conducted by first collecting LIDAR and GPS data in two test areas and then processing the data to produce centerline, curb, and guardrail location information. Then, the calculated curb and guardrail location information could be compared, where applicable, to reference data deemed to be accurate. Then, the offset between these two data sources can be calculated to better understand the systems efficacy.

3.2 Test Areas

Data was collected for evaluation in two test areas chosen due to the availability of high accuracy map data that could be used as a reference against which the collected data could be evaluated. The reference data is described in greater detail later in this section.

3.2.1 *Minnesota State Highway 65*

The first test area is a 3.8 mile stretch of MN 65 between I-494 in Fridley, MN and Anoka County Road 10 in Spring Lake Park, MN. This road is an urban divided highway with two through lanes in each direction. The extent of this test area is shown in Figure 3.1 superimposed on orthorectified aerial photography from MnDOT and map data from ESRI and their data partners.
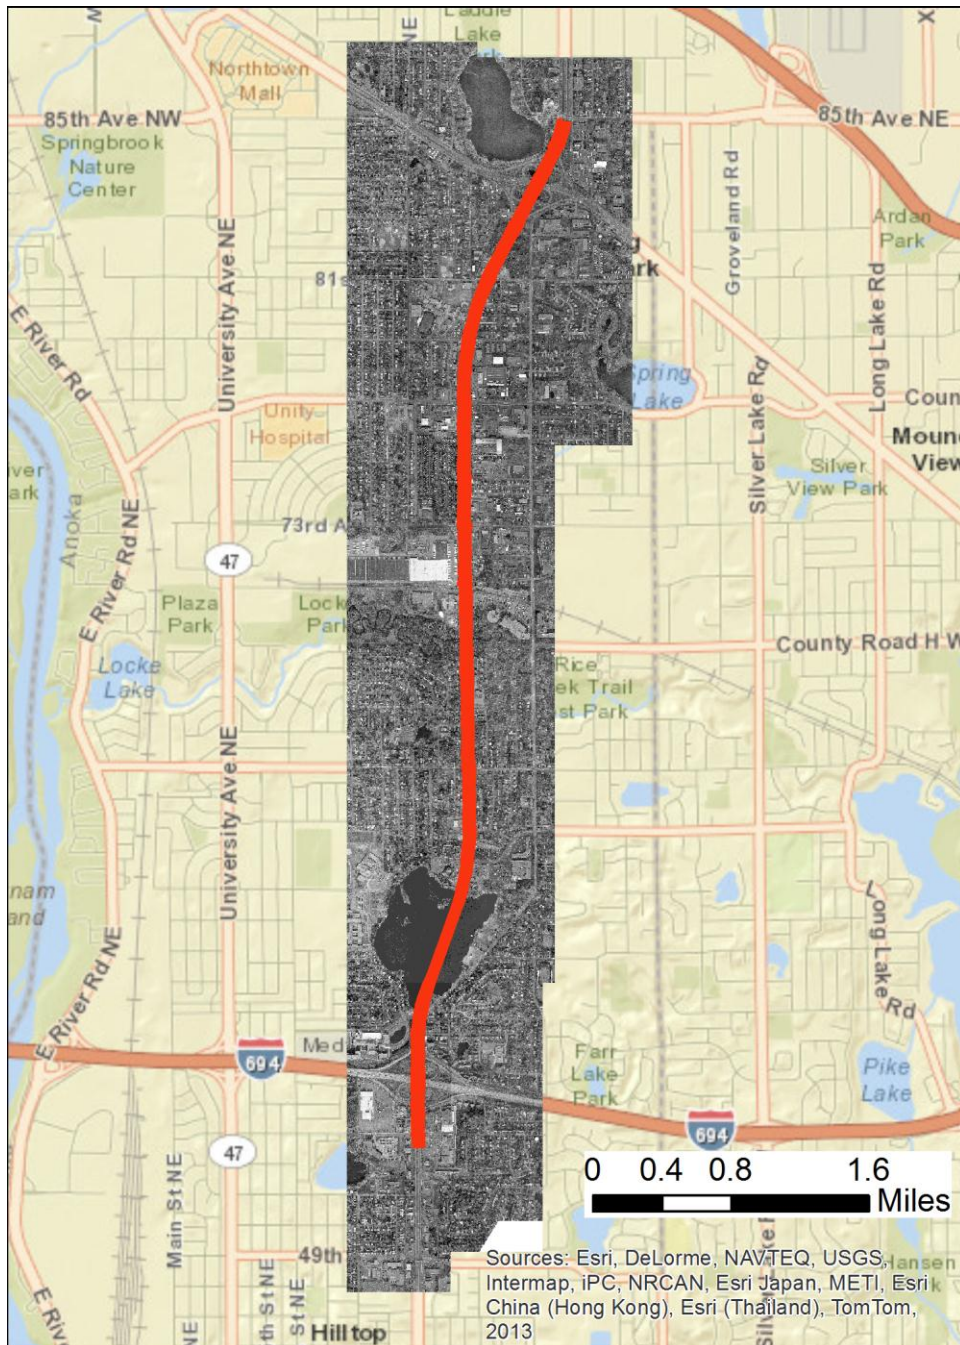

Figure 3.1: MN 65 test area data extent (Base map from ESRI)

3.2.2 United States Highway 63

The second test area is a 2.8 mile stretch of US 63 between US 61 in Red Wing, MN and Wisconsin State Highway 35 in Hagar City, WI. This road is an undivided two lane highway that includes two major bridges crossing the Mississippi River and Wisconsin Channel. The extent of this test area is shown in Figure 3.2 superimposed on imagery from Wisconsin DOT.

Figure 3.2: US 63 test area data extent

3.3 Reference Data

The reference data was provided by Continental Mapping, a contractor from Madison, Wisconsin. They drove the test areas with a Riegl VMX-250 mobile LIDAR scanning system in the summer of 2013 and then processed the resulting point cloud to extract lane and road boundary information delivering the data in March 2014. The raw point cloud data has a specified horizontal accuracy of 5.39 cm for US 63 and 8.50 cm for MN 65. The road features extracted from the point cloud have a stated accuracy of 10 cm or better.

The reference data contains, among other features, curbs and guardrails. These elements are each represented by a line string, which is a series of connected points. For curbs, the line string reports the location of the bottom of the face of the curb. For guardrails, the line string reports the location of the center of the guardrail as seen from above. Figure 3.3 and Figure 3.4 were provided by the contractor illustrating the reporting locations for these two shapes. Each figure shows an oblique view and a top down view. The solid red line marks where the feature's position is reported. A dashed red line is superimposed on the images to show approximately where the algorithm discussed in this report would report the location of these features.
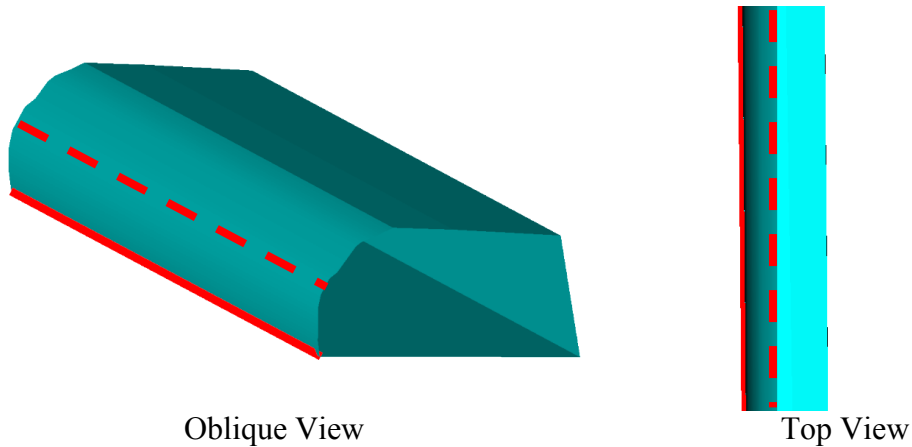

Oblique View   Top View

Figure 3.3: Reference data curb shape

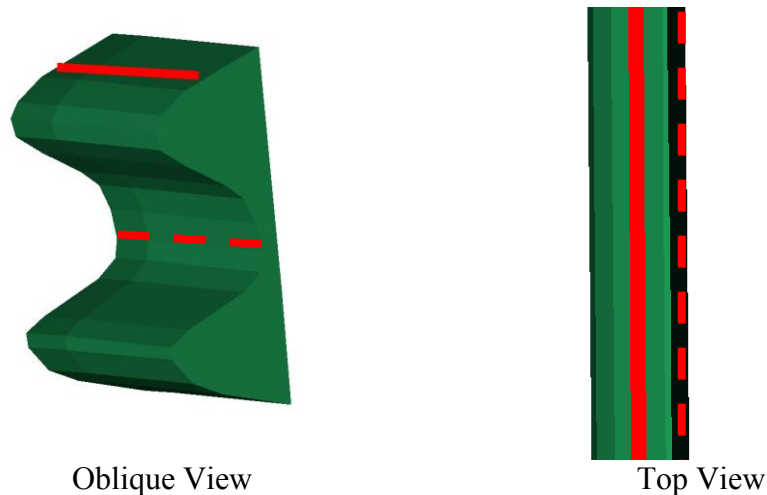Oblique View   Top View

Figure 3.4: Reference data guardrail shape

4   Results

4.1   Centerline Calculation

Road centerlines were calculated based on vehicle path data for both test areas. As described above, the accuracy of the calculations is dependent on the geometry of the road. For the entire MN 65 test area, the road is an urban divided multi-lane highway with occasional protected right and left turn lanes. This makes the centerline calculation described above inappropriate. Although it was possible to make the calculation, the road centerline was placed in the median. This meant that the centerline could not be compared to any available reference data. Furthermore for divided roads, a more meaningful centerline is re-defined so that there is one road centerline per direction. Because data was collected only in the right most lane for each pass, it was not possible to calculate such a centerline.

The US 63 test area consists mostly of relatively simple undivided highway. Figure 4.1 shows the extent of the data collected marked with four locations to be discussed below. These images contain lane markings from the reference data shown as red lines and the calculated centerlines as blue points. Again, the base imagery is orthorectified aerial photography from the Minnesota and Wisconsin DOTs.
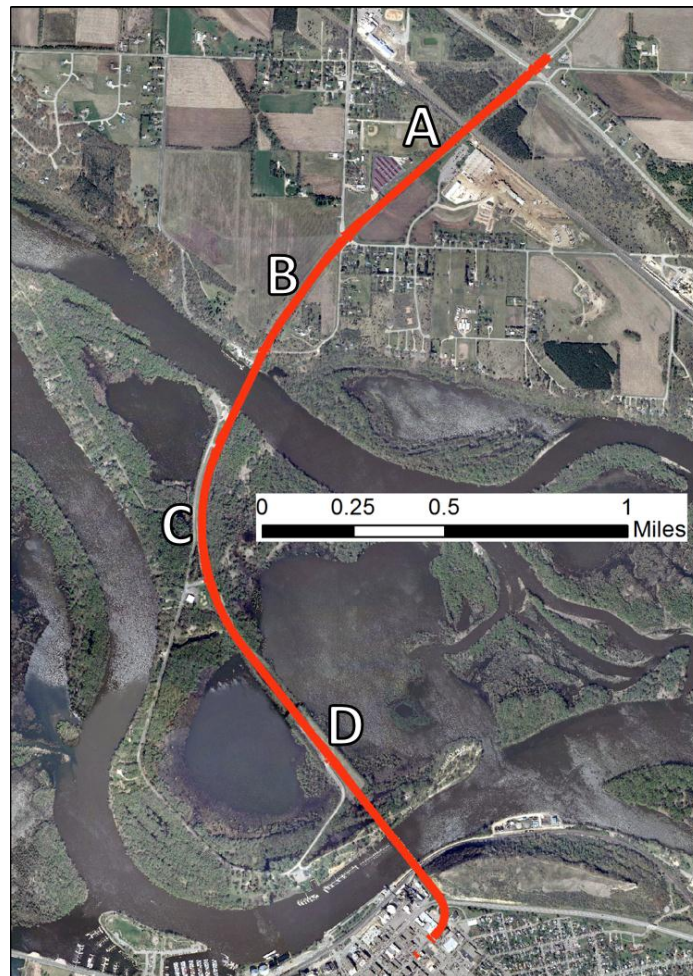

Figure 4.1: US 63 test area marked with centerline evaluation locations

Location A was one such area where the roadway geometry allowed for a meaningful centerline calculation. We used this 300 meter subset of the roadway to calculate statistics on the accuracy of the calculated centerline data.

These results represent data collected from 4 trips, a trip consisting of one pass in each direction. The accuracy of the calculation was determined by comparing the distance of the calculated centerline points to the reference data. The accuracy statistics are shown in Table 4.1.

Table 4.1: Summary of centerline calculation errors for location A (300 meter road segment)

	Trip 1	Trip 2	Trip 3	Trip 4	Total
# of samples	256	244	261	260	**1021**
Mean error [m]	0.041	0.054	0.078	0.051	**0.056**
Standard deviation [m]	0.025	0.048	0.045	0.028	**0.040**

Figure 4.2 and Figure 4.3 show this location with increasing magnification. Figure 4.4 shows a section in the middle of the road segment whose accuracy is representative of the rest of location A.

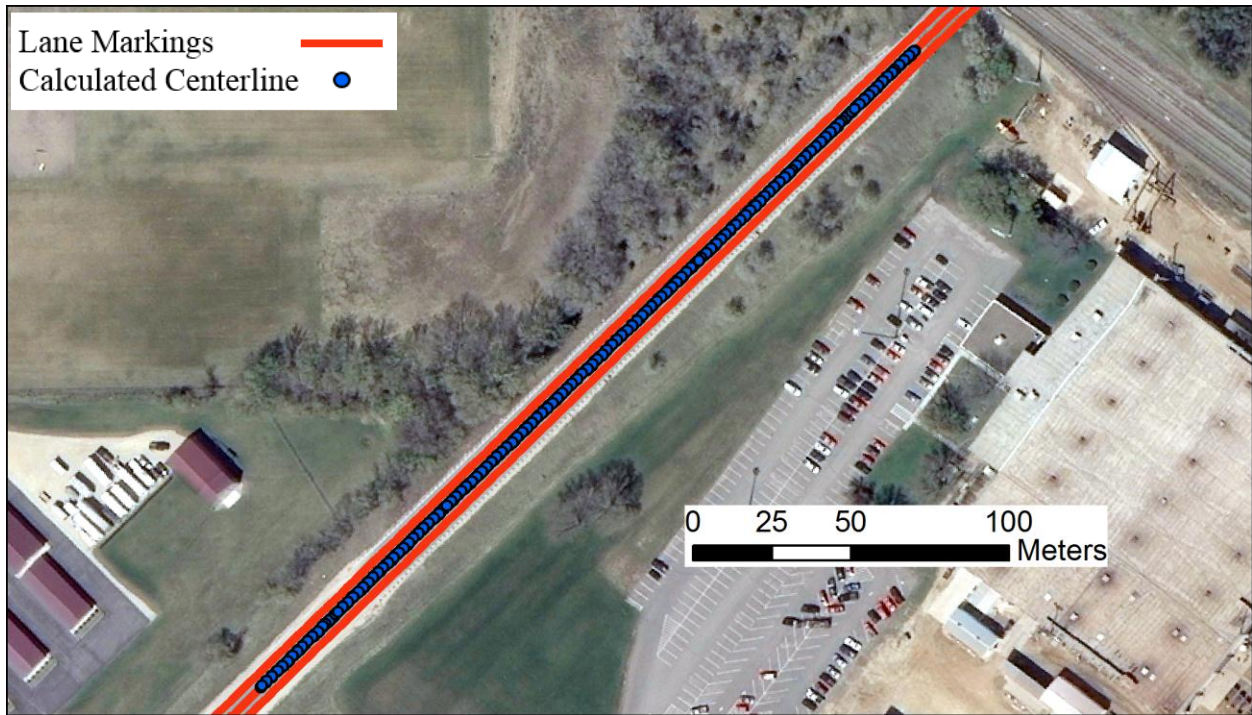Figure 4.2: Centerline (US 63) – Location A data extent
(Data from 4 trips)

Figure 4.3: Centerline (US 63) – Location A magnified
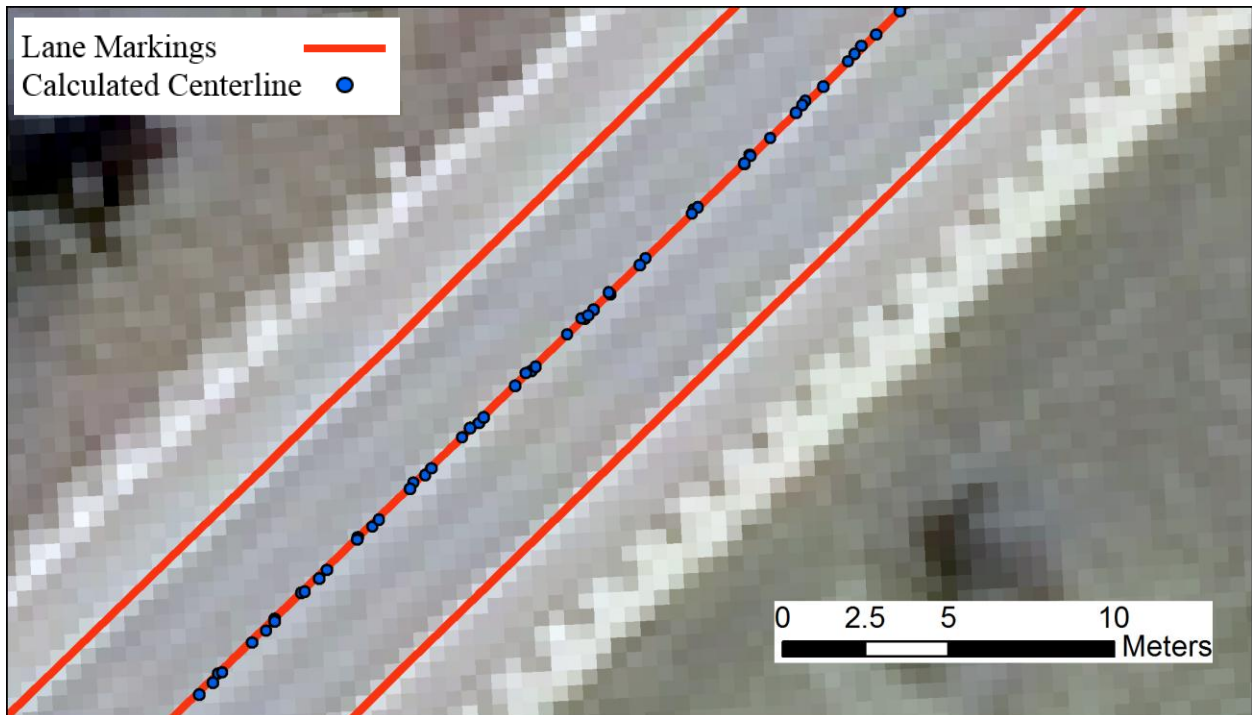(Data from 4 trips)

Figure 4.4: Centerline (US 63) – Location A magnified again
(Data from 4 trips)

Figure 4.5 shows location B and Figure 4.6 shows location C. These locations represent road segments where the roadway is still undivided, but whose centerline is characterized by a double yellow line. In these figures, the calculated centerline generally falls within the space between the double paint stripes.
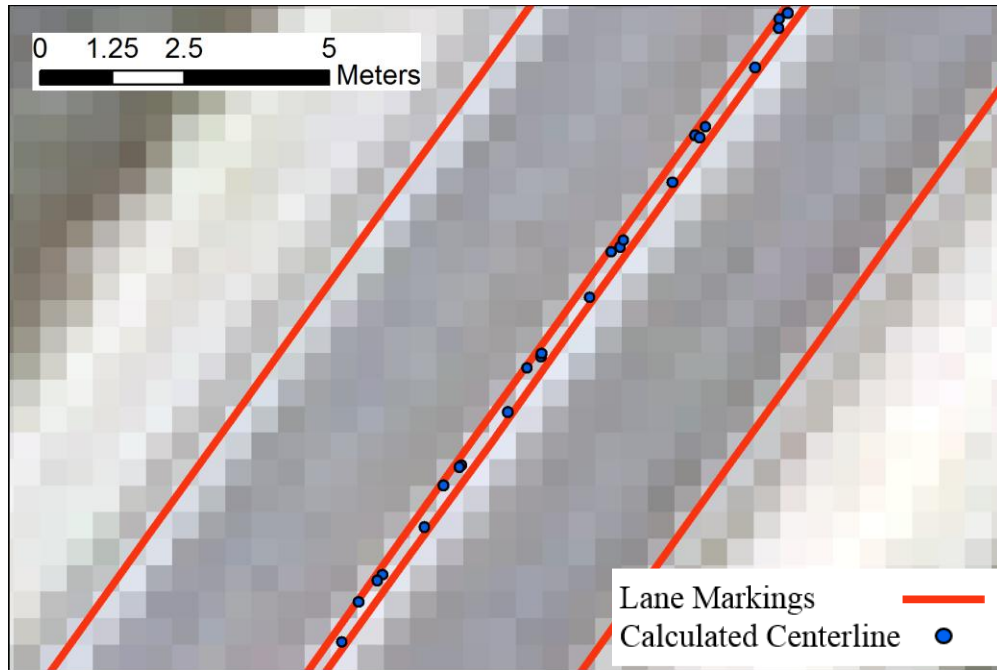

Figure 4.5: Centerline (US 63) – Location B
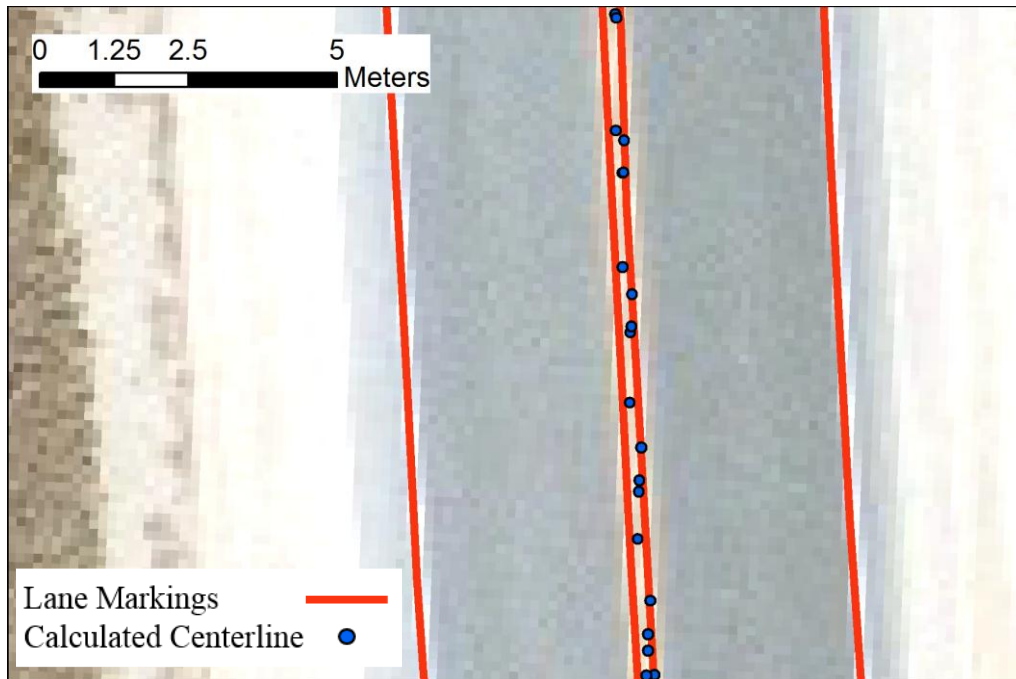(Data from 4 trips)

Figure 4.6: Centerline (US 63) – Location C
(Data from 4 trips)

Some areas of US 63 also contained complex geometry which caused the method described above to be inappropriate. An example of such an area is location D shown in Figure 4.7 and a magnified version in Figure 4.8.
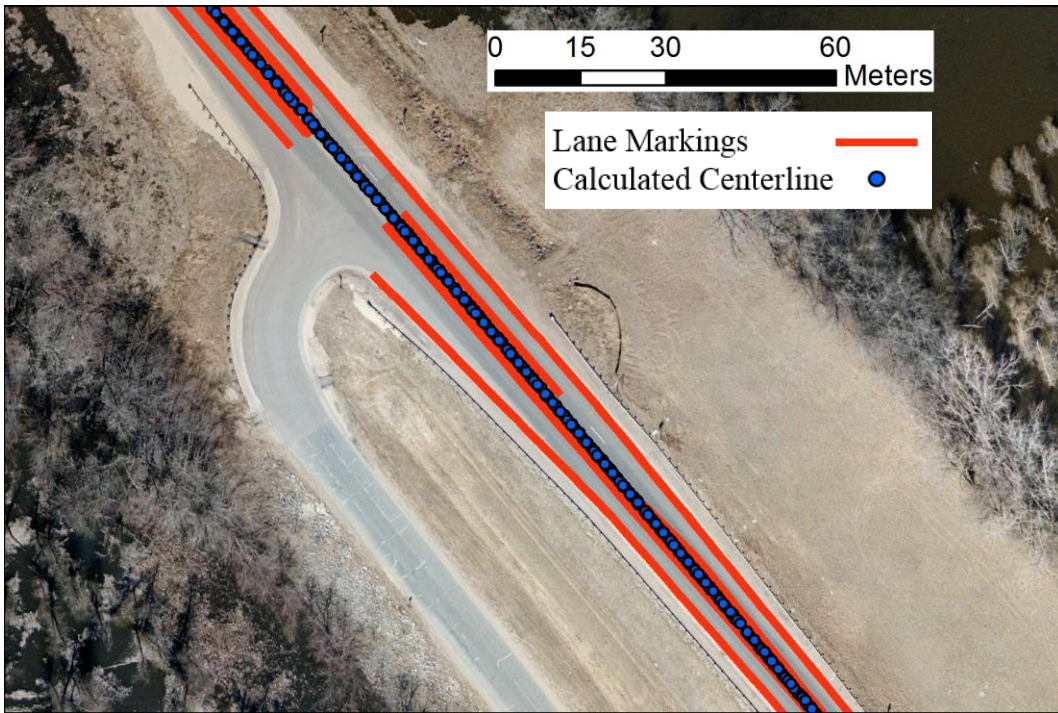

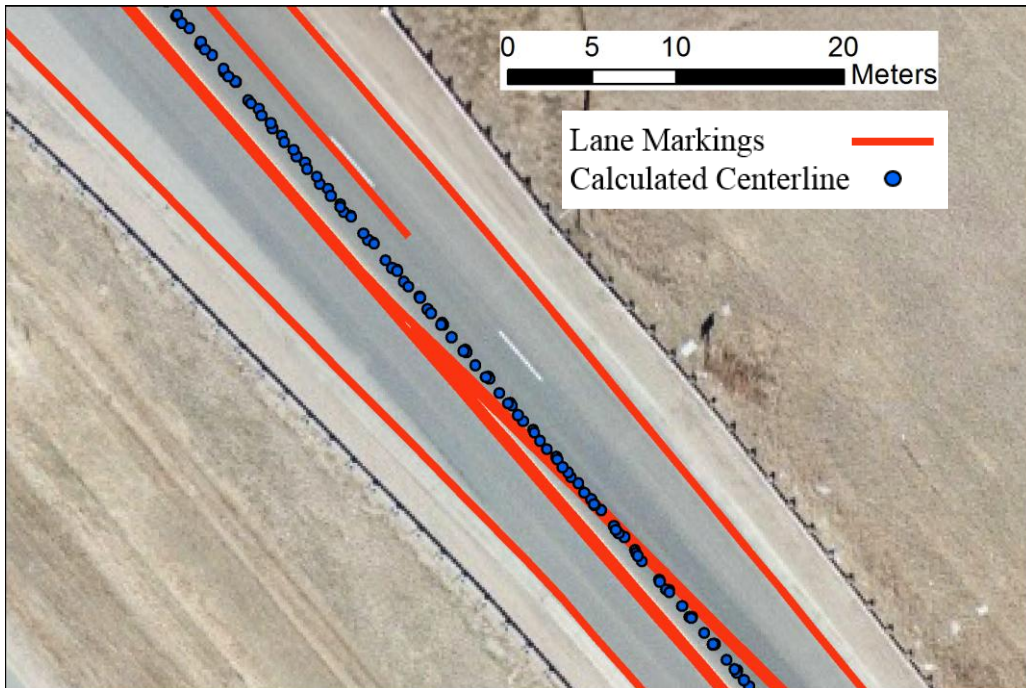Figure 4.7: Centerline (US 63) – Location D
(Data from 4 trips)

Figure 4.8: Centerline (US 63) – Location D magnified
(Data from 4 trips)

4.2   Quantitative Evaluation of Road Feature Detection

The quantitative results presented in this section consist of the means and standard deviations of the offsets between the collected roadside features (in our case, the curbs and guardrails) and the reference data. These calculations were done by projecting or transforming both the calculated and reference data into the Minnesota South State Plane FIPS 2203 coordinate system. These are aggregate measures calculated for each road feature and for each pass. A pass consists of travel along the entire length of the test area in a single direction.

The mean and standard deviation data is presented as the mean in meters followed by the standard deviation for that value in parentheses. Features are numbered as to match their designation in the reference data. These numbers were assigned by Continental Mapping and are arbitrary. Similar feature numbers do not imply similar feature characteristics or location. Feature numbers that are skipped in the tables represent features that were not detected by the system on any pass.

Occasionally, the system would detect multiple features that corresponded to a single roadside feature in the reference data. This situation is illustrated in Figure 4.9.

Feature from Reference Data
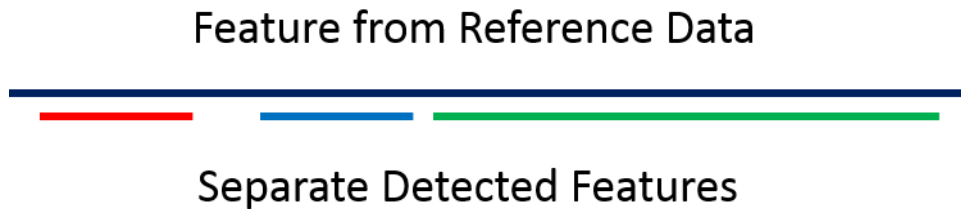
Separate Detected Features

Figure 4.9: Illustration of incorrect road feature segmentation

In this case the reported values are the combined mean and standard deviation and the entry is marked with a " * ". Omissions in the table represent a feature that was not collected for a given pass. Table 4.2 and Table 4.3 show guardrail and curb data from MN 65. Table 4.4 and Table 4.5 show guardrail and curb data from US 63.

Table 4.2: Guardrail Data for MN 65

Guardrail Number	Pass 1	Pass 2
4	0.054 (0.187)	0.135 (0.231) *
8	0.108 (0.612) *	0.044 (0.041)
9	0.063 (0.027) *	0.064 (0.031) *

Table 4.3: Curb Data for MN 65

Curb Number	Pass 1	Pass 2
2	0.068 (0.033) *	0.065 (0.034) *
10	0.310 (1.189) *	0.100 (0.253)
17	0.071 (0.015)	0.065 (0.011)
28	0.070 (0.084)	0.148 (0.290)
32	0.069 (0.026)	0.074 (0.102)
33	0.157 (0.222) *	0.046 (0.014) *
41		0.081 (0.308)
42	0.022 (0.016)	0.018 (0.016)
48	0.045 (0.014)	0.052 (0.015)
49	0.048 (0.027)	0.161 (0.388)

Table 4.4: Guardrail Data for US 63

Guardrail Number	Pass 1	Pass 2	Pass 3	Pass 4
0	0.057 (0.026)	0.058 (0.022) *	0.055 (0.025) *	0.054 (0.026)
1	0.035 (0.022) *	0.031 (0.020) *	0.035 (0.022)	0.031 (0.025) *
2	0.052 (0.019) *	0.050 (0.017) *	0.045 (0.019)	0.050 (0.028) *
3	0.060 (0.055)	0.059 (0.054)	0.058 (0.056) *	0.064 (0.055)
4	0.123 (0.009)	0.119 (0.011)	0.120 (0.011)	0.118 (0.010)
5	0.059 (0.036)	0.062 (0.038)	0.061 (0.038)	0.073 (0.113)
6	0.015 (0.009)	0.020 (0.007)	0.018 (0.010)	0.021 (0.011)
7	0.077 (0.320)	0.074 (0.030)	0.074 (0.031)	0.071 (0.029) *
8	0.030 (0.018) *	0.025 (0.014)	0.022 (0.012)	0.020 (0.012)
9	0.020 (0.009)	0.020 (0.011)	0.019 (0.008)	0.015 (0.010)
10	0.106 (0.064)	0.078 (0.014) *	0.078 (0.012)	0.074 (0.011)
11	0.065 (0.014)	0.065 (0.011)	0.060 (0.011)	0.060 (0.013)
12	0.025 (0.013) *	0.019 (0.012) *	0.020 (0.013) *	0.019 (0.010) *

Table 4.5: Curb Data for US 63

Curb Number	Pass 1	Pass 2	Pass 3	Pass 4
11	0.073 (0.015)	0.069 (0.015) *	0.069 (0.015) *	0.068 (0.016) *
12	0.052 (0.023) *	0.058 (0.124) *	0.039 (0.024) *	
25			0.022 (0.016)	
30	0.078 (0.026)	0.079 (0.026)	0.073 (0.026)	0.071 (0.031)
31	0.046 (0.079) *	0.080 (0.075) *	0.059 (0.078) *	
32		0.235 (0.185)		
35	0.096 (0.068)	0.081 (0.007)		
36	0.078 (0.037)		0.109 (0.012)	0.084 (0.300) *

These results show that in general, the system is capable of accurately determining the location of the features it detects. Most of the mean errors are around 5 to 8 cm which is within the 10 cm goal. It's important to note that these values do not give any information about whether the detected feature was closer to or further from the vehicle than the reference road feature. With

more information about which side of the reference data the features are detected, it may be possible to determine the systematic error for each feature to be used for system calibration which would further increase the accuracy. The data also shows that there are some features that aren't properly captured meaning that single features are detected as multiple features or they aren't collected at all. This leaves some room for improvement as discussed later in this report.

4.3   Qualitative Evaluation

The qualitative evaluation of the data consisted of mapping and comparing the road features collected by the system with those identified in the reference data. The images shown in this section consist of the reference data and the calculated results from all passes. These two data sources are superimposed on orthorectified aerial photography which is used for visual context. The accuracy of the imagery is limited to 12 to 18 inches per pixel. The reference data is represented as line strings colored light green for guardrails and orange for curbs. The collected data from the system described in this report is represented as points colored green for guardrails and red for curbs. The test areas are shown marked with letters that designate the location of the subsequent figures.

4.3.1   MN 65

Figure 4.10 shows the southern end of the MN 65 test area from which 4 locations are examined more closely. This image includes map data from ESRI.

Figure 4.10: MN 65 data extent with figure locations (Base map from ESRI)

31

Figure 4.11 shows an example of accurate guardrail detection at location A. There is a small discrepancy at the ends because guardrail end treatments do not have the same cross section as the rest of the guardrail. Because the sensors always faced the outside (right side) of the road, interior features (located on the left of the vehicle) are not detected. The reference data was collected with hardware that did allow for this collection and therefore is capable of capturing these features.
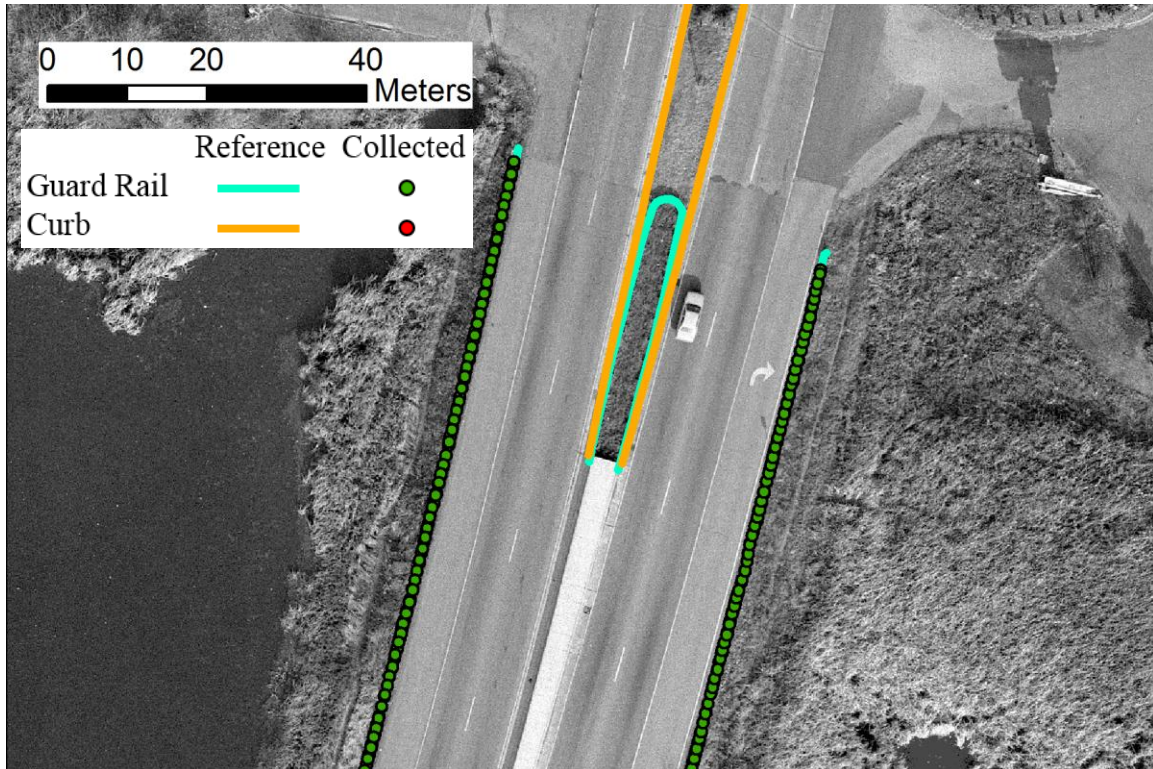

Figure 4.11: MN 65 – Location A
(Data from 2 passes in each direction)

Figure 4.12 shows location B. This location is a section of bridge where guardrails line the edge of the road. Figure 4.13 shows a zoomed in area of location B whose errors are representative of the entire location. The errors shown in this figure are generally 5 to 7 cm.
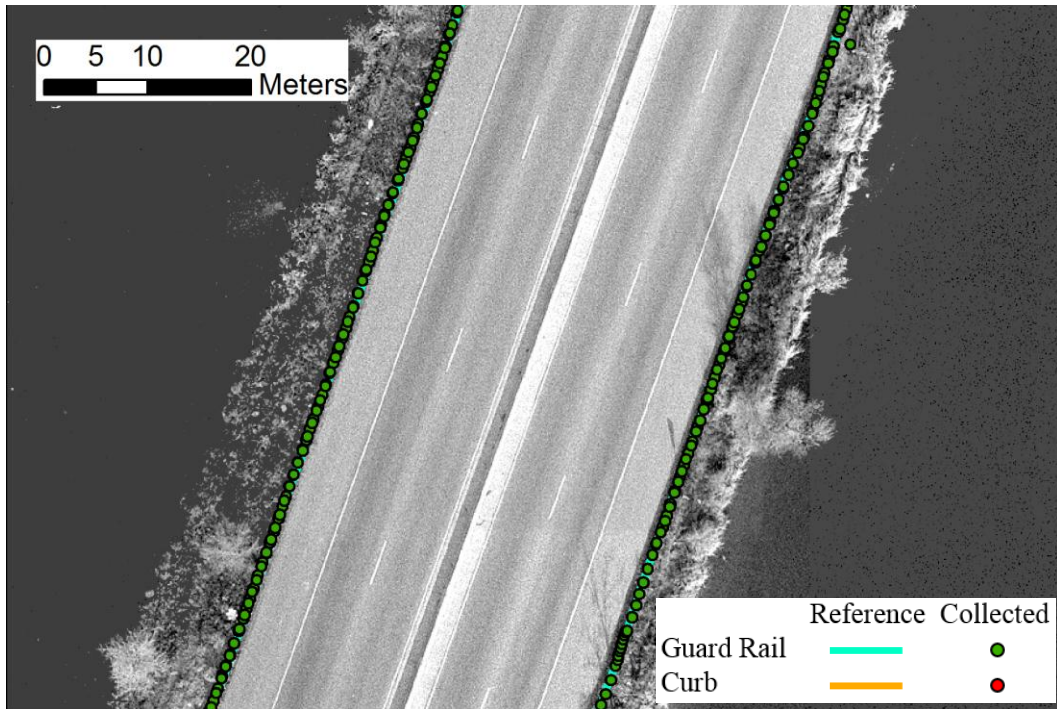

Figure 4.12: MN 65 – Location B
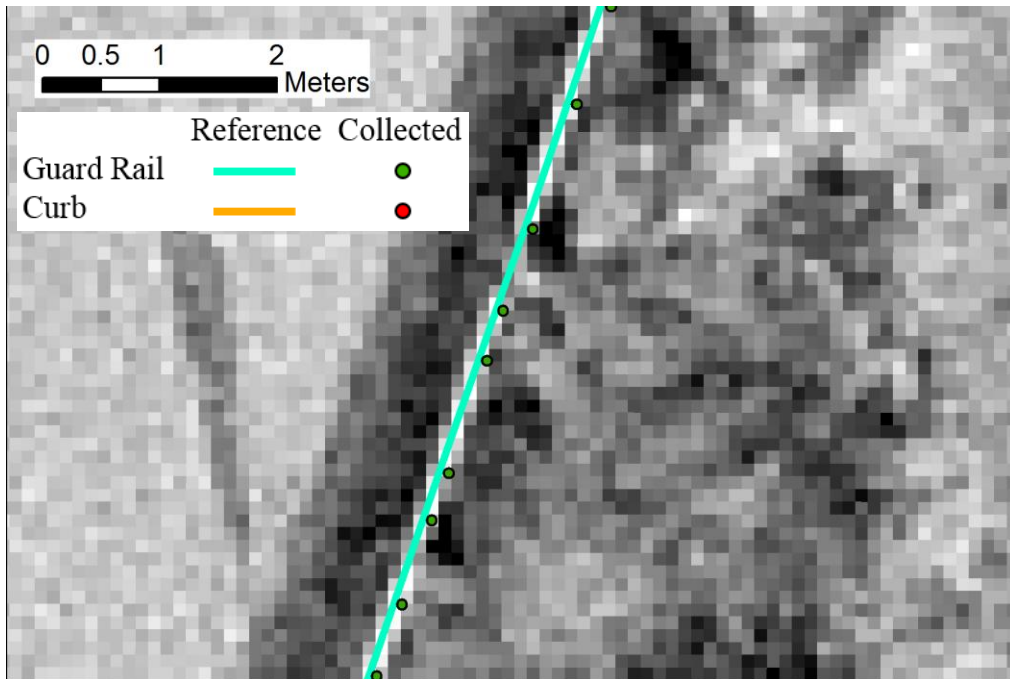(Data from 2 passes in each direction)

Figure 4.13: MN 65 – Location B magnified
(Data from 2 passes)

Figure 4.14 shows location C, an urban intersection. Road features are collected here with minor errors and gaps in the features. Again, interior road features located next to the median are not collected. Due to the thresholding as described above, small curb features such as traffic islands were also not detected. Occasional misplaced curb shapes are detected on the southern half of the intersection. Figure 4.15 shows the north side of the intersection and Figure 4.16 shows the south.
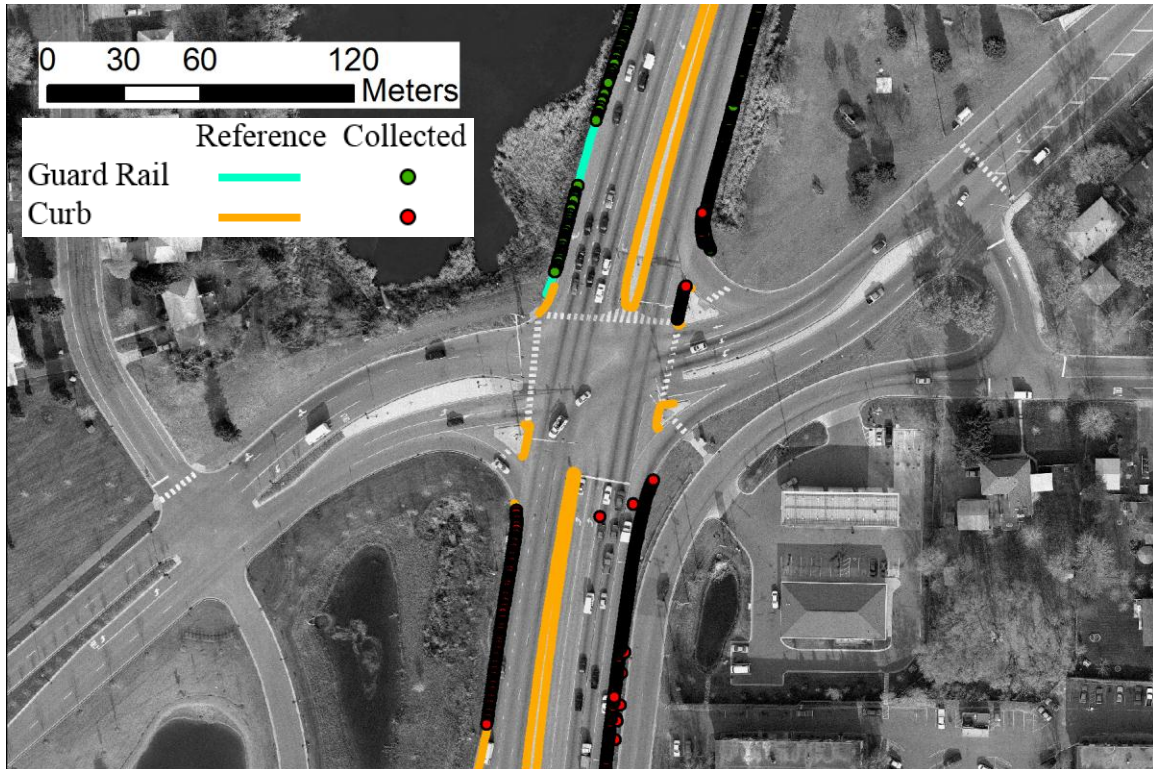

Figure 4.14: MN 65 – Location C
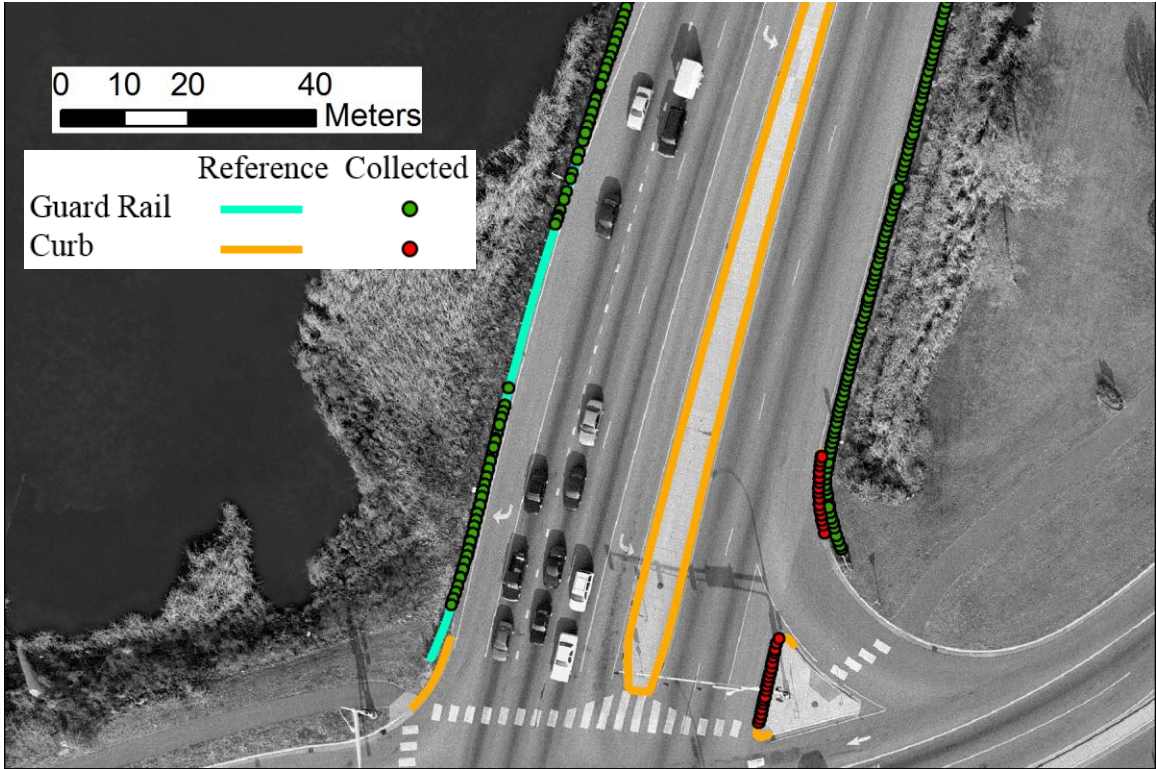(Data from 2 passes in each direction)

Figure 4.15: MN 65 – Location C north side of intersection
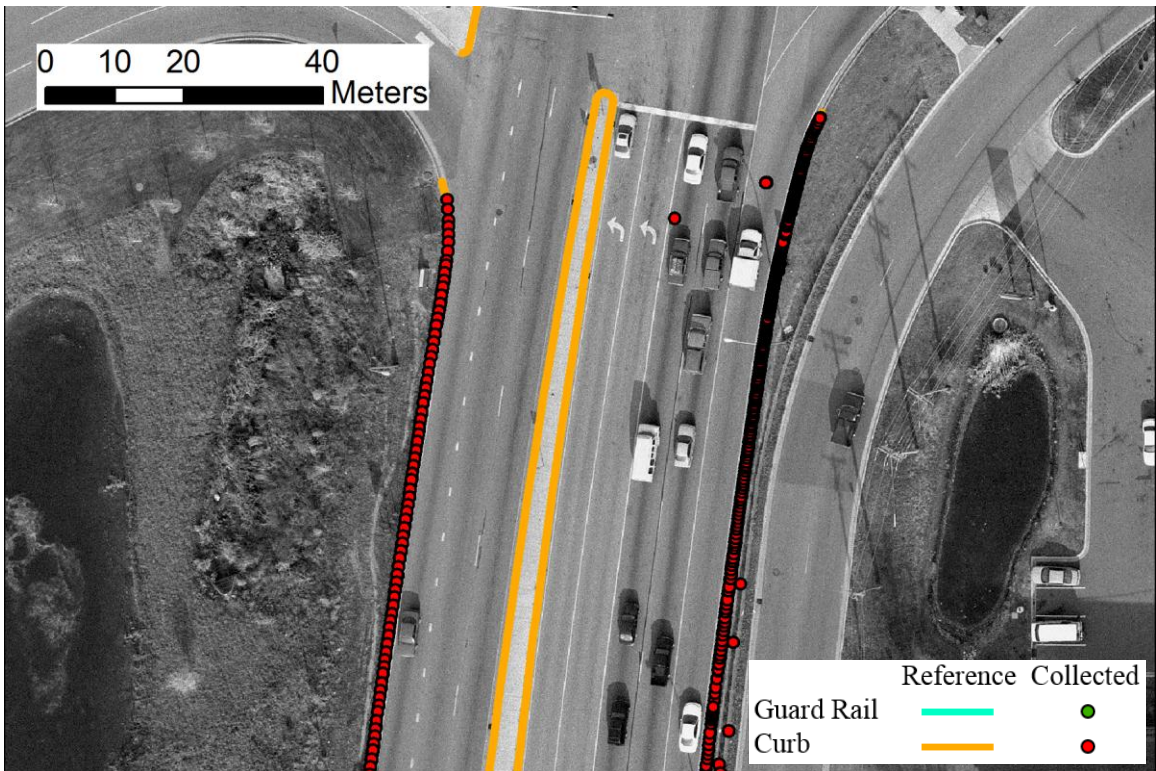(Data from 2 passes in each direction)

Figure 4.16: MN 65 – Location C south side of intersection
(Data from 2 passes in each direction)

Figure 4.17 shows location D, an area where an on ramp splits from the MN 65 towards I-694. Again, there is high accuracy in the areas where features were detected with the occasional misplaced point. This also illustrates an area where the system sometimes does not detect curbs when they are present. This anomaly will need further investigation.

Figure 4.17: MN 65 – Location D
(Data from 2 passes in each direction)

4.3.2   US 63

Figure 4.18 shows the extent of the US 63 test area. From this area, 6 locations are examined more closely below.
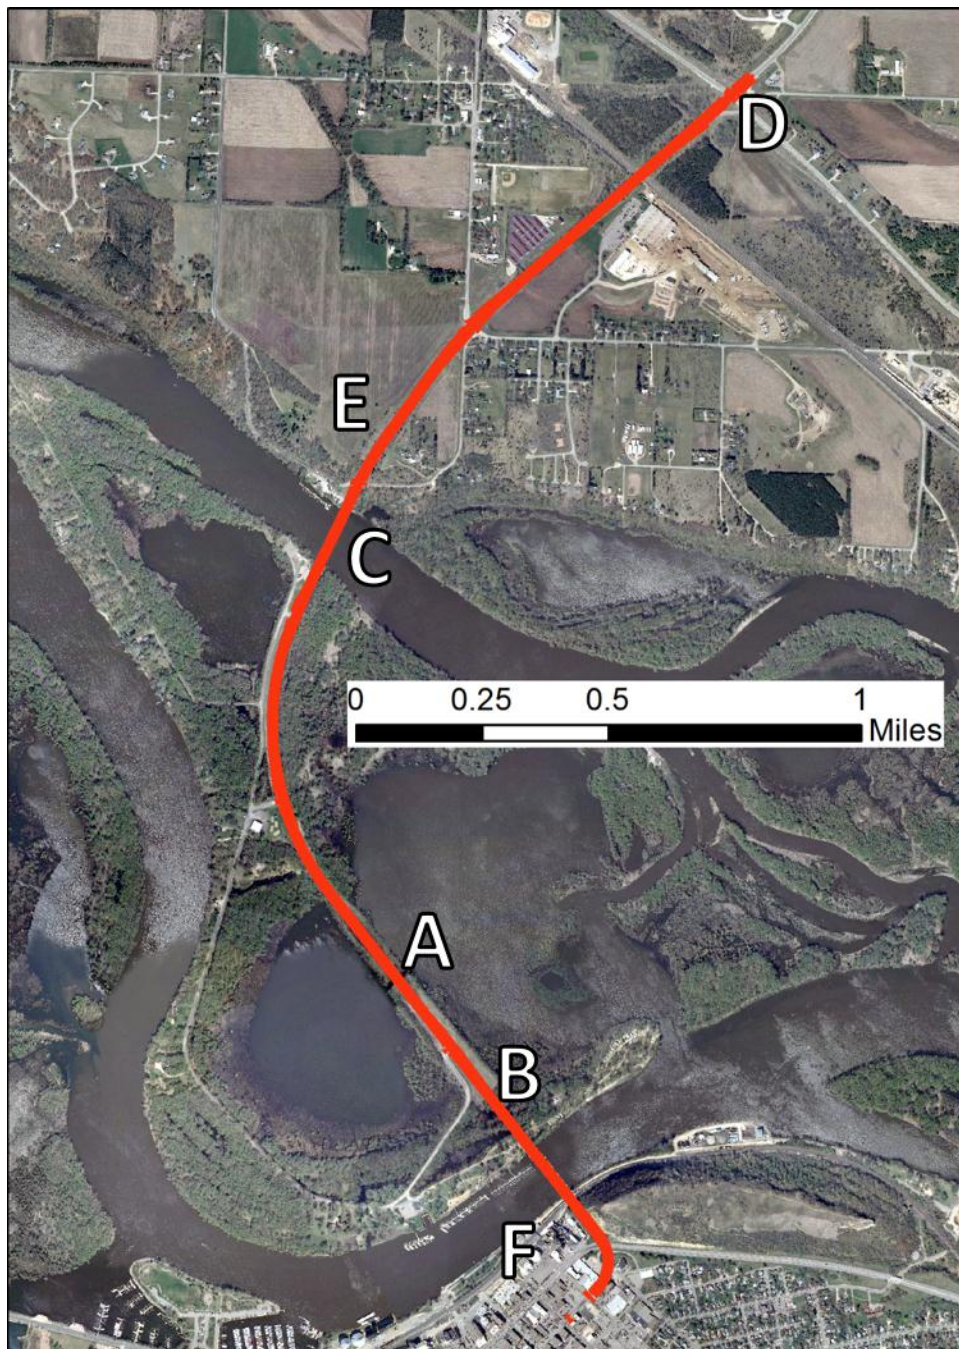

Figure 4.18: US 63 data extent with figure locations

Figure 4.19 shows location A which is a small bridge over a small channel. This is an example of very accurate guardrail detection. There is however a slight discrepancy associated with the placement of the ends of the guardrail due to the end treatments' shape.
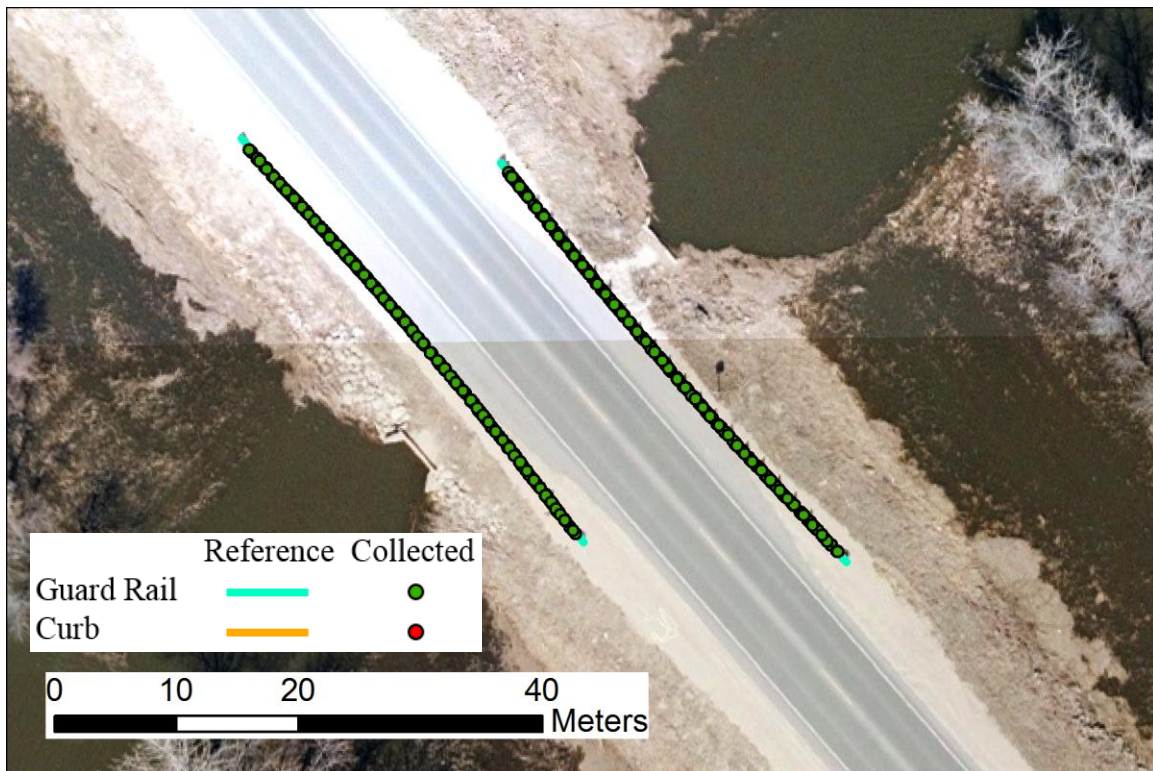

Figure 4.19: US 63 – Location A
(Data from 4 passes in each direction)

Figure 4.20 shows location B which is a bridge over a marshy area just north of the Mississippi river. Here, the detected curb data lines up so well with the reference data, the orange line representing the curb is not visible. In this section, the error is generally less than 3 cm.
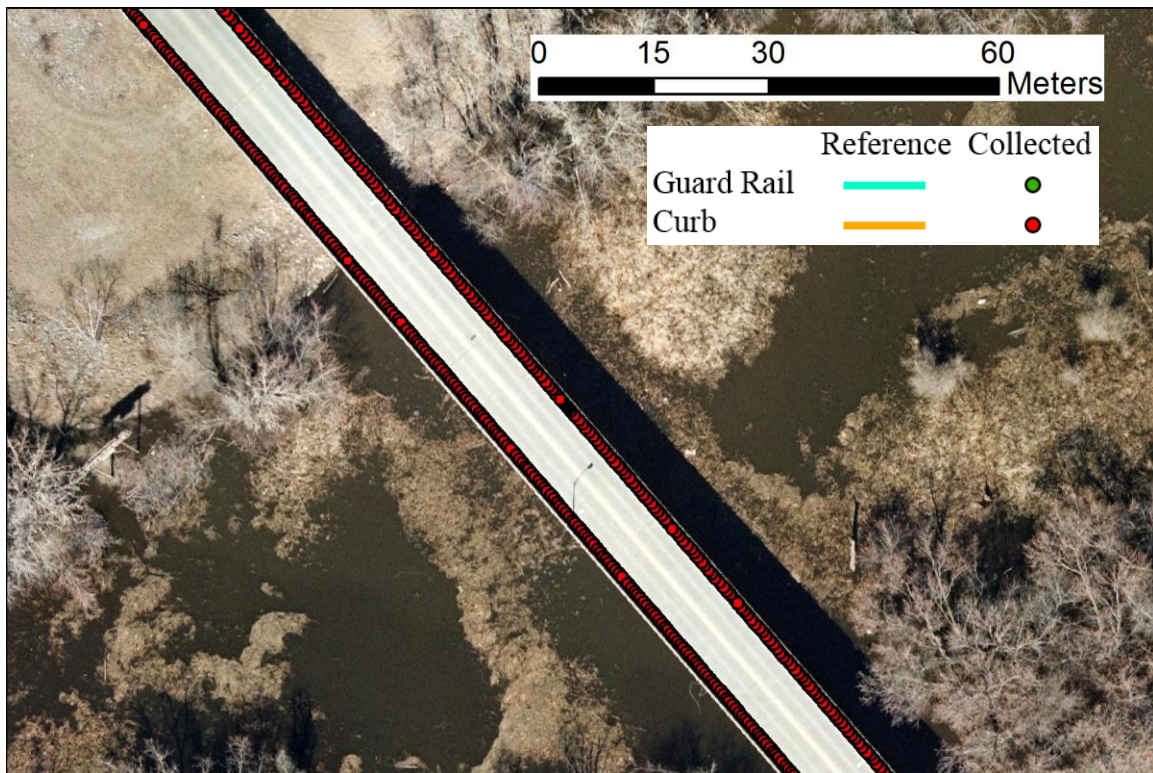

Figure 4.20: US 63 – Location B
(Data from 4 passes in each direction)

Location C is the bridge over the Wisconsin Channel. This is another example of highly accurate road feature detection including their beginning and ends of features. Figure 4.21 shows the south end of the bridge and Figure 4.22 show a magnified view. Figure 4.23 shows the north end of the bridge. For reference, the offsets shown in the zoomed in images range from approximately 0 to 6 cm.
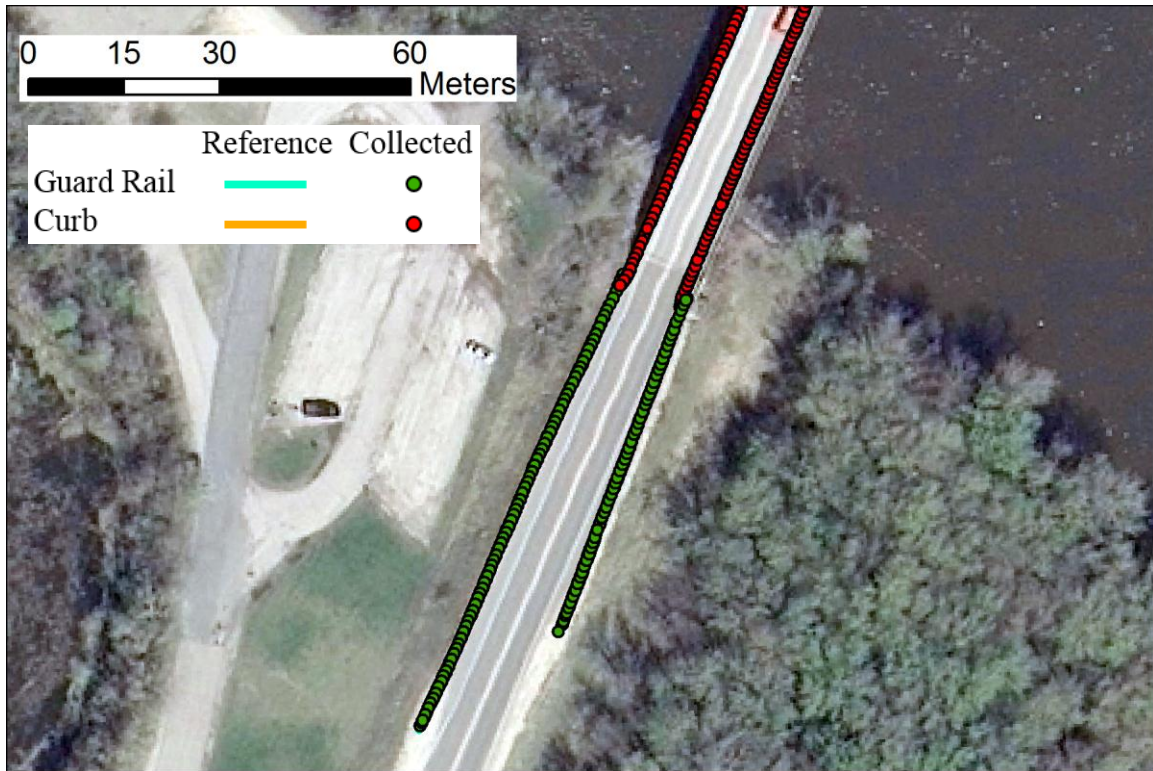

Figure 4.21: US 63 – Location C south end of bridge
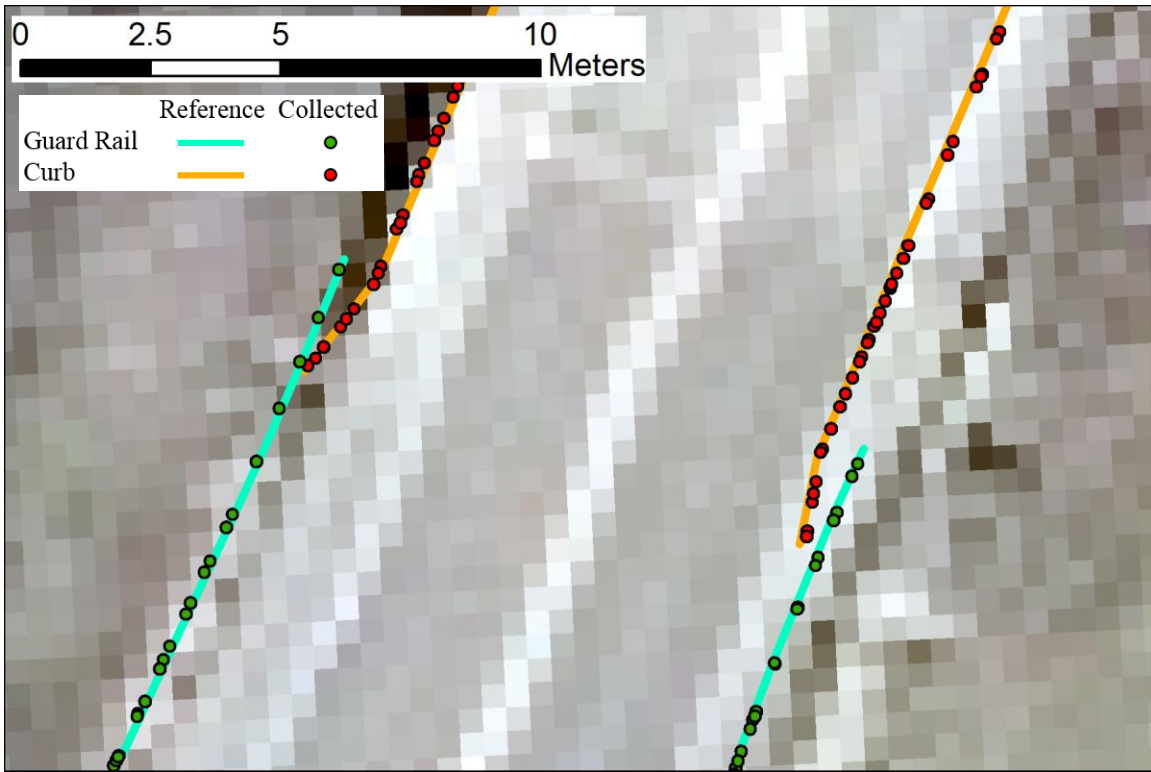(Data from 4 passes in each direction)

Figure 4.22: US 63 – Location C south end of bridge magnified
(Data from 4 passes in each direction)
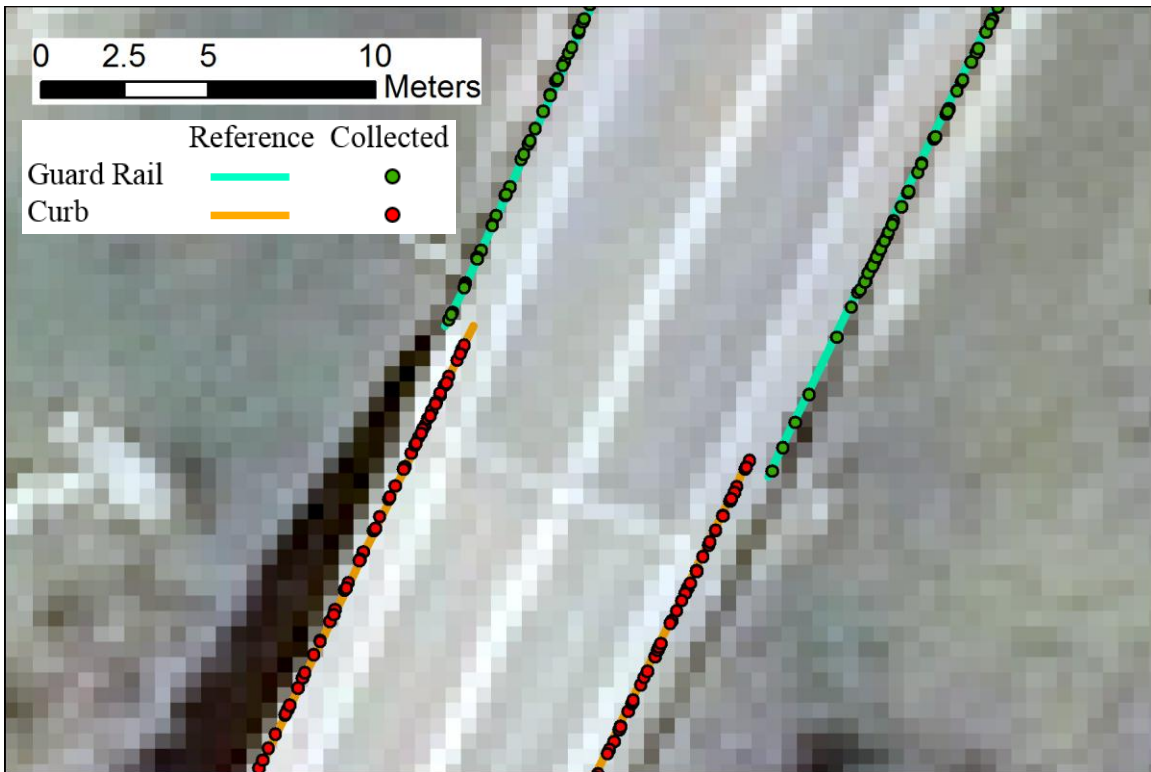
Figure 4.23: US 63 – Location C north end of bridge
(Data from 4 passes in each direction)

Figure 4.24 shows location D which is positioned immediately before an intersection where the right turn lane leaves the through lanes similar to an on ramp. The vehicle never traveled in the protected right turn lane. Note that the system does not detect the guardrail once it is multiple lane widths away from the vehicle. However, the entire guardrail was extracted from the raw reference data because that data collection effort included travel on the right turn lane.
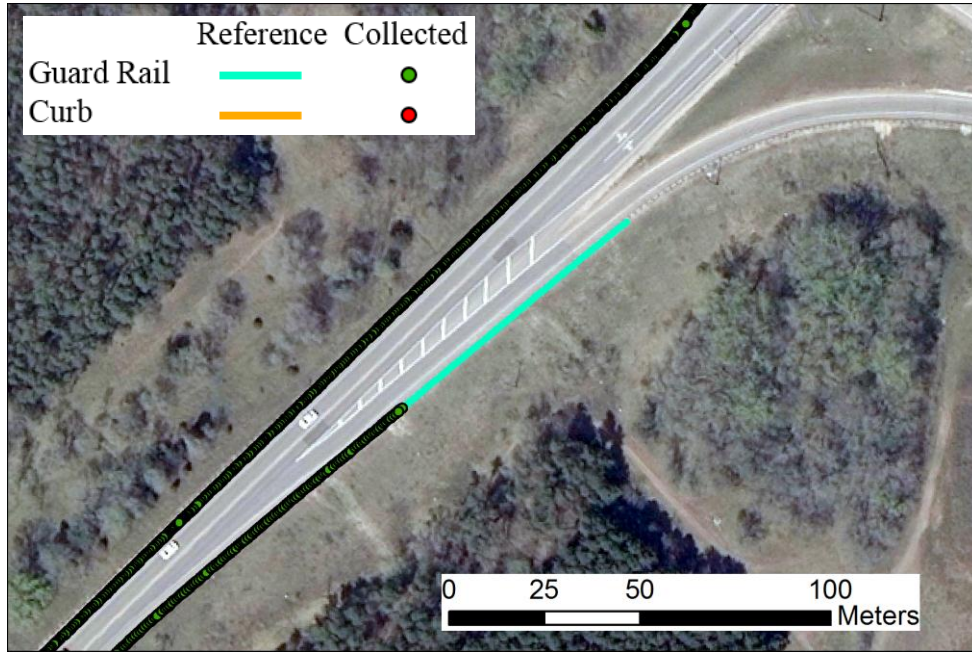
Figure 4.24: US 63 – Location D
(Data from 4 passes in each direction)

Figure 4.25 shows a magnified view of location D. The image shows a zoomed in section of guardrail in the middle of the north side of the road.
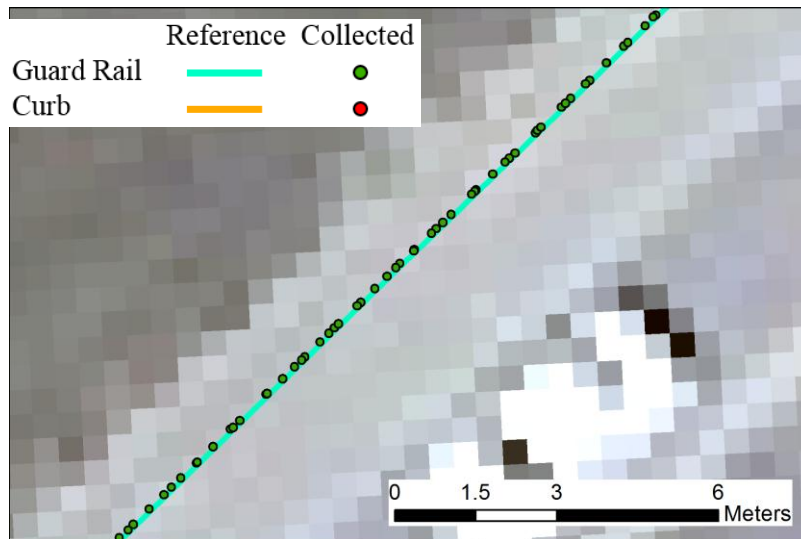
Figure 4.25: US 63 – Location D magnified
(Data from 4 passes)

42

Figure 4.26 shows location E which illustrates an area when a roadside feature is erroneously detected. The system detects curbs anytime the curb shape is present in the data including times when the ground near the road has a step shape. This is one of only two areas where this anomaly occurred in the US 63 test area.

Figure 4.26: US 63 – Location E
(Data from 1 pass)

Figure 4.27 shows the extent of location F which is where US 63 crosses the Mississippi River and enters Red Wing, MN. The presence of a tall bluff is noted in the image with a blue oval marking its approximate location.
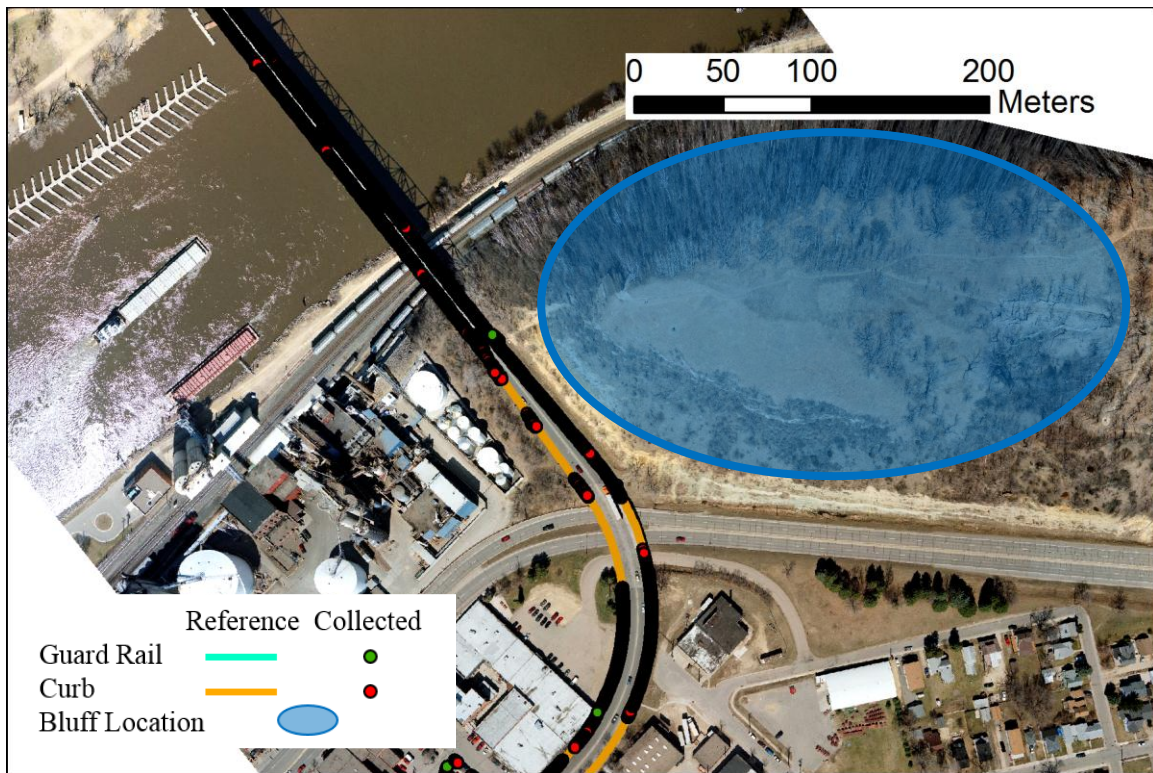
Figure 4.27: US 63 – Location F
(Data from 4 passes in each direction)

Figure 4.28 shows the bridge over the river as it begins to connect to Red Wing. Here, the system correctly detects the presence of curbs, but does not accurately locate them. Issues in this area likely stem from partial sky occlusion due to the bridge geometry and the bluff.

Figure 4.29 shows US 63 after the river bridge as it passes over US 61 in Red Wing. On the northern edge of this figure, there are some errors with the positioning of some curb segments which is likely due to the issues described above. This image also show an example of gaps where the system did not detect curbs that were present in the reference data.

Figure 4.28: US 63 – Location F Mississippi River Bridge
(Data from 4 passes in each direction)

Figure 4.29: US 63 – Location F bridge over US 61
(Data from 4 passes in each direction)

5   Discussion and Conclusions

5.1   Centerline Calculation

For the test area selected, the algorithm was capable of calculating the centerline with approximately 6 cm accuracy. It is noted that this was a section of the road that was selected for its straightforward geometry and therefore 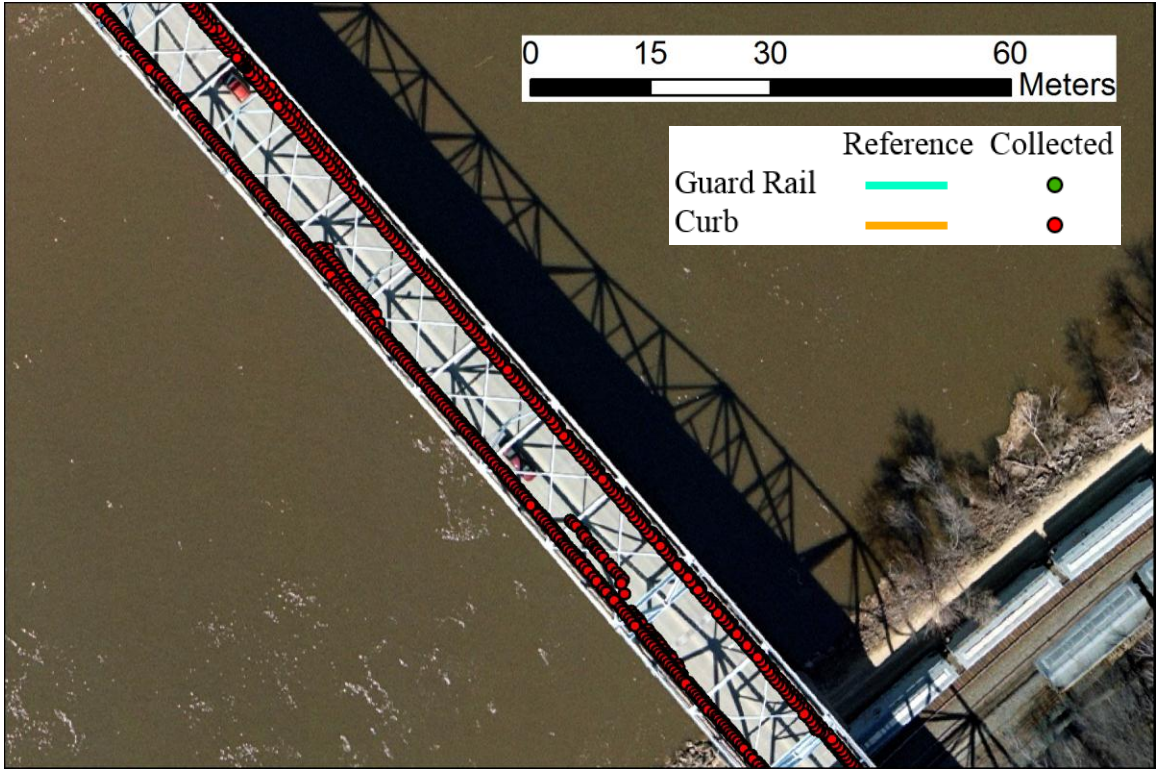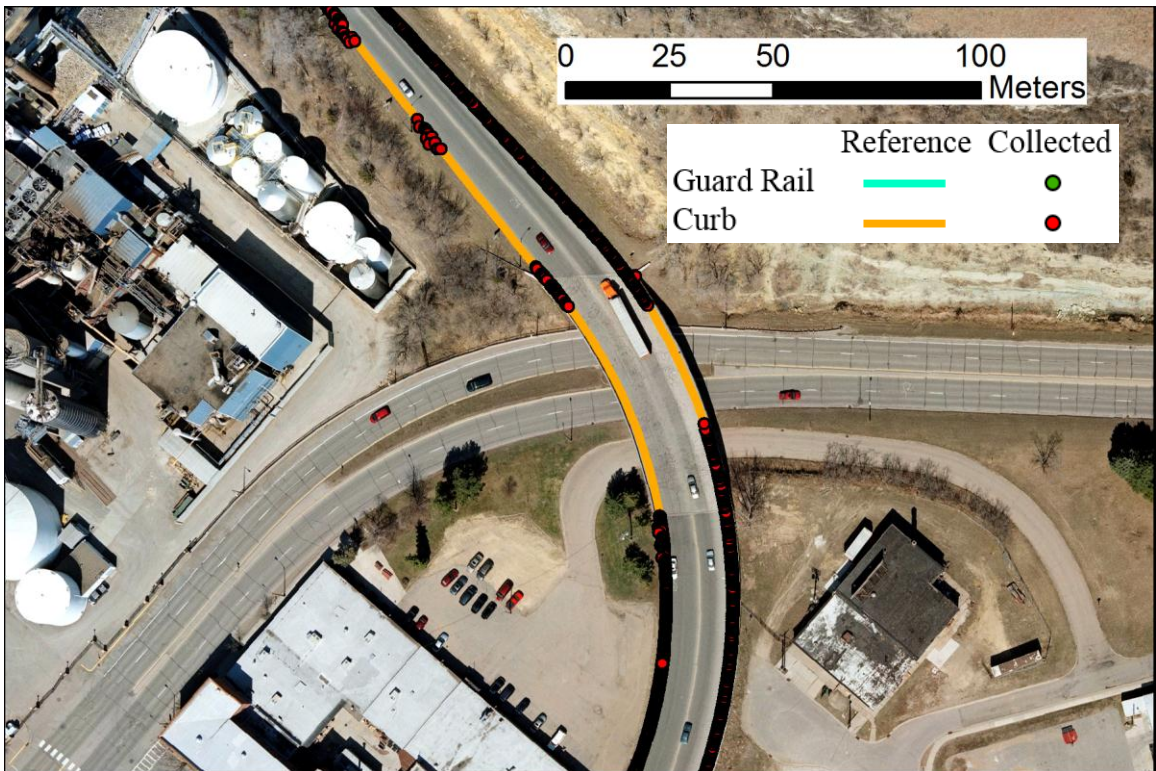this accuracy is not representative of all road segments, but is likely representative of symmetrical, undivided roads. The results reiterate the fact that this technique's accuracy is highly dependent on the road geometry. However for most rural roads, this is likely an acceptable assumption to make.

5.2   Road Feature Detection

The results of the road feature detection evaluation show that in general, if the system detects a feature, it will do so accurately to 10 cm or better. From time to time, the system failed to detect existing road features and occasionally it would detect features that were not present. In areas with poor GPS coverage the system is not capable of accurately locating itself and therefore determining the position of detected road features.

5.3   Sources of Residual Error and Anomalies

5.3.1   Sensor error

Most of our results were within the 0.1 meter specification. What is described here is the source of errors less than 0.1 meters. One such error source stems from the accuracy of the sensors themselves. There is a relatively small error (around 1 to 3 cm) associated with position fixes provided by the GPS depending on the distance to a reference station, the position of the satellites, and local geometry and interference. Additionally, position fixes are only calculated at 10 Hz so position data must be interpolated for most LIDAR scans. As discussed above, this interpolation assumes that between position fixes, the vehicle is moving in a straight line at a constant speed. Determining the position of the vehicle is a critical function that is a prerequisite for calculating the location of all detected features and centerlines. Because of this, any error in the positioning system propagates to the final calculated locations. To reduce the error in the positioning system it would be necessary to get a GPS receiver with a higher accuracy or to augment the positioning system with an IMU or DMI.

The LIDAR also has errors in the raw scan data it collects. However, the stated accuracy of individual scans is around 0.5 to 3 cm. This means that any road feature detected correctly in the scan, may still be off by this amount which would then be further compounded by any error in the GPS antenna location. To reduce the error in the LIDAR, it would be necessary to replace the scanner with a unit that has higher accuracy.

5.3.2   Timing ambiguity and synchronization

Another factor that affects the quality of the data is the ability to tie the sensor data together. For example, finding a feature cross section within a raw data scan is not useful unless it can be correlated to a real world location. To do this, the time that sensors make individual readings is recorded along with the raw data. Then it's possible to correlate the LIDAR data to the GPS data in post-processing in order to determine the feature locations. However, there are errors in determining the times that sensor readings occur. It is trivial to record the times when sensor messages are received by the computer, but it is not always straightforward to record the times when the data was observed by the sensor. This is because the internal clocks on the sensors

themselves can be difficult to synchronize accurately. To mitigate these errors, it would be necessary to obtain a higher end GPS receiver and LIDAR sensor that would be capable of synchronizing their internal clocks to each other.

5.3.3   Algorithm Error

The most notable system errors were failure to detect existing features or detecting features that weren't present. This stems from the algorithm's sensitivity to detecting features in the raw scan data and the need to then filter the results. The efficacy of the algorithm is dependent on the values set to threshold the data. For example there are cutoffs to determine when a shape is detected in the raw scan and for when multiple located features should be combined into a single feature or deleted. Any time a threshold is set, there may be features that are included and excluded erroneously. To reduce these types of errors, additional work could be performed to further tune the algorithms to ensure that feature cross sections are accurately detected in the raw scans and that the system correctly determines which detected shapes correspond to a single, continuous feature.

5.3.4   Feature Occlusion

As discussed above, a possible system error is a failure to detect an existing roadside feature. One possible source of this error is feature occlusion. This means that if the line of sight between the LIDAR and the feature is blocked, it is not possible to detect that feature. This is most common in areas where grass or other greenery grows over or around the roadside feature.

Depending on traffic conditions and road geometry, there may be other vehicles in between the data collection vehicle and the edge of the road. The data collection vehicle is driven in the right-most through lane which may not always be the furthest right lane if there is a protected right turn or other temporary lane. These anomalies are difficult to avoid without requiring multiple passes to account for these conditions.

5.3.5   Error vs Offset

When discussing the system errors, it is important to consider that the reference data isn't perfectly accurate. All the errors calculated in this report are generated by comparing the collected data to the reference data. If there are errors in the reference data, this would affect the validity of the results. The stated accuracy of the reference data is 10 cm or better, so as offsets get below 10 cm, it's difficult to determine whether errors are in the collected data or the reference data. To better account for these issues, higher accuracy reference data would be required.

5.4   Future Work

The results show that the system proposed in this report can successfully detect curbs and guardrails as well as calculate a centerline for roads with simple geometries. However, the results also indicate that there are additional roadside features that could be detected by adding sensors or further developing the algorithms used to post process the raw data.

5.4.1   Additional feature detection

Currently, the processing software can only detect curbs and guardrails. However, it is possible to detect any roadside feature that is large enough to show up in the scan data and that has a cross

section that is sufficiently unique from other shapes appearing in the data. For example, the system may be able to detect and differentiate between different types of barriers including guardrails, jersey barriers, bridge railings, etc. For road sections where there is a large step between the pavement and the ground underneath, the system could identify these unsafe conditions that should be reported to a maintenance group.

Bridge and overpass heights can also be determined. If a portion of the LIDAR scanner's field of view includes the area directly above the vehicle, the distance between the bottom of the bridge and the top of the road can be calculated. However, it is important to note that care would need to be taken to account for the GPS outage caused by traveling under the bridge.

5.4.2   Vision system
A limitation of the LIDAR scanner used here is that it can only detect relatively coarse (at least 5 cm high) geometric features. This would exclude detecting features that are identifiable by visual cues such as lane markings. If the system is augmented with video cameras, it could then also extract this type of feature.

Lane markings would be especially useful when trying to determine road centerlines because if the centerline was identifiable by a paint stripe, it could be detected directly as opposed to being calculated from the vehicle path. Non-centerline lane markings would also be very useful for applications like driver assist systems where understanding the geometry of the lane is critical.

Using a sensor suite incorporating vision to locate the position of road signs would also be useful from an asset management standpoint. This would allow for easier detection of when signs are stolen or are damaged. It could also be useful for adding information about the presence of safety signs in road safety and crash databases. The position of mile reference markers could also be collected which would allow for better data integration with linear referencing systems.

5.5   Next Phase
The next phase of this project is to further investigate methods to improve the system and mitigate issues in areas where there is room for improvement. Development will continue to improve the algorithms to detect roadside features within the scan data and to cluster detected shapes into such features. Video data collected by a camera will augment the system's current capabilities in order to detect lane markings including the fog line and the road centerline. Additional test areas will be identified to better understand the system's capabilities in areas where map data generated by the system could reduce lane departure crashes. The data generated with the system will be shared with GIS staff at the county and state level to determine the applicability of the data and to better understand the needs of these groups to better inform future iterations of the system.

5.6   Final Conclusions
The results of this work show that it is possible to design and implement a relatively inexpensive mobile data collection system that is capable of collecting data to determine the location of a road centerline for roads with straightforward geometries and roadside features such as curbs and guardrails. With the sensors utilized in this project the system was capable of locating these features generally better than 10 cm.

This system is not as accurate or fully featured as commercially available mobile LIDAR scanning systems, but this is reflected in its price. This system can perform its function at a much lower price point than expensive commercial mobile LIDAR scanning systems. Additionally, this system is modular allowing for different sensors to be used as informed by the application. The modularity of the system means it is a platform that can be scaled to the needs of a data collection effort in a cost effective and efficient way.

6   References

[1] CH2M HILL. Final Report for the Minnesota County Roadway Safety Plans. Minnesota Department of Transportation, 2014.

[2] National Cooperative Highway Research Program, *Guidelines for the Use of Mobile Lidar in Transportation Applications*, (Washington D.C.: Transportation Research Board, 2013, NCHRP Report no. 748), http://www.trb.org/Main/Blurbs/169111.aspx.

[3] Mobileye (Internet), *OEM - Mobileye*, (Accessed March 2014), http://www.mobileye.com/markets/oem/.

[4] Minnesota Department of Public Safety – Office of Traffic Safety, *Minnesota Motor Vehicle Crash Facts 2012*. (St. Paul: Minnesota Department of Public Safety, June 2013), https://dps.mn.gov/divisions/ots/reports-statistics/pages/crash-facts.aspx.

[5] L. Alexander, P.M. Cheng, A. Gorjestani, B. Newstrom, C. Shankwitz, and W. Trach, *Driver Assistive Systems for Rural Applications: A Path to Deployment, Volume 1*, (St. Paul: Minnesota Department of Transportation, August 2005, Report no. Mn/DOT 2005-30), http://www.cts.umn.edu/Publications/ResearchReports/reportdetail.html?id=1022.

[6] L. Alexander, P.M. Cheng, M. Donath, A. Gorjestani, B. Newstrom, C. Shankwitz, and W. Trach, *Bus Rapid Transit Technologies: Assisting Drivers Operating Busses on Road Shoulders Volume 1*, (Minneapolis: Center for Transportation Studies, December 2003, Report no. CTS 04-12), http://www.cts.umn.edu/Publications/ResearchReports/reportdetail.html?id=959.

[7] Trimble (Internet), *Trimble MX2 Datasheet*, 2013, (Accessed March 2014), http://www.trimble.com/imaging/Trimble-MX2.aspx.

[8] Riegl (Internet), *Riegl VMX-250 Datasheet*, September 2012, (Accessed March 2014), http://www.riegl.com/nc/products/mobile-scanning/produktdetail/product/scannersystem/6/.

[9] K. Yen, B. Ravani, and T. Lasky, *LiDAR for Data Efficiency*, (Olympia, WA: Washington State Department of Transportation, September 2011, Report no. WA-RD 778.1), http://www.wsdot.wa.gov/Research/Reports/700/778.1.htm.

[10] K. Yen, K. Akin, A. Lofton, B. Ravani, and T. Lasky, *Using Mobile Laser Scanning to Produce Digital Terrain Models of Pavement Surfaces*, (Sacramento, CA: California Department of Transportation, November 2010, Report no. F/CA/RI/2008/xx), http://ahmct.ucdavis.edu/pdf/UCD-ARR-10-11-30-01.pdf.

[11] Trimble (Internet), VRS Installations around the world, (Accessed April 2014), https://www.trimble.com/infrastructure/vrs-installations.aspx.

[12] MnDOT (Internet), MnCORS GPS Network, (Accessed April 2014), http://www.dot.state.mn.us/surveying/cors/.

[13] Trimble (Internet), Trimble R7, (Accessed June 2014),
http://www.trimble.com/Survey/trimbler7gnss.aspx.

[14] Trimble (Internet), GNSS Antennas for Heavy Construction, (Accessed June 2014),
http://construction.trimble.com/products/marine-systems/gnss-antennas-for-construction.

[15] Sick (Internet), *Sick LMS511 Online Data Sheet*, (Accessed June 2014),
https://www.mysick.com, search for: LMS511-20100.

[16] Velodyne (Internet), *Velodyne HDL-64E S2 Datasheet*, (Accessed June 2014),
http://velodynelidar.com/lidar/hdlproducts/hdl64e.aspx.

[17] Riegl (Internet), *Riegl VQ-250 Datasheet*, (Accessed June 2014),
http://www.riegl.com/nc/products/mobile-scanning/produktdetail/product/scanner/22/.

[18] W. Zhang, "LIDAR-Based Road and Road-Edge Detection", *2010 IEEE Intelligent Vehicles Symposium 2010*, June 2010, 845 - 848.

[19] A. Mancini, E. Frontoni, P. Zingaretti, "Automatic Road Object Extraction from Mobile Mapping Systems", *IEEE/ASME International Conference on Mechatronics and Embedded Systems and Applications (MESA), 2012,* July 2012, 281 - 286.

[20] Highway Performance Monitoring System (HPMS), (Accessed June 2014),
http://www.fhwa.dot.gov/policyinformation/hpms.cfm.

[21] Office of Highway Policy Information, *Highway Performance Monitoring System Field Manual.* (Washington D.C.: Federal Highway Administration, March 2014),
http://www.fhwa.dot.gov/policyinformation/hpms/fieldmanual/.

[22] E. Arpin, C. Shankwitz, and M. Donath, *A High Accuracy Vehicle Positioning System Implemented in a Lane Assistance System when GPS is Unavailable*, (Minneapolis: Center for Transportation Studies, July 2011, Report no. CTS 11-18),
http://www.cts.umn.edu/Publications/ResearchReports/reportdetail.html?id=2042.