

Vehicle Localization using an Automotive-Fixed Laser

A thesis

Submitted by

Timothy Fleck

In partial fulfillment of the requirements

For the degree of

Master of Science

In

Mechanical Engineering

TUFTS UNIVERSITY

August 2013

Adviser: Jason Rife, Ph.D.

Abstract

Advanced Driver Assistance Systems (ADAS) offer safety advantages to drivers in many situations. A modern Lane-Departure Warning System (LDWS), for example, can alert a driver when an unintended highway lane departure is imminent, preventing potentially deadly crashes or collisions. Enhanced LDWS for urban environments will require very accurate lateral positioning information, even in environments in which visual features (road-lane markers) may be sparse and in which buildings may occlude or corrupt GPS signals. To achieve such a capability, this thesis proposes a low-cost automotive positioning concept for urban environments, a concept we call Registration using Automotive-Fixed Laser (RAFL). The approach registers the car location on a reference map using data from a laser rangefinder and a dead-reckoning system. The concept can estimate both longitudinal and lateral position. A preliminary road test of our concept demonstrated RAFL accuracy can exceed that of conventional GPS processing (using WAAS-corrected L1 C/A measurements) in high multipath urban environments.

Acknowledgements

Jonah Kadoko, an undergraduate in the mechanical engineering department at Tufts, calibrated our equipment and collected the data for our experiment.

Contents

Abstract	ii
Acknowledgements	iii
List of tables	vi
Chapter 1: Introduction	1
Chapter 2: Registration with an automotive-fixed laser	7
Introduction	7
Section 1: RAFL Concept	7
Section 2: RAFL Algorithm.....	12
Section 3: Experimental System	17
Section 4: Experiment Results and Analysis	19
Chapter 3: Concepts for improving our method	24
Introduction	24
Concept 1: Map improvements.....	25
Concept 1.1: Adjusting the map using ground-based measurements.....	26
Concept 1.2: Adding confidence levels to reference map segments	30
Concept 2: Analyzing the sensor input	36
Concept 2.1: Median filtering of the laser input	36
Concept 2.2: Shape matching algorithm	42
Concept 2.3: Obstacle detection algorithm	45
Concept 3: Improving the comparison function.....	49

Symmetric cost functions for error minimization	50
Asymmetric cost functions for error minimization	58
Section 4: Effects of combined improvements	63
Chapter 4: Conclusion.....	69
Contributions	69
Future Work	70
Impact.....	72

List of tables

Table 1: Improvements to our method presented in Chapter 2 24

Table 2: Confidence values assigned to landmarks 33

Table 3: Z score values for different cost functions 58

Table 4: Improvements to our method presented in Chapter 2 63

Table 5: Effects of concepts for improvement 67

List of figures

Figure 1: Concept block diagram	8
Figure 2: Reference map overlaid on satellite image	9
Figure 3: Data acquisition and variables	10
Figure 4: Laser measurement versus time. Since the rangefinder points in the negative un direction, its value is always negative.....	11
Figure 5: Laser measurement (blue) overlaid on the reference map (green) prior to precise registration.	12
Figure 6: Experimental setup.....	18
Figure 7: GPS (red) and RAFL (blue) position estimates. The green line is the reference map.	20
Figure 8: Estimates of lateral position	21
Figure 9: Lateral offsets.....	22
Figure 10: Laser rangefinder reconstruction versus satellite map	27
Figure 11: Reference map adjusted after consideration of disparity with rangefinder data	29
Figure 12: Effect of revising the map with ground data.....	30
Figure 13: A brick wall close to the road	31
Figure 14: A house separated by a hedge, driveway, road and parking lot	31
Figure 15: A statue surrounded by plants	31
Figure 16: A satellite photo showing features of varying confidence.....	32
Figure 17: Buildings near 500-meter mark, with their confidence factors (green) and distance markers (blue).....	34
Figure 18: Effect of adding a confidence factor to the reference map.....	35
Figure 19: Unfiltered laser measurements.....	37

Figure 20: This chart shows a modest level of filtering. The filtered laser measurement is shown in black, with spikes removed shown in red. The values of n and c are 5 and 20, respectively.	39
Figure 21: This chart compares filtering with $c = 10$ to filtering with $c = 20$. Shown in red are spikes removed at $c = 10$, but not at $c = 20$. The value of n is 5.	40
Figure 22: This chart compares a 5-point median filter ($n = 5$) to a 15-point median filter ($n=15$). The 5-point filter removes the blue spikes. The 15-point filter removes the red spikes. Spikes removed in both cases are not shown. Both filters used a value of 20 for c	40
Figure 23: Filtered and unfiltered range values.....	41
Figure 24: Results with (green) and without (red) filtering, compared to GPS (blue)	42
Figure 25: Points identified as linear.....	44
Figure 26: Effects of shape matching (green).....	45
Figure 27: Laser rangefinder readings.....	46
Figure 28: Likely obstacles detected by the rangefinder are identified in red.....	48
Figure 29: Position estimates with (green) and without (red) obstacle detection.....	49
Figure 30: The quadratic cost function for various values of a	51
Figure 31: An improved cost function.....	52
Figure 32: Inverse quadratic plot with $b_1 = 10$, $b_2 = 0.25$, $b_3 = 5$	54
Figure 33: Inverse quartic plot with $b_1 = 10$, $b_2 = 0.1$, $b_3 = 5$	55
Figure 34: Shifted error function plot with $b_1 = 5$, $b_2 = 1$, $b_3 = 2$, $b_4 = 0$	56
Figure 35: Sample measurements (red) and reference map (green) at 55 seconds	57
Figure 36: Asymmetrical error function cost function.....	60
Figure 37: Performance comparison of cost functions. BLUE-GPS. RED-estimates using the quadratic cost function. MAGENTA-estimates using the inverse quadratic cost function. GREEN-estimates using the asymmetrical error function cost function.	62
Figure 39: The estimated trajectory using the inverse quadratic cost function.....	65
Figure 40: The estimated trajectory using the full suite of improvements.....	66

Chapter 1: Introduction

Advanced Driver Assistance Systems (ADAS) play an increasingly important role in modern transportation. These technologies promise to enable drivers and passengers to travel more quickly and safely with less effort on the part of the driver. ADAS are any systems that process sensor input to help guide a human driver. They include technologies for navigation, lane keeping, cruise control, maintaining a safe following distance, and automated braking to avoid an accident. ADAS also have the capability to warn drivers about unsafe situations in time for the driver to react, and further development of technologies used in ADAS can even be used to maneuver a vehicle autonomously. Their potential benefit is enormous, as each year over 1.3 million people [1] die from traffic-related injuries, with tens of millions more injured but surviving [2]. Over 90% of traffic accidents involve human error [3], the effect of which could be mitigated with ADAS technology.

Nearly all ADAS, and the task of fully autonomous driving, require determination of vehicle location. Collision avoidance systems, navigation systems, intelligent braking assistance, adaptive cruise control, and lane departure warning systems all depend on accurate estimates of vehicle position relative to the road, to other vehicles, and to other landmarks or road users.

Lane departure warning systems (LDWS), which are already in place in some consumer vehicles, calculate when a vehicle is moving or about to move out of its current travel lane. By alerting the driver in these situations, LDWS can prevent unintentional lane departures, which can lead to a crash into a roadside object, a collision with an oncoming vehicle, or a lateral collision with a vehicle moving in the

same direction in an adjacent lane. According to a survey by Kuehn *et al.* of accidents reported to German insurers between 2002 and 2006, over 25% of all vehicle accidents involved unintentional lane departures [4]. An accurate, robust, and affordable LDWS could save many lives. In an overview of advanced driver assistance systems, Lu notes that a LDWS or lane-keeping assistant (which is similar to LDWS, but may involve direct control of steering) based on absolute positioning would need an accuracy of around 0.3 m [5]. The next four paragraphs describe different solutions proposed for providing accurate estimates of position.

Global Navigation Satellite System (GNSS) receivers, already installed in many cars for navigation purposes, are not generally accurate enough for LDWS. Lu estimates the accuracy of the WAAS-enabled code phase at 2-4 meters, depending on the system used [5]. Using the carrier phase of GPS, Bajikar was able to determine location of various survey nails with a standard deviation of 2cm and 35cm in the lateral and longitudinal directions [6], respectively. However, the carrier phase does not give an absolute measurement. It needs to be supplemented with the code phase, which does not have the required accuracy, or by another technology that can determine absolute position. Both code phase and carrier phase can be degraded by multipath errors [7], which are common in urban environments. For these reasons, GNSS technologies are unlikely to provide a sufficiently accurate and robust LDWS position estimate.

A system relying on road-based magnetic or RFID markers can provide the needed accuracy. This system requires two components—beacons along a roadway, and a detector on each vehicle. Chan has shown that a sensor can detect the position of magnetic tape placed along a travel lane within 5-8 cm [8], and the system is simple.

However, due to large upfront costs, there are no plans in place for widespread installation of markers, tape or beacons.

Vision-based locating is promising, but currently unreliable, especially for urban applications. Some video-based locating systems work by analyzing brightness gradient in an image to identify lane markers [⁹], which are generally brighter than the surrounding pavement. This type of locating system generates estimates of lateral position only. LDWS based on vision are currently available on consumer vehicles [¹⁰]; however, these systems are only meant for use on highways. A method developed by McCall and Trivedi can detect a car's lateral position within 9-14 cm, although the accuracy varies with road conditions [¹¹]. In poor weather or on inadequately marked roads, the systems may not function at all. As such, LDWS systems are not yet available to operate in urban and suburban environments, where lane markers are not ubiquitous. Moreover, considerable post-processing would be needed to make these algorithms robust and to account for obstacles, outliers, and other conditions that occur more frequently in urban environments than in more structured highway environments.

Some systems utilize laser rangefinders to determine the location of the vehicle compared to known landmarks. Unlike GNSS they work best in urban canyons, and they have higher accuracy [¹²]. Rather than detect specially installed markers, the rangefinders detect whatever objects are already in place, relying on a map of these objects. And unlike vision systems, they function equally well in every road condition, with any or no visible light, and in all but the most adverse weather. The rest of the thesis will focus on specific ways laser rangefinders can be used in ADAS applications, particularly in positioning.

At its simplest a laser rangefinder consists of an emitter, a detector, and a clock. Infrared pulses are emitted, reflect off a target object, and return to the detector. The clock measures the time between when the pulse is emitted and when it is detected, and the distance traveled is calculated from this time and the speed of light. Some applications employ one or more of these emitter-detector pairs mounted on a rotating turret on a vehicle roof so that they can take distance readings in multiple directions. Sometimes the lasers rotate in a vertical plane as well. These types of setups are referred to in this thesis as Scanning Laser Rangefinders (SLRs).

Scanning Laser Rangefinders are often used for Simultaneous Localization and Mapping (SLAM), where an autonomous vehicle navigates in a previously unknown environment. Xu et al. [13] used a laser scanner to map an indoor environment. After a series of points were collected at one vehicle location, they were arranged into line segments. A robot could determine its position at a later time by looking for line segments of the same length. Other demonstrations of laser-based SLAM have taken place in tunnels, in controlled outdoor environments, and in urban settings. Soloviev [14] used a line segment comparison method to guide an autonomous vehicle in alleyways between buildings. Building edges were arranged into line segments, which were matched to previous line segments as before, accounting for scanner tilt that would be more common in a less controlled environment. Madhavan [15] and Joerger [16] used kalman filters to look at point landmarks instead of line segments. Madhavan, rather than performing SLAM, created a map in one pass through a mining tunnel, and then stored the map and used it at a later time with the same laser scanner. Also using a stored map, a team at Stanford and Google [17] developed an autonomous vehicle that could navigate city streets with traffic and pedestrians, using a suite of advanced

sensors. These included a 64-beam LIDAR and two other smaller scanning LIDAR scanners, as well as four cameras, six radars, and a combined INS and GPS unit.

This thesis explores the possibility of using a vehicle-fixed laser rangefinder for vehicle localization, rather than the SLRs more commonly employed for robotics and autonomous vehicle research. We call our approach Registration using an Automotive-Fixed Laser (RAFL). The main motivation for RAFL is cost. In 2011 prices, the cost of a low-grade scanning laser rangefinder is about a tenth of the cost of a new four-door sedan, making their implementation in consumer automobiles unlikely. High-end, turret-based scanning laser rangefinders, such as those used by the Stanford/Google vehicle, may cost more than five to ten times as much as the car. By comparison, laser rangefinders without the scanning mechanisms have about one-twentieth the cost of the lowest-end scanning-beam laser rangefinders. In addition, fixed-beam lasers generate a smaller data set, requiring less computing power to process. If a LDWS using a fixed-beam laser can be developed, it could be commercialized sooner and more widely.

This thesis describes three contributions to this field:

1. The development of the RAFL method, and the mathematical framework for it
2. The implementation and demonstration of our method in a road test
3. A summary of concepts for improving the method further

We have developed a positioning system using a vehicle-fixed laser rangefinder, dead reckoning, and a map. We then collected data while driving through an urban environment, and wrote an algorithm to provide position estimates based on the data. Although more work is needed to confirm the accuracy estimates and to test the

method on different data sets, we believe we have a capable of estimating position with meter-level accuracy.

The remainder of the thesis is organized into three chapters, each divided into four subsections, and a brief conclusion after the second chapter. The next chapter introduces our RAFL concept and the key components of it. Then it describes the mathematical algorithm at the core of our method. The third chapter builds on the first chapter, and explains how improvements to the reference map, to the sensor filtering, and to the algorithm can improve the robustness and of our method in non-ideal situations. The conclusion summarizes the results from both prior chapters, and suggests further opportunities for research and development.

Chapter 2: Registration with an automotive-fixed laser

Introduction

This chapter demonstrates the feasibility of locating a car with a fixed laser rangefinder and a dead-reckoning system. Our system combines a few affordable elements to replicate the behavior of SLR systems. Because it depends on a registration between a map created from collected data, and a reference map, we call our system Registration with an Automotive-Fixed Laser (RAFL).

The first section of this chapter explores the RAFL concept in detail, and explains the working principles. The second section develops a mathematical model for the system, and an algorithm for performing RAFL with the inputs used. The third section shows how a method for collecting data and using our method to calculate the vehicle positions at which the data was taken. We did not attempt to test our method in real-time. The fourth section analyzes the results of using our algorithm on a small stretch of road with a well-developed reference map. We compared our estimates to data from a video and a GPS receiver for independent verification of our results.

Section 1: RAFL Concept

Our concept couples the laser rangefinder with dead reckoning to estimate vehicle position along a path with a surveyed profile of roadside buildings. A block diagram for the system is illustrated in Figure 1. By taking continual laser range

measurements and using the vehicle's velocity profile to align these measurements along a path, we are able to construct a rough map of the roadside at the height of the laser rangefinder. Comparing this profile to a satellite image of the surrounding area, we determine the car's location relative to the buildings in the image.

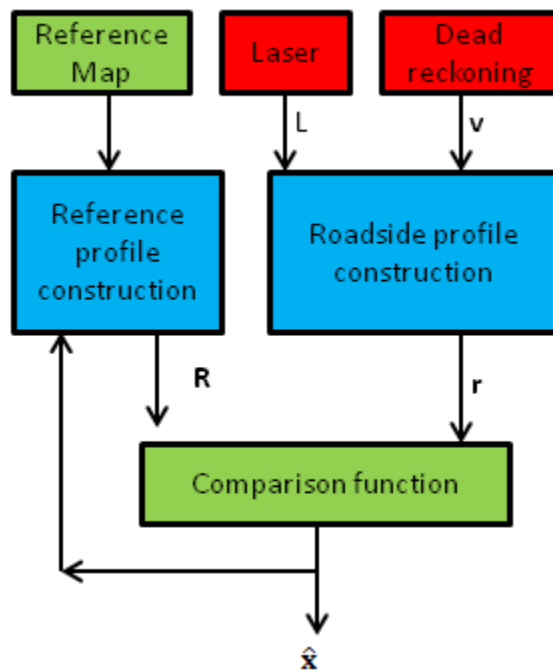


Figure 1: Concept block diagram

In this thesis, the reference map was created from a satellite image of the area. Building corners are manually selected in an image, and the lines between them are automatically drawn to represent building edges. A reference map overlaid on the satellite image is shown in Figure 2. In concept, this process for building reference maps from satellite data could be automated. Alternatively, an even more accurate roadside profile might be generated by replacing the satellite data with data collected

by specialized survey vehicles equipped with accurate (and expensive) positioning equipment.



Figure 2: Reference map overlaid on satellite image

By driving a vehicle equipped with a RAFL system (featuring dead reckoning sensors and an automotive-fixed laser) along a pre-surveyed route, a measured roadside profile can be generated in real time and compared to the surveyed roadside profile.

Figure 3 shows how the data are collected. A car traveling along a road takes measurements of a building profile by shining a laser to the right. Because an automotive-fixed laser is employed, no information is collected when the vehicle is

stationary. Therefore, new laser range information can only be collected when the car is moving. Given many laser data points, with their origins related by dead reckoning, information about objects detected from different vantage points can be assembled into a map.

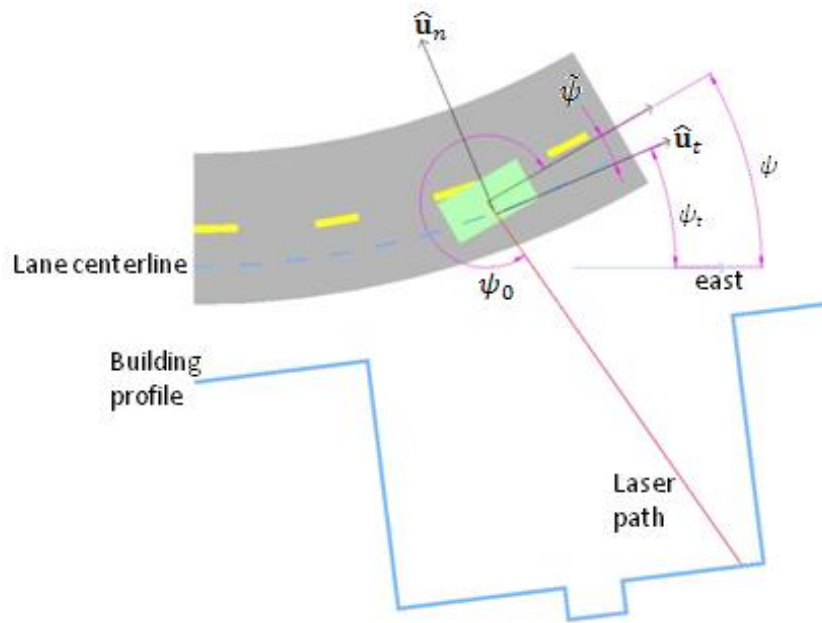


Figure 3: Data acquisition and variables

The raw laser range measurements are collected and stored in sequence with their associated time steps as shown in Figure 4. A dead reckoning system computes the displacement at each time. Figure 5 shows the data from Figure 4 rotated and superimposed on top of the reference map. Aligning (or registering) the two graphs more precisely gives an estimate of position. The registration algorithm for calculating a position estimate is presented in the next section.

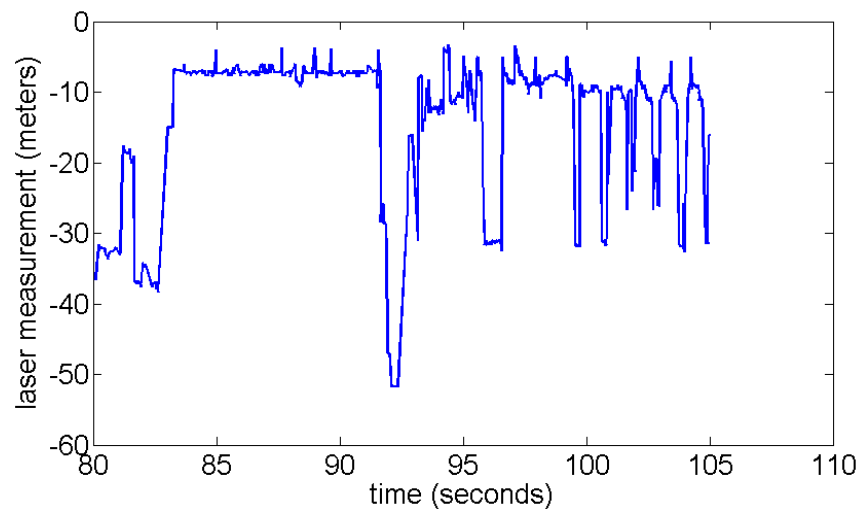


Figure 4: Laser measurement versus time. Since the rangefinder points in the negative \hat{u}_n direction, its value is always negative.



Figure 5: Laser measurement (blue) overlaid on the reference map (green) prior to precise registration.

Section 2: RAFL Algorithm

An optimization-based strategy is used to compute vehicle position by determining the position which best aligns a measured roadside profile to a reference roadside profile (obtained from a map database).

The amount of data used to construct the measured roadside profile is selected such that the “length” of the roadside profile (i.e. the distance traveled along the road) is always constant. The length of time required to travel an arbitrary distance p along the road direction is T . The following equation relates the distance p to the time duration T given that the vehicle velocity \mathbf{v} and the direction of the road (described by the unit vector $\hat{\mathbf{u}}_t$) may vary over time. The equation is computed backwards from the present time, t^* .

$$p(T; t^*) = \int_{t^*-T}^{t^*} \mathbf{v} \cdot \hat{\mathbf{u}}_t dt \quad (1)$$

Inverting this equation, it is possible to compute how much time T has elapsed for the vehicle to travel a specified distance p_{ref} .

$$T = p^{-1}(p_{ref}; t^*) \quad (2)$$

Because the reference length of the roadside profile p_{ref} is held constant, and because data are assumed to be sampled at a roughly constant rate, more data points are used to construct roadside profiles when a vehicle is moving slowly than when it is moving quickly. (By comparison, if roadside profiles were constructed for a fixed T , a car stuck in traffic or at an intersection for the entire time T would have only one point in its roadside profile.)

Because data are sampled discretely, equations (1) and (2) need to be converted into a discrete form. In this thesis, for simplicity, we use the forward Euler approximation to discretize continuous integration. The forward Euler approximation of (1) is given below.

$$p(K; k^*) = \sum_{k=k^*-K}^{k^*} \mathbf{v}(k) \cdot \hat{\mathbf{u}}_t(k) \Delta t(k) \quad (3)$$

The number of time steps used in constructing the measured roadside profile is $K+1$, where K is obtained from the following equation.

$$K = \arg \min \left(\left| p_{ref} - \sum \mathbf{v}(k) \cdot \hat{\mathbf{u}}_t(k) \Delta t(k) \right| \right) \quad (4)$$

The first step in constructing the roadside profile is to estimate the approximate trajectory of the vehicle during the previous $K+1$ time steps. Dead-Reckoning Sensor (DRS) data are used to reconstruct this trajectory. DRS alone cannot identify the absolute position $\mathbf{x}(k)$ of the vehicle as a function of time; however, the DRS data can provide measurements of position displacement $\mathbf{d}(k)$ over the time window of interest. Here we define position displacement as zero at the current time (k^*) and negative at preceding time steps, according to the following formula.

$$\mathbf{d}(k) = \mathbf{x}(k) - \mathbf{x}(k^*) \quad (5)$$

The vector \mathbf{x} is the position of the vehicle, and its components are chosen to be coordinates along the $\hat{\mathbf{u}}_t$ and $\hat{\mathbf{u}}_n$ directions. The velocity vector \mathbf{v} , used in equations (1) through (4), is the time-derivative of \mathbf{x} . The coordinate system moves along the road with the car, so it is different at each time t^* . For convenience, we assume the origin of the coordinate system is always shifted (in the $\hat{\mathbf{u}}_n$ direction) to the center of the lane, regardless of the vehicle's lateral position.

The DRS system is also assumed capable of estimating the vehicle heading at each time step during the window of interest, in an earth-fixed (North-East) coordinate

system. Heading relative to the local-road coordinate system is denoted $\tilde{\psi}$. Absolute heading ψ relative to East can be converted to local-road coordinates by introducing an offset angle ψ_t , which accounts for the road direction.

$$\tilde{\psi}(k) = \psi(k) - \psi_t(k) \quad (6)$$

The offset ψ_t can be looked up in a map database (as a function of the predicted car position).

The measured roadside profile can be obtained by adding laser ranges to each point along the vehicle's trajectory $\mathbf{d}(k)$. If the laser is arranged at an angle ψ_0 from the car's front, and measures a distance $L(k)$ to the nearest roadside object, the measured roadside profile $\mathbf{r}(k)$ can be constructed from the following formula.

$$\mathbf{r}(k) = \begin{bmatrix} L(k)\cos(\tilde{\psi}(k) + \psi_0) \\ L(k)\sin(\tilde{\psi}(k) + \psi_0) \end{bmatrix} + \mathbf{d}(k) \quad (7)$$

Again, the coordinates of the roadside profile vector $\mathbf{r}(k)$ are assumed to be aligned with the road coordinate system. The first coordinate is assumed to be the tangential component, parallel to $\hat{\mathbf{u}}_t$; the second is the normal component, parallel to $\hat{\mathbf{u}}_n$.

The measured profile $\mathbf{r}(k)$ can be compared to a reference profile $\mathbf{R}(k; \hat{\mathbf{x}}, \hat{\psi})$. The reference profile is generated from the estimated absolute position $\hat{\mathbf{x}}$ and estimated absolute heading $\hat{\psi}$. The best estimates are those for which the reference profile \mathbf{R} matches the measured profile \mathbf{r} most closely.

In this thesis, the reference profile \mathbf{R} is constructed from a map database using ray tracing methods. The map database is assumed to comprise the coordinates of all building corners along the right side of the road. The database is assumed to list

building corners in terms of absolute coordinates (latitude and longitude); however, it is assumed trivial to convert into a local-road coordinate system (with components in the $\hat{\mathbf{u}}_t$ and $\hat{\mathbf{u}}_n$ directions). The ray-tracing algorithm is used to determine the distance from the lane center to a building's face, interpolating between building corners. The reference profile \mathbf{R} is constructed by taking the estimated position $\hat{\mathbf{x}}$, a laser heading $\tilde{\psi} + \psi_0$, and uses the ray tracing function to find an expected building face position at the current time step. The displacement $\mathbf{d}(k)$ is added to $\hat{\mathbf{x}}$ to infer the past trajectory of the vehicle. In this manner, the ray tracing function \mathbf{f} can be used to generate the reference profile over a window consisting of the current and K previous data points

$$\mathbf{R}(k; \hat{\mathbf{x}}, \tilde{\psi}) = \mathbf{f}(\mathbf{d}(k) + \hat{\mathbf{x}}, \tilde{\psi}(k) + \psi_0) \quad (8)$$

The two profiles can be compared with a cost function J , with the lowest cost function indicating the best fit. In this sense, we have defined an optimization problem with the following form.

$$\begin{aligned} \min_{\hat{\mathbf{x}}} \quad & J(\mathbf{R}(k; \hat{\mathbf{x}}) - \mathbf{r}(k)) \\ \forall k \in & [k^* - K, k^*] \end{aligned} \quad (9)$$

s. t.:

$$|\hat{\mathbf{x}}_t(k) - [\hat{\mathbf{x}}_t(k-1) + \mathbf{v}_t(k-1)\Delta t]| < B_t$$

$$|\hat{\mathbf{x}}_n| < B_n$$

In this optimization problem, the bounds B_t and B_n refer to the half-widths of the search area (in directions tangential to and normal to the road path, respectively). The tangential search area is centered around an initial best guess, which is calculated

from the estimate in the previous step and a velocity term to account for the car's movement between time steps. In this thesis, the variable B_n is set to cover the full lane width, regardless of the car's previous lateral position; wider search areas might be considered for roads with multiple lanes in the same direction.

It should be noted that the cost function in equation (9) is a function of a time-varying vector defined over times from k^*-K to k^* . An example cost function, the quadratic cost function, for a vector \mathbf{a} of dimension M , evaluated over time steps from k^*-K to k^* , has the following form.

$$J_Q = \sum_{k=k^*-K}^{k^*} \sum_{m=1}^M a_m^2(k) \quad (10)$$

Section 3: Experimental System

We tested the viability of our concept with a road test. Our experiment allowed us to assess the accuracy of RAFL using physical equipment. The experiment also allowed us to see how well our satellite-derived map corresponded to features seen from the ground. Although we did not have access to ground truth substantially more accurate than our results, we were still able to roughly determine our positioning accuracy.

We collected data using a sensor package consisting of an Opti-logic RS100 Laser Rangefinder and a Garmin GPS 18 receiver (shown in Figure 6). The GPS receiver provided both positions and velocities. GPS velocities were used for dead reckoning. GPS positions were not used in our algorithm, so we could compare our results to the GPS positions to estimate the accuracy of our method. Although wheel odometry, an inertial measurement unit, or gyros might have provided more reliable dead-reckoning measurements in an urban environment, GPS velocity data were selected as the DRS

measurements for early-stage testing to simplify system integration. Both the GPS receiver and the laser rangefinder were connected to a laptop computer in the car. Their outputs were synchronized with timestamps. Data were stored and the RAFL algorithm was tested in post-processing.

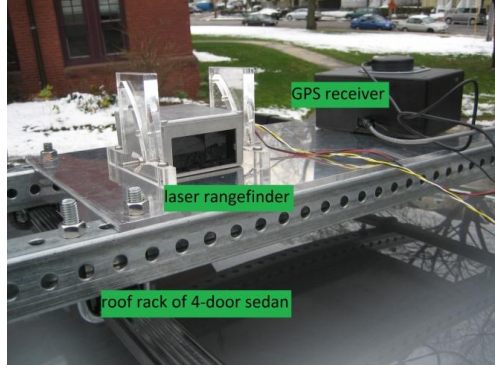


Figure 6: Experimental setup

The laser rangefinder was oriented at a right angle to the vehicle's fore-aft axis; in our coordinate system this corresponds to $\psi_0 = 270^\circ$. As a simplification, we analyzed only data for which the road was straight, allowing approximation of the vehicle's heading ψ as equivalent to the road direction ψ_t . Thus, in analyzing our experimental data, $\tilde{\psi}$ was zero, as specified by equation (6). With these assumptions, equation (7) became

$$\mathbf{r}(k) = \begin{bmatrix} 0 \\ -L(k) \end{bmatrix} + \mathbf{d}(k). \quad (11)$$

In other words, the laser measurement was assumed parallel to the $\hat{\mathbf{u}}_n$ direction. Moreover, the lookup function in equation (8) was dependent only on the car's position and not its orientation.

$$\mathbf{R}(k; \hat{\mathbf{x}}) = \mathbf{f}(\mathbf{d}(k) + \hat{\mathbf{x}}) \quad (12)$$

The reference map that provided the basis for the lookup function was created using the satellite mode of Google Maps. We were able to determine the latitudes and longitudes of building corners to pixel-level precision (11cm), and connect the corners to generate building faces. We then transferred the latitude and longitude locations of each building corner to the local-road coordinate system.

Although we processed all of our data well after collecting it, there is no reason it could not be analyzed in real time, as it would need to be for the method to be utilized. There is room to optimize the code to make it faster, and the reference map could either be stored on an onboard computer, or collected on an as-needed basis from a centralized location with a wireless communication device. The methods for this are beyond the scope of this thesis, but we are confident they are not significant barriers to implementation.

Section 4: Experiment Results and Analysis

This section describes data collected in experimental trials conducted in Medford, MA in the vicinity of the Tufts University campus. Analyses of those data sets are also discussed. Ten trials were conducted, each over a 1.25-mile stretch of urban roads. In our analysis we will focus on a relatively small subset of this data, a swatch of 0.5-mile of road over which the vehicle was traveling in nearly a straight line, according to videos taken during the experiment. Both our method and the GPS give position estimates in line with our expectations. Figure 7 shows the positions from one of these

ten trials overlaid on a satellite map of the road. The estimates in the lateral direction are very close. The longitudinal coordinates agree less well.

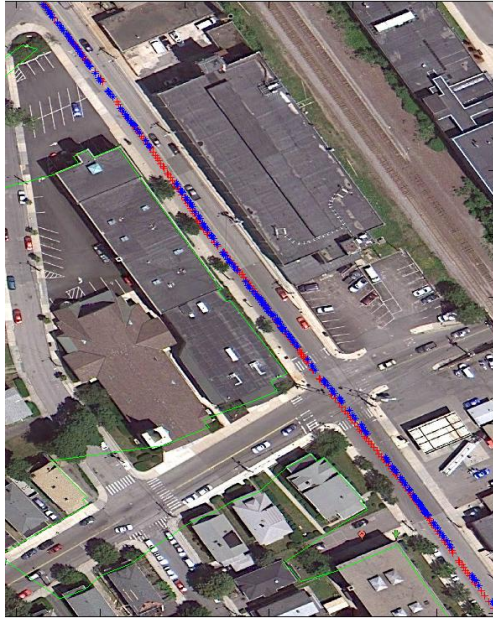


Figure 7: GPS (red) and RAFL (blue) position estimates. The green line is the reference map.

Figure 8 shows the estimates of lateral position x_n generated by both the algorithm and by the code phase of the GPS receiver. It is evident that both the GPS and the algorithm show similar movements in lateral position at the same time. Both show a rightward (negative) movement of the car between 85 and 90 seconds, followed by a relatively sharp leftward reversal until about 95 seconds. Further confirmation that our method is effective comes from the video. North of the intersection, there are no cars parked on the side of the road. South (the direction the car is traveling) of the intersection, there are cars parked along the curb, and at about 88 seconds our vehicle moved left to allow more clearance between them and the side of our car.

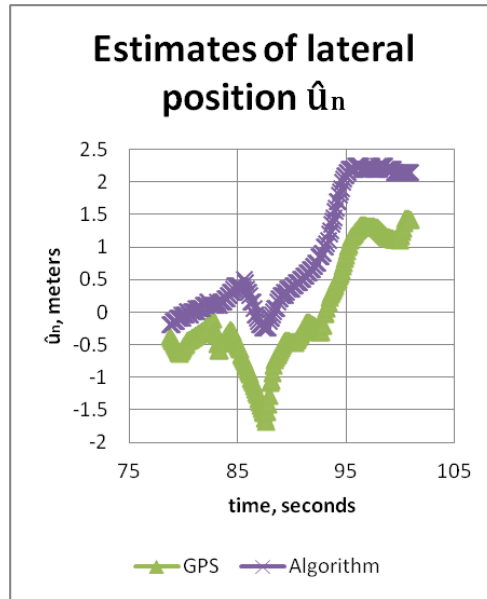


Figure 8: Estimates of lateral position

Since the true position of the vehicle at any point during the trial is unknown, it is not possible to determine the accuracy of our method through this data alone. The offset between the GPS readings and our algorithm’s lateral measurements stays between 0.2 and 1.4 meters for the duration of the trial, even as the values range between -2 and +2.5 meters from our assigned center line, as shown in Figure 9. The average offset was 84 cm, and the standard deviation of the offset values was 12 cm. These preliminary results suggest that our method may be capable of estimating lateral position with meter-level accuracy. More research will be needed to determine whether these trials on this stretch of road represent an unusually favorable situation, and to compare our algorithm’s estimates to actual positions.

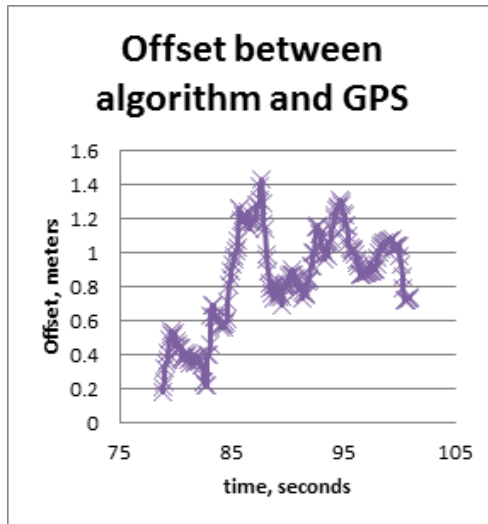


Figure 9: Lateral offsets

The levels of precision of our instruments and map place a limit on the resolution of our method. The reference map resolution is 11 cm per pixel, and latitudes and longitudes are given by Google Maps out to 6 decimal places, which is also 11 cm. If the colors of a building's roof are similar to the colors on the ground nearby, it is often difficult to resolve the building edge to within a single pixel. Registration errors, e.g. how well the coordinates given by the map tool match the feature's true absolute coordinates, are unknown. Moreover, an aerial view is not exactly the same as a slice taken at the exact altitude of our laser rangefinder, e.g. if the roof overhangs the building. The aerial map limits our resolution to 20-30 cm at best, and unknown registration errors limit our method's accuracy as well.

We have demonstrated a method for determining vehicle position using a laser rangefinder fixed to the vehicle. Along with a reference map and a dead reckoning sensor, we have estimate our vehicle's position at each point in time in several trials over one stretch of road. By comparing these estimates to the GPS readings during this

time, and to a video we took, we have confidence in the RAFL concept, and it seems likely that our method has achieved meter-level accuracy, at least for a representative roadway lined with large buildings.

Although the algorithm described in this chapter performs well in the presence of buildings with flat façades close to the road, the algorithm's performance is degraded for more complicated roadsides characterized by non-permanent obstacles (e.g. tall roadside vehicles), irregularly shaped objects (e.g. statues), and non-uniform façades (e.g. houses with steps and porches). Algorithm accuracy also suffers due to occasional large noise events (ranging "spikes"). In order to maintain the desired meter-level accuracy under a wider range of conditions, the next chapter presents ideas that adapt the basic RAFL concept to better manage diverse roadside environments and rare, impulsive errors.

Chapter 3: Concepts for improving our method

Introduction

Although the RAFL method can produce a quality estimate of position in ideal conditions, its accuracy can be compromised by noise in the laser rangefinder signal, obstacles in the path of the beam, irregularly shaped targets, and discrepancies between the aerial map and the way objects appear from the road. The accuracy of our method in these non-ideal conditions can be improved by 1) increasing the accuracy of the reference map, 2) filtering the laser rangefinder input, and 3) replacing the cost function from equation (10) with a function better suited to RAFL. Each of these three concepts has several elements, which are summarized in Table 1. The first three sections in this chapter will introduce concepts for making these improvements, and will also evaluate their effectiveness on one short (400-meter) straight stretch of road. A fourth section shows the effects of combining these improvements.

Table 1: Improvements to our method presented in Chapter 2

Section 1: Map Improvements	Section 2: Filtering the sensor input	Section 3: Improving the cost function
Adjusting the map using ground-based measurements	Median filtering of the laser input	Replacing the quadratic cost function with a cost function less affected by outliers
Adding confidence to reference map segments	Shape matching algorithm	Asymmetrical adjustments
	Obstacle detection algorithm	

Concept 1: Map improvements

This baseline RAFL method works very well for flat, featureless walls parallel to, and close to, the road, because satellite imagery provides an excellent model for roadside features. This section investigates approaches that work more reliably in cases when satellite imagery is less useful, as in the case of the facades of houses, which are rarely completely planar. Gutters, eaves, porches, steps, oriels, and other features may make modeling the roadside profile (at about 1.5 m height) challenging. Irregularly shaped objects, and small to medium-sized objects, like street signs and telephone poles, can also cause problems with a map based on an aerial view. Furthermore, the distinction between ground and a non-ground feature can be blurry when both are the same color, often a dark gray. Ultimately, inaccurate maps lead to the wrong sample swath being selected, and an inaccurate estimate.

The proposed concept for improving the roadside profile model consists of two steps. Firstly, ground-based data collection can be used to adjust the segment endpoints. The first subsection will discuss ways to revise the map from either video or from laser ranges taken from a vehicle. It will also show the effects of the revisions to the map. Secondly, the map's utility can be increased by adding information not given by the map coordinates. Each segment could have information about the type of object it represents, and the likelihood that the object will show up accurately during a laser rangefinder measurement. The second subsection will describe the ways these likelihood values are assigned, and the impact of adding these likelihood values to the algorithm.

Concept 1.1: Adjusting the map using ground-based measurements

This subsection investigates the possibility of improving the accuracy of the map with data collected from a camera and a laser rangefinder fixed to a ground-based vehicle. The first attempt at an accurate reference map used only aerial photographs available from Google Maps. It is possible that an aerial photograph does not capture the relevant details, or that a mapmaker (human or computer) could misinterpret the photo. Revision with ground data could help make the map more accurate and useful. This subsection proposes a change to the map based on visual observations and laser measurements, and describes the change's effect on our algorithm's accuracy.

A map created from laser rangefinder measurements on the ground has the advantage of being constructed in the same way as the sample swaths to which it is being compared. The camera and rangefinder used to create the map have the same line of sight as the laser rangefinder used in the RAFL process. However, this method is limited by the accuracy of all the equipment in use. Uncertainties in the vehicle position and rangefinder mounting make accurate and absolute locating of map points difficult. For this reason, we used the ground-based data only for checking for errors in and making modifications to the aerial map.

Figure 10 shows the motivation for editing the map with our laser rangefinder data.



Figure 10: Laser rangefinder reconstruction versus satellite map

The smooth green line shows the manual tracing of the buildings in the aerial view. The red, cyan, and blue lines show the buildings as traced by the laser rangefinder. We assumed that the laser rangefinder was mounted at the exact location of the GPS unit we mounted on the car, and that it always pointed in the direction perpendicular to the travel lane. As can be seen, all three laser-constructed maps agree with each other better than they agree with the aerial map. There are several possible explanations for this.

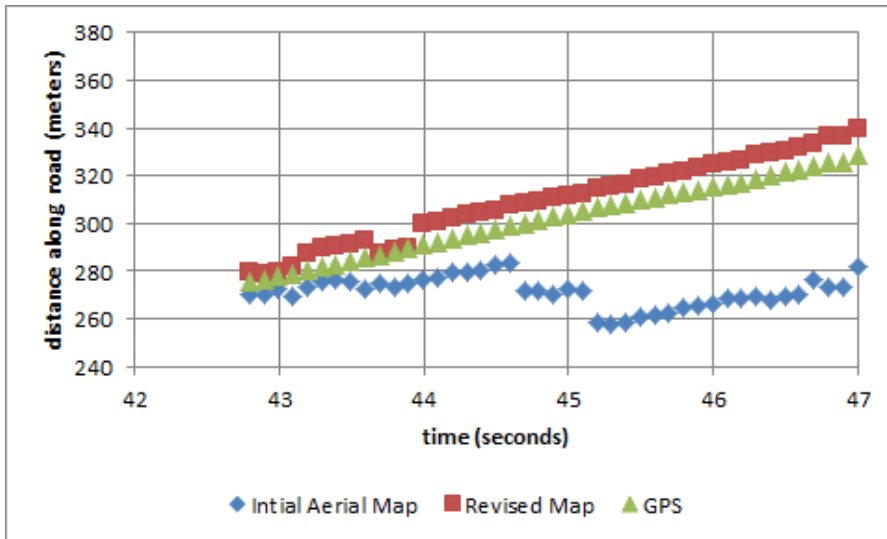
It is possible that the aerial map inadequately depicts the features at the proper height, in the direction they would be viewed from a passing vehicle. Irregular shapes especially, and any non-prismatic shapes, will look different from a satellite compared to how they look from the road. A second possibility is the inaccuracy of the positions created by the laser rangefinder. The GPS unit has unknown errors, especially in a crowded area. The heading of the car, the mounting angle of the laser rangefinder, and the reading of the laser rangefinder all limited precision, so the reconstructed scan will not be completely accurate. A third possibility is obstacles. Since these three measurements were all taken within 15 minutes, they will pick up the same temporary obstacles. For example, it is almost certain that they are detecting the blue vans at the top of the image, which happen to be in slightly different positions on the day of the satellite photo and the day of our experiment.

Despite the opportunities for error in reconstructing the path, the laser data suggested a modification was needed. It is reasonable to conjecture that some object under the tree in the center of the image is blocking the laser path, and a personal inspection of the area confirmed that the aerial map was incorrect. The tree near the center of the image is not in the center of an open area. It is actually in the corner of the area, but the fence at the back of the area is partially blocked by the tree, and where not blocked it is too narrow to show up clearly in the satellite image. Moreover, in Figure 10, part of the building roof had been mistaken for similarly-colored gravel. Figure 11 shows the revised map. One can see that the reference map now agrees much better with the red, cyan, and blue laser measurements.



Figure 11: Reference map adjusted after consideration of disparity with rangefinder data

These changes take place between 272 and 302 meters from the Boston Ave and College Ave intersection, the origin of the local coordinate system. The video shows the car passed through this area around 42 seconds, without any sudden stops or reversals. We ran our RAFL algorithm in this area, with a swath length of 30 meters. As can be seen in Figure 12, both maps yield similar results near when the car enters the area. However, as the car travels forward, the 30-meter swath includes more of the affected area, and less of the previous area where both maps agreed. The algorithm based on the initial incorrect map shows increasing irregularities in this region. If the length of the sample swath is increased past 35 meters, the effect diminishes significantly, since the affected area never accounts for the entire swath.



Comment [TF1]: Data Set SetB,
Inverse4Norm(1,0.5,0.5), sw = 30, rw = [-37.5 to
12.5], medfilter(5,200), no truckwatch, no
shapewatch, t=22.79345:0.1: 74.79345.
Map7NC vs. Map8NC

Figure 12: Effect of revising the map with ground data

Concept 1.2: Adding confidence levels to reference map segments

This subsection presents the concept of weighting areas of our reference map differently. Each segment of the map is assigned a weight, so that the cost function becomes a weighted summation. This subsection explains why some parts of the map should be weighted more, how the weighting is determined, how the weights, or confidence levels, can be implemented mathematically, and their impact on the results of the RAFL algorithm.

Adding confidence levels to the map could improve accuracy, by reducing the impact of fuzzy or uncertain points on the cost function, and increasing the impact of well-defined points. Since the aerial images taken from Google Maps are two-dimensional, each x and y location on the map has only one pixel assigned to it. From here the mapmaker extracts only binary information about height: is there a permanent object there at least 1.5 meters high or not? However, the reality is much more

complex: a single x-y location could contain a solid porch from ground level to 80cm high, then a lattice railing from 80 to 190cm, then plants on top of the railing, then empty space between the plants and the roof, and then the roof. If the roof hangs over the edge of the porch, then the porch, railing, plants, and house will all be modeled as a solid block at the front edge of the roof. Since the height of the laser rangefinder is about 1.5 m, it will likely not give the reading expecting from the solid block model.

The three images in Figure 13 through Figure 15 are all taken from along the section of our route where the reference map was created. Each location has a lateral distance assigned to it on the reference map, but the distance the laser rangefinder will actually read is increasingly difficult to determine.



It is not necessary to take pictures of the area from street level to determine which structures have the highest confidence, although that could be helpful. From the aerial map, one can determine which structures are irregular (plants, statues, etc.),

which structures are small or curved (mailboxes, posts, streetlamps), which are somewhat regularly shaped residential buildings, and which are large rectangular institutional buildings, with increasing levels of confidence. For example, in Figure 16 the long school (1), houses on St. Clement's Rd. (2), and the tree/statue (3) can all be distinguished as such.

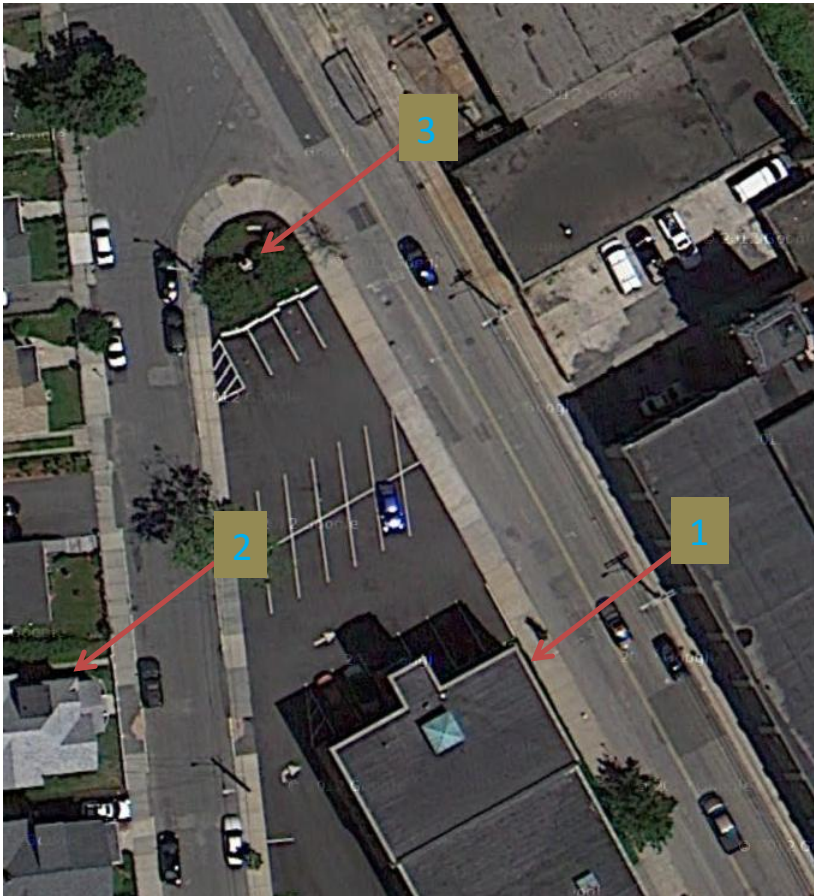


Figure 16: A satellite photo showing features of varying confidence

Choosing a confidence level for a particular segment is a complex decision involving local knowledge, feature identification, and mental recreation of common three-dimensional features from photographs. For this reason, the confidence levels

are assigned by a human. We used a 1-10 scale, with the value of 10 being assigned to the points with highest confidence. Values were assigned according to Table 2.

Table 2: Confidence values assigned to landmarks

Value assigned	Landmark Description
1-2	Building faces nearly perpendicular to the travel direction of the vehicle.
3-4	Irregular objects, such as statues and plants.
5-6	Faces of houses, sheds, and garages, or fences. At a moderate distance from the road, parallel or nearly parallel to the road.
7-8	Faces of houses, sheds, and garages, or fences. Close to the road, parallel or nearly parallel to the road.
9-10	Rectangular buildings very close to the road

These confidence values become scaling factors assigned to each map location. Our previous map was a 4-column table: each line segment on the map was stored as 2 points, each of which has 2 coordinates. The confidence factor for the segment was added as a fifth column in the table. The ray tracing function identifies the segment being chosen, and scales to the confidence factor assigned to that segment.

$$\min_{\hat{\mathbf{x}}} J(\mathbf{R}(k; \hat{\mathbf{x}}) - \mathbf{r}(k)) \quad (9)$$

becomes

$$\min_{\hat{\mathbf{x}}} J\left(C(\mathbf{R}(k; \hat{\mathbf{x}})) \cdot (\mathbf{R}(k; \hat{\mathbf{x}}) - \mathbf{r}(k))\right) \quad (13)$$

Figure 17 shows a few line segments and their confidence factors traced on top of a satellite photo of the area.

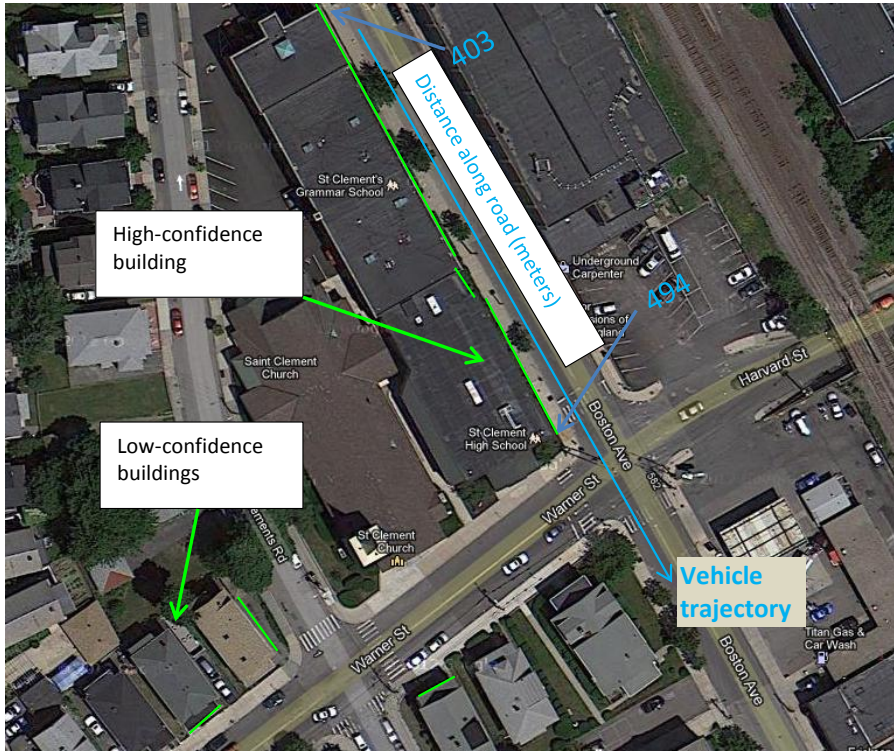
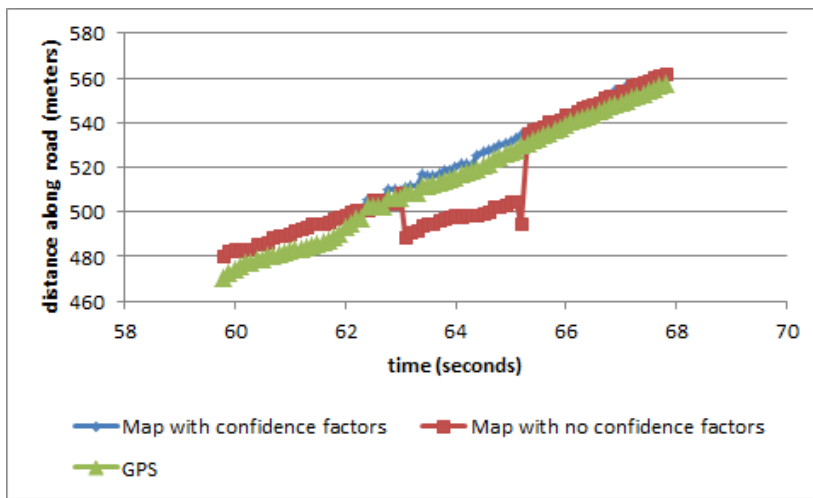


Figure 17: Buildings near 500-meter mark, with their confidence factors (green) and distance markers (blue)

To test the effects of adding these confidence factors, we ran the algorithm twice on the same data set, with the same parameters and conditions. The only change we made was the confidence factors: the first time through, every line segment on the map had a confidence factor of 10, and the second time through, we used our own set of factors, some of which are shown in Figure 17. We used a longer sample swath (80 meters), so that at most points along the route, our reference swath would include both high- and low-confidence points of the reference map. The confidence factors were added after the corrections from ground data described in the previous subsection. However, this area is 200 meters away, so the changes do not interact.



Comment [TF2]: Data Set SetB.
 Erfnorm(b=1,c=2). rw choices are irregularly spaced
 between -40 and 40. See "Thesis Stuff" spreadsheet.
 sw = 80 meters . n= -3:0.25:3. rfilter(5,20), no
 shapematch, no truckwatch.
 Map 8NC (red) vs. Map8C (blue)

Figure 18: Effect of adding a confidence factor to the reference map

Figure 18 shows the results from the two runs of our algorithm from the same subset of data. The green line, representing GPS measurements, is added for reference, and at least roughly tracks the true location and velocity of the vehicle, as determined by video. The blue and red lines represent location estimates from the algorithm.

Before around 500 meters, the algorithm gives the same result both times. However, once the car passes St. Clement High School, the algorithm with confidence factors shows a steady course, and the algorithm without confidence factors shows a reversal in direction. Adding this confidence factor improves the robustness in this situation. Without the confidence factors, the algorithm is equally weighting the part of the reference map including the buildings along Warner St. Attempting to match this data is difficult, since the cross street is highly likely to have obstructions. Since the building faces along Warner Street are perpendicular to our vehicle's path, a small change in vehicle orientation creates a large change in the length of the laser path. With the confidence factors, the algorithm relies more heavily on the data taken while next to the school, which is less error-prone.

Concept 2: Analyzing the sensor input

The second concept for making RAFL performance more robust involves analyzing raw measurement data to reduce sensitivity to impulsive noise and time-varying roadside profiles (e.g. cases when cars may be parked alongside the road). Noise and obstructions can be partially eliminated or ignored, leaving only the more reliable of the laser measurements. The first step in the analysis describes a conventional way of doing this, through median filtering of the laser input. The second and third steps describe additional methods of input analysis. A confidence factor, similar to the map confidence factor, is introduced, and its effect on the performance of the algorithm is evaluated.

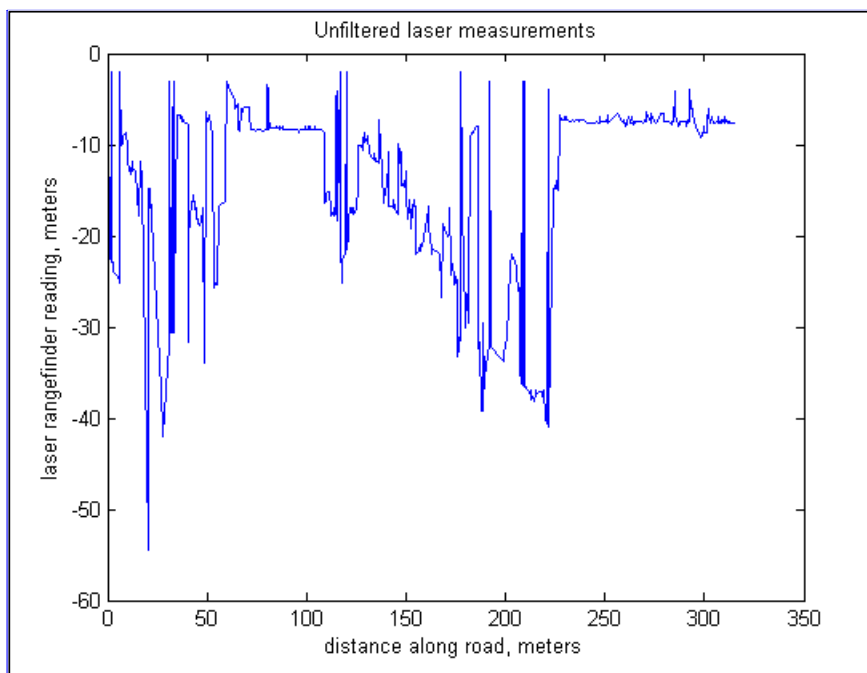
Concept 2.1: Median filtering of the laser input

Removing outliers is the simplest of the three methods of sensor input filtering, since it does not require any input from the DRS, but can work with the input from the laser rangefinder alone. Figure 19 shows a graph of the unfiltered measurements, with several spikes, which represent outliers to be filtered out.

Before applying the median filter, we removed values at the minimum and maximum value of the rangefinder. Maximum or minimum values could result from a default value due to interference or problems in detection, from a view that looks out over the horizon with no obstacle, or from an object with unusual reflective properties such that the laser beam does not get back to the detector. It is also possible that there is an object at the exact distance of the minimum or maximum value, but the frequency with which these values occur, and the lack of surrounding points with approximately

the same values, suggest that the points at the minimum or maximum value can be removed.

After removing these absolute maxima and minima, we also removed some local extrema that are local outliers. Local outliers are within the range of the rangefinder, yet well outside the range of the points around them. The spikes shown in Figure 19 could be noise in the measurements, or measurements of narrow features. The assumption that these are not very narrow features takes into account both the speed of the vehicle and the frequency with which laser rangefinder measurements are taken.



Comment [TF3]: Made from List2, a simple plot of r vs. se

Figure 19: Unfiltered laser measurements

We chose to use a median filter because the median value is less distorted by outliers than the weighted or unweighted average of surrounding values. Since we are more interested in exploring how filtering applies to our method than in developing

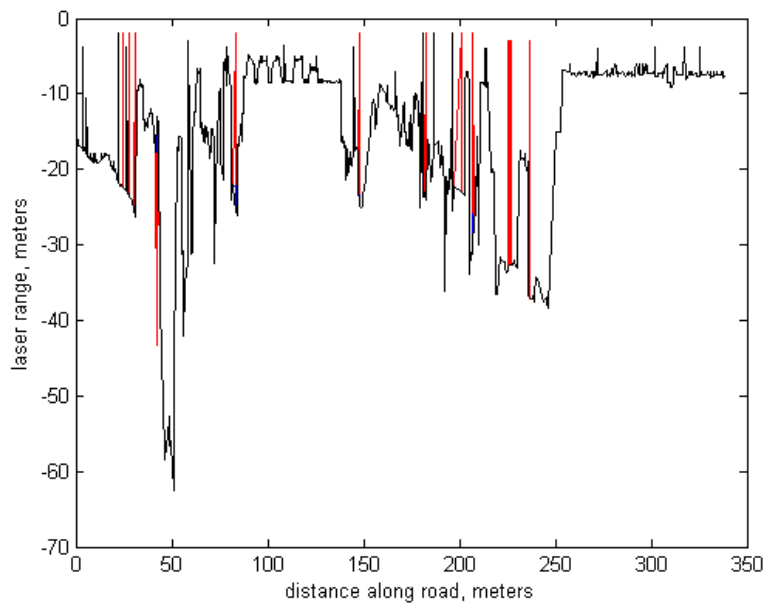
novel filtering techniques, we only tried the median filter. Based on our sampling time and vehicle speed, we used this formula for filtering.

$$Filtered\ L(i) = \begin{cases} (L(i) - M_n(i))^2 < c^2, & L(i) \\ (L(i) - M_n(i))^2 > c^2, & M_n(i) \end{cases} \quad (14)$$

$L(i)$ is the laser measurement at index i . $M_n(i)$ is the median laser measurement of n points centered around $L(i)$, where n is an odd number.

This filter replaces some of the laser measurements with the median measurement of the surrounding points. Where the laser measurements are constant or monotonic, then $L(i) = M_n(i)$. Where $L(i)$ is a local extremum, it will not be equal to $M_n(i)$. If the difference is greater than the cutoff value c , then the laser measurement at that point will be replaced with the median. If the value of c is very high, almost no values will be replaced. If the value of c is close to 0, nearly every local extremum will be replaced.

The width n of the median filter can also be adjusted. Increasing n will eliminate some 2-point or 3-point spikes, but it will not eliminate 1-point spikes if they are too close together. For this application, values of n between 5 and 21 were considered. The next three figures show the effects of filtering of the laser measurements over this range.



Comment [TF4]: Data set ListA. No shape matching or obstacle detection.

Figure 20: This chart shows a modest level of filtering. The filtered laser measurement is shown in black, with spikes removed shown in red. The values of n and c are 5 and 20, respectively.

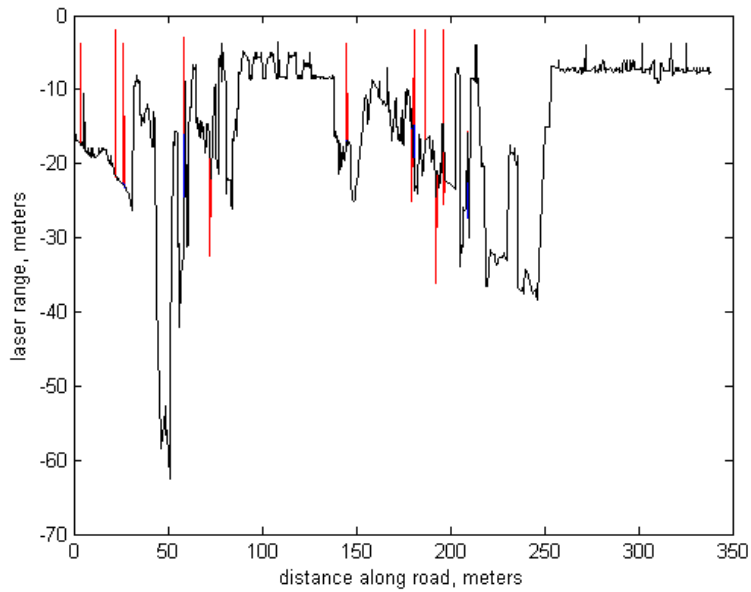


Figure 21: This chart compares filtering with $c = 10$ to filtering with $c = 20$. Shown in red are spikes removed at $c = 10$, but not at $c = 20$. The value of n is 5.

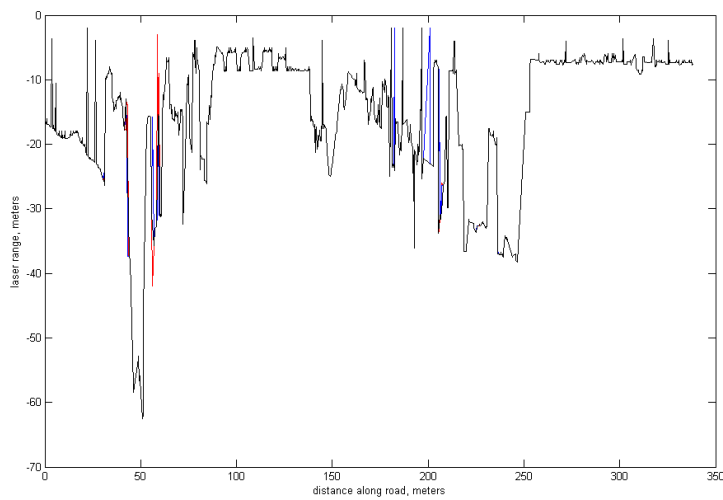
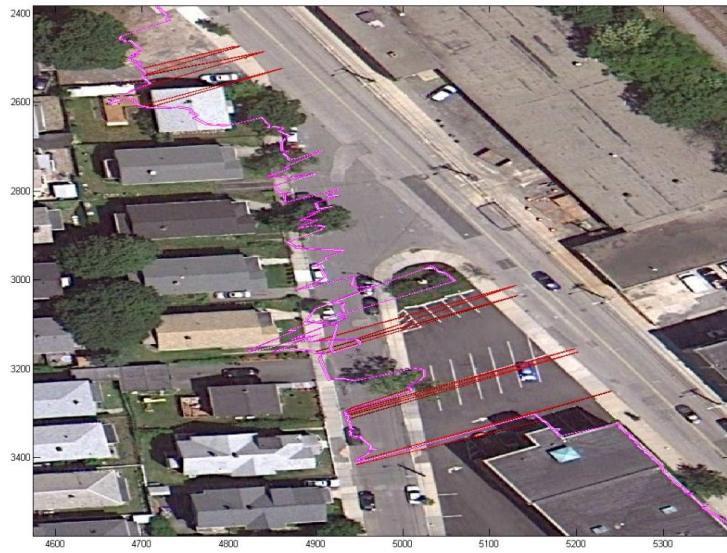


Figure 22: This chart compares a 5-point median filter ($n = 5$) to a 15-point median filter ($n=15$). The 5-point filter removes the blue spikes. The 15-point filter removes the red spikes. Spikes removed in both cases are not shown. Both filters used a value of 20 for c .

Figure 23 shows laser rangefinder values overlaid on a map. The red is unfiltered, and the magenta/white shows the laser measurements filtered using Equation (14) with $n = 9$ and $c^2 = 50$.

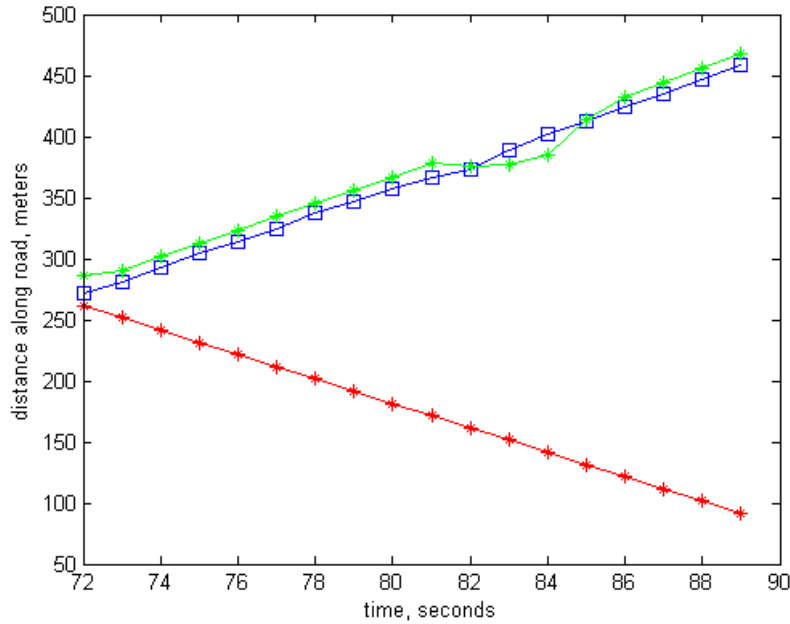


Comment [TF5]: This is a close-up from List2 mapped on top of the BARM. Red lines are r vs. se ; magenta lines are rf vs. se with $rf = \text{medfilter}(r, 9, 50)$

Figure 23: Filtered and unfiltered range values

Filtering the laser rangefinder input before use in the algorithm can sometimes make a dramatic difference in the position estimates. Figure 24 shows the effect of smoothing for $n = 5$ and $c = 5$ over a 17-second stretch of data. Blue is the GPS reading, red is the algorithm estimate from unfiltered laser data, and green is the algorithm estimate from the filtered data. Without filtering, the algorithm never finds a high-quality match for a sample swath, and ends up with completely erroneous result. The filtered data matches the reference map better, so it is able to provide reasonable position estimates with the RAFL algorithm.

For the single data set considered, good results were obtained for a filter width n of 5 and for a filter threshold of 5. This is among the most aggressive filters we have considered, and is needed due to the large amount of noise in this particular data set



Comment [TF6]: Data set ListA, Map 8NC, sw=70, rw=40 (-10 to +30), no shape matching, no obstacle detection, Inverse4norm v1.3 [1,1,.5], n=3:0.25:3
No filtering vs. rfilter(5,25)

Figure 24: Results with (green) and without (red) filtering, compared to GPS (blue)

Concept 2.2: Shape matching algorithm

The second method relies on larger grouping of points to determine which correspond to real-world physical objects, and which are associated with sensor noise and obstacles. Points that match one of a set of predetermined shapes are given higher confidence, and points that do not match any of the shapes are given lower confidence, regardless of whether these shapes match up with anything in the reference map.

The higher confidence for these points can be incorporated into the formula for the cost function much as the confidence factor for the map. Each laser rangefinder point can have an associated confidence value, $C_L(k)$, where k is the time step.

$$\min_{\hat{\mathbf{x}}} J(\mathbf{R}(k; \hat{\mathbf{x}}) - \mathbf{r}(k)) \quad (9)$$

becomes

$$\min_{\hat{\mathbf{x}}} J(C_L(k) \cdot (\mathbf{R}(k; \hat{\mathbf{x}}) - \mathbf{r}(k))) \quad (15)$$

or, if combined with the confidence factor from the reference map

$$\min_{\hat{\mathbf{x}}} J(C_L(k) \cdot C_M \cdot (\mathbf{R}(k; \hat{\mathbf{x}}) - \mathbf{r}(k))) \quad (16)$$

$C_L(k)$ values can range from 0, at a point that does not match a shape, to 10, at a point that definitely matches a shape. These points will have more impact on the cost function. Low-confidence points will be relatively neglected, and the fit will be judged good or bad based on how well the high-confidence points fit the map. C_L values are assigned to a group of points that match a particular shape thought to be representative. In our reference map created from an aerial view, almost all the target surfaces are straight building edges, so the only shape we matched was a straight line. To calculate whether a group of n points are part of a line, we transformed the coordinate system so that the origin was at the first point, and the x axis passed through the n^{th} point. The deviation from this line could be calculated from the y values. The most general formula for the deviation is given below.

$$D = \frac{\sum |y_i|}{n} \quad (17)$$

The particular equation we chose to use is given in (18).

$$D = \frac{\sum y_i^2}{n^2} \quad (18)$$

Confidence is inversely proportional to D . Based on trial and error, we established this equation empirically.

$$C = \begin{cases} D < 0.004, & 2 \\ D \geq 0.004, & 1 \end{cases} \quad (19)$$

Sets of 5, 10, 20, and 40 points are selected for evaluation with Equation (18).

Figure 25 shows the graph of laser data plotted against distance. The points highlighted in red are parts of lines with $C = 2$. If a point is in at least one line marked $C = 2$, the point's C value is 2; otherwise it is 1. This method is best applied after the laser data has been filtered, since filtering smoothes the data and makes selecting a line more likely.

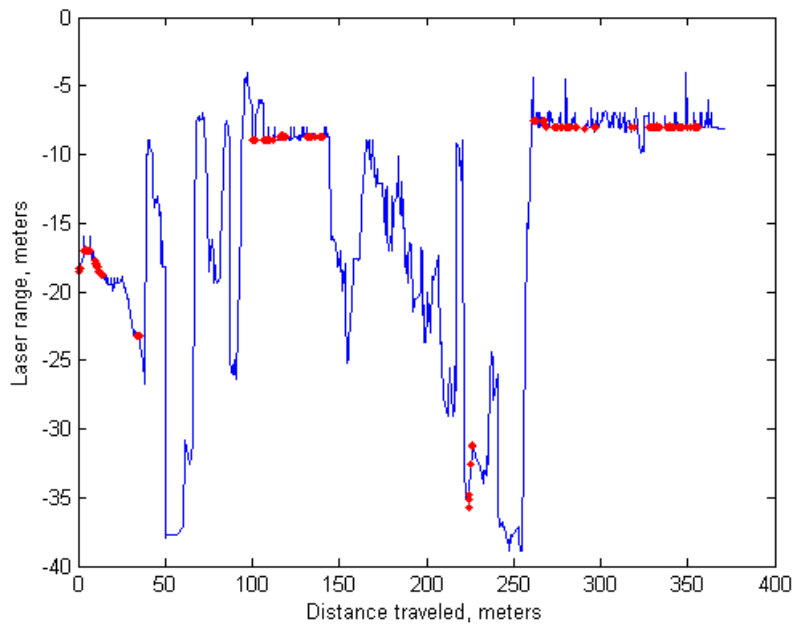


Figure 25: Points identified as linear

This method shows improvement in the overall performance of the algorithm, as can be seen in Figure 26. Although both see a hiccup in the location estimate near

the center of the image, only the estimate with shape matching (green) continues with a reasonable estimate of lateral position, in keeping with video evidence that the car did not switch lanes at this time.



Comment [TF7]: Data set ListA, Reference Map 8NC, Inverse4Norm v1.3, sw=70m, rw=[-10 to +30] from previous estimate, n=-3:0.1:3, rfilter(5,25), no obstacle detection, t=72 to 89 s.

No shapewatch vs. shapewatch 1.0 w/IsLine 1.0

Figure 26: Effects of shape matching (green)

Concept 2.3: Obstacle detection algorithm

Just as the confidence factor can be increased when points show a promising shape, the confidence can be decreased when the shape of the points seem to be in the shape of an obstacle. For example, Figure 27 shows five features that are likely vehicles parked alongside a road in front of a building. Video taken simultaneously with the velocity and range data confirms that there were in fact five vans parked along the side of the road in about this location, and that they were tall enough to block the laser rangefinder.

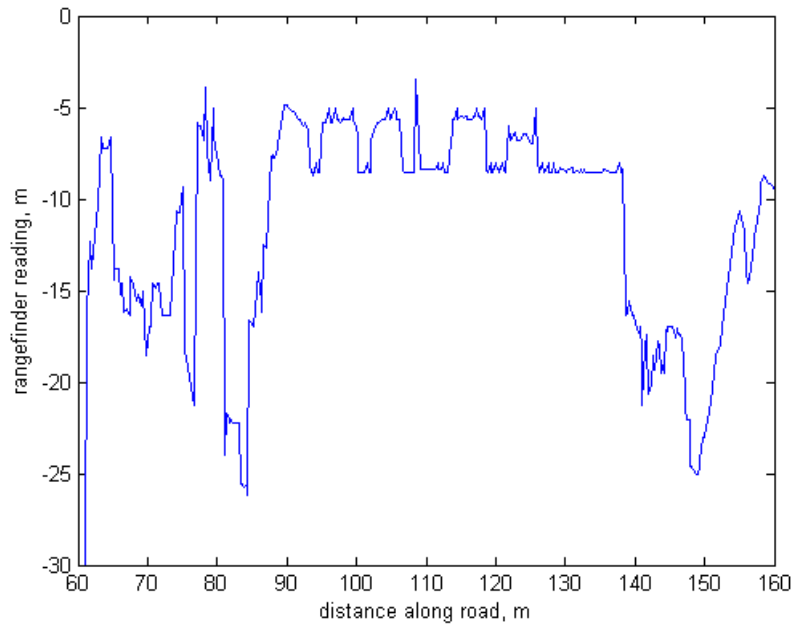


Figure 27: Laser rangefinder readings

Many specialized algorithms could be written for various objects and locations; we focused on cars parked along the curb. We defined a likely obstruction as a set of points meeting these criteria:

- 1) Each obstruction begins at a point between 2 and 7 meters from the laser rangefinder, and this point is at least 1 meter closer to the rangefinder than the point preceding it.
- 2) Each obstruction ends at a point between 2 and 7 meters from the laser rangefinder, and this point is at least 1 meter closer to the rangefinder than the point immediately following it.
- 3) The distance between these two points is between 2 and 7 meters.
- 4) The points are all roughly in a line parallel to the road. That is, all the points have the same laser rangefinder reading.

This can be expressed mathematically in these equations. Each set of points (s_i, r_i) begins with $i = a$ and end with $i = b$. Recall that all r values are negative. The number of points is $n = b - a + 1$. The set is an obstruction if and only if all of these criteria are met.

$$r_a - r_{a-1} > 1 \quad (20)$$

$$-7 < r_a < -2 \quad (21)$$

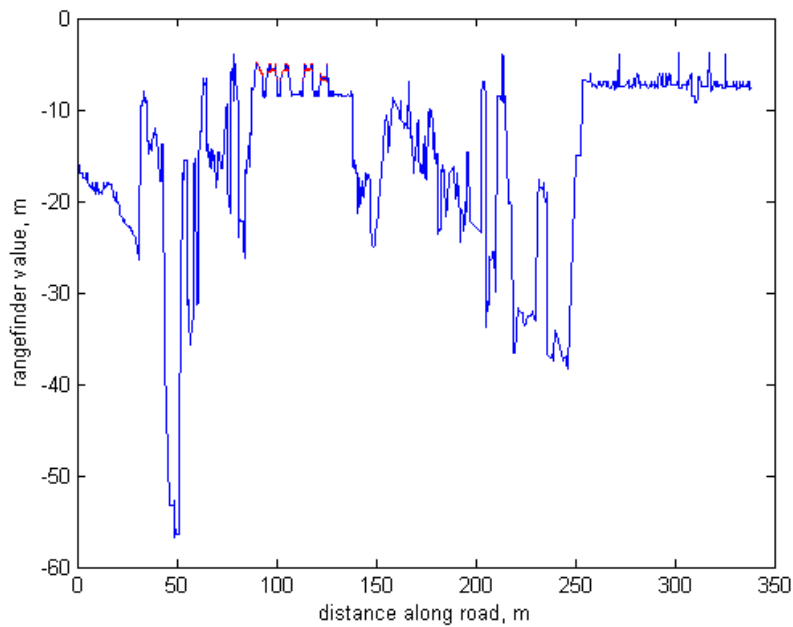
$$r_b - r_{b+1} > 1 \quad (22)$$

$$-7 < r_b < -2 \quad (23)$$

$$var(\mathbf{r}) = \sum_{i=a}^b \frac{(r_i - \bar{r})^2}{n} < 0.25 \quad (24)$$

$$2 < s_b - s_a < 7 \quad (25)$$

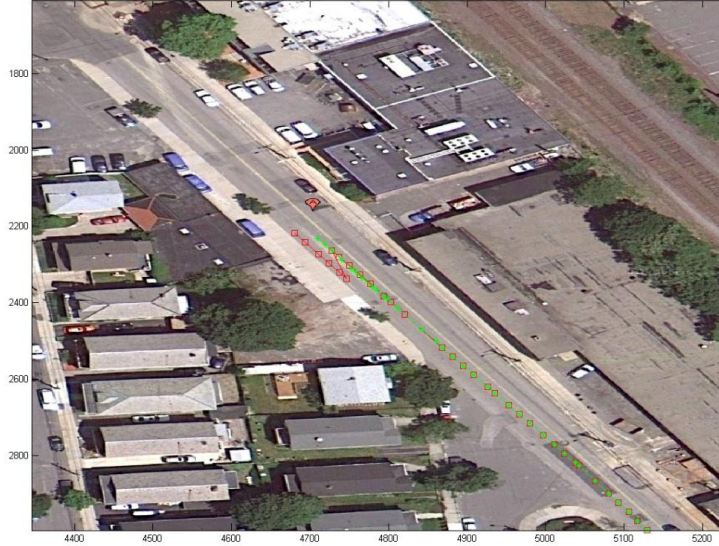
These criteria yield five obstacles in our data set, as shown in Figure 28.



Comment [TF8]: Data set ListA.
Medfilter(r,5,15) applied. Truckwatch 3.0.

Figure 28: Likely obstacles detected by the rangefinder are identified in red.

Moderate to aggressive smoothing of the laser rangefinder data can augment the usefulness of the obstacle detection algorithm. The effect of the algorithm is shown in Figure 29. Immediately after passing the trucks, the algorithm fails unless the obstacles are detected and removed from the sample swath. Video evidence corroborates the assumption that the car did not make the Z-shaped motion described by the set of red squares.



Comment [TF9]: Data set ListB, Map 8C, sw=70, rw=[-8 to +24], n=-3:0.1:3, rfilter(5,25), no shape matching, inverse4Norm v1.3 [1,1,0.5], t=41:0.25:50.

No truckwatch vs. truckwatch 3.0

Figure 29: Position estimates with (green) and without (red) obstacle detection

Concept 3: Improving the comparison function

Changes to the function that compares the sample swath to the reference swaths can also mitigate the problems of sensor noise and obstacles. The quadratic cost function used in Equation (10) is intuitive and computationally efficient, but is not optimized for the RAFL situation. Subsection 1 describes the shortcomings of the quadratic cost function, and the characteristics that a cost function optimized for our situation would have. The first subsection briefly explores a few alternative cost functions used in the comparison, that have improved outlier rejection.

Based on our observations of trial data, we realized that the observed laser data often gave smaller measurements than expected at a particular location, but rarely a larger measurement. The second subsection postulates reasons for this asymmetry, and

demonstrates an adjustment to the cost function that takes the asymmetry into account.

Symmetric cost functions for error minimization

Mathematically, the RAFL algorithm solves the optimization problem first presented in equation (9).

$$\begin{aligned}
 \min_{\hat{\mathbf{x}}} \quad & J(\mathbf{R}(k; \hat{\mathbf{x}}) - \mathbf{r}(k)) \\
 \quad & \forall k \in [k^* - K, k^*] \\
 \text{s.t.} \quad & \\
 & |\hat{\mathbf{x}}_t(k) - [\hat{\mathbf{x}}_t(k-1) + \mathbf{v}_t(k-1)\Delta t]| < B_t \\
 & |\hat{\mathbf{x}}_n| < B_n
 \end{aligned} \tag{9}$$

The cost function in the top equation could be one of several functions. The quadratic cost function, presented earlier, is among the simplest:

$$J = \sum_{k=k^*-K}^{k^*} \sum_{m=1}^M a_m^2(k) \tag{10}$$

For a value of M=1 and K=0, the value of the cost function would vary like this:

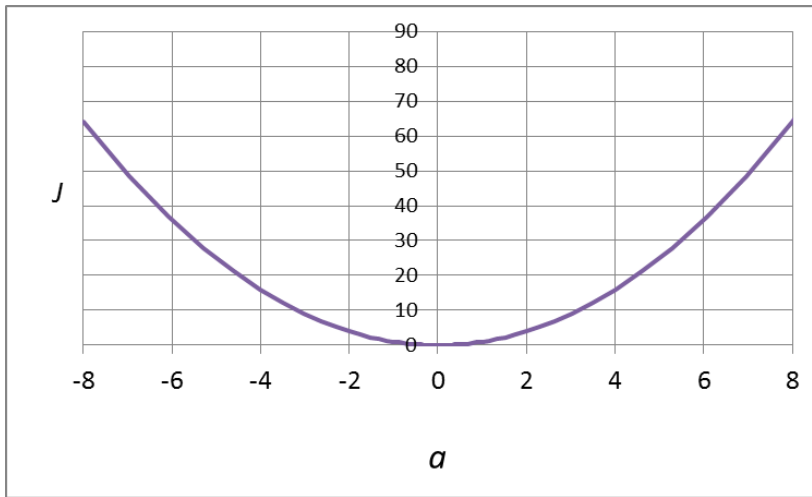


Figure 30: The quadratic cost function for various values of a

The quadratic cost function is easy to implement, but it has its flaws when incorporated as part of our model, particularly for large discrepancies between measured r values and reference R values. The ideal cost function would have these properties:

- Balance the advantages of summations and averages. As mentioned earlier, the laser rangefinder samples at a fixed frequency, but the time length of the collection period varies inversely with the car velocity, as shown in Equation (4). So some sample swaths will have more points than others. A summation would tend to give less neutral values to larger swaths, which could be advantageous since more points bring more certainty. On the other hand, if all the points in the set of swaths are not evenly balanced around 0, larger swaths could be given an automatic advantage or disadvantage. An average would not have this bias.

- Ignore differences between $R(k)-r(k)$ values that are below the accuracy of the laser rangefinder, or smaller than needed for the application. A match of 1 cm should not have a penalty 10 or 100 times higher than a match of 1 mm.
- Give the same penalty to all $R(k)-r(k)$ values above a certain threshold. A point 100 m off is no worse than a point 10 m off; either indicates that laser target is definitely not the target at that map position.
- Give increasing penalties for moderate errors. Moderate errors (a few centimeters to a few meters) may indicate the n-coordinate estimate is wrong, especially if there are equal errors at several points within the swath. Within this range, larger errors should have larger function values.

A better curve might look something like this:

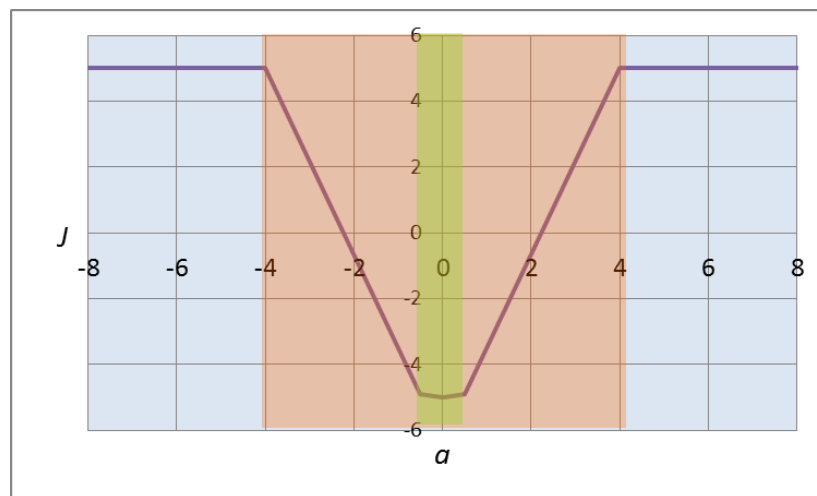


Figure 31: An improved cost function

In the green region, closest to 0, the J values are close to flat. In the mid-range red region, they start increasing, with $a = 2$ giving a value of $J = 0$. In the blue region, there is a penalty associated with a poor fit between the sample swath and the

reference map at a particular location. Within the blue region, the penalty is relatively inelastic with respect to the value of a .

One potential drawback is that the cost function is no longer convex. However, in our implementation we evaluated every possible value of \hat{x}_t and \hat{x}_n (with a finite step size) in the search space. Since this method does not use derivatives or local minima, the impact of non-convexity on our results is negligible. Additionally, the cost function argument a is not necessarily linear with \hat{x}_t , so the cost function may not be convex with respect to the inputs in any case.

There are several mathematical possibilities for producing the relation shown in Figure 31. This subsection presents three of them: 1) the inverse quadratic cost function, 2) the inverse quartic cost function, and 3) the shifted error function cost function. All are presented and discussed as summations; however, they could be turned into averages simply by dividing by the product $(K+1)M$. The tradeoffs between summations and averages are the same for each.

Inverse Quadratic

$$J = \sum_{k=k^*-K}^{k^*} \sum_{m=1}^M \left[-\frac{b_1}{1 + b_2 a_m^2(k)} + b_3 \right] \quad (26)$$

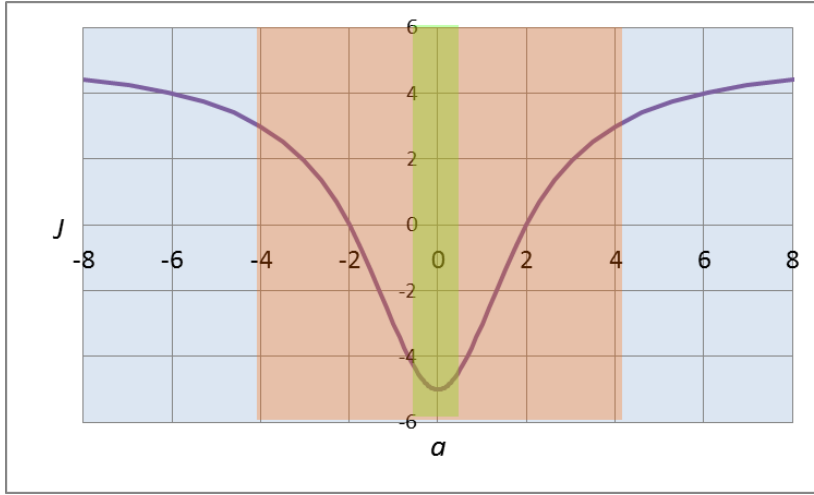


Figure 32: Inverse quadratic plot with $b_1 = 10$, $b_2 = 0.25$, $b_3 = 5$

The inverse quadratic cost function mostly resembles the improved cost function, and is smoother. However, it is not as flat in the green region, and the slope varies more in the red region. The 1 in the denominator prevents the graph from approaching negative infinity at $a = 0$. The coefficient b_1 controls the distance between the maximum value (as a approaches infinity or negative infinity) and the minimum value (at $a = 0$). The coefficient b_2 controls how stretched the graph is along the x-axis: the larger b_2 is, the steeper and narrower the graph. The coefficient b_3 shifts the entire graph up and down; this can determine whether the cost function generally produces negative or positive values. The inverse quartic cost function, presented next, is very similar.

Inverse Quartic

$$J = \sum_{k=k^*-K}^{k^*} \sum_{m=1}^M \left[-\frac{b_1}{1 + b_2 a_m^4(k)} + b_3 \right] \quad (27)$$

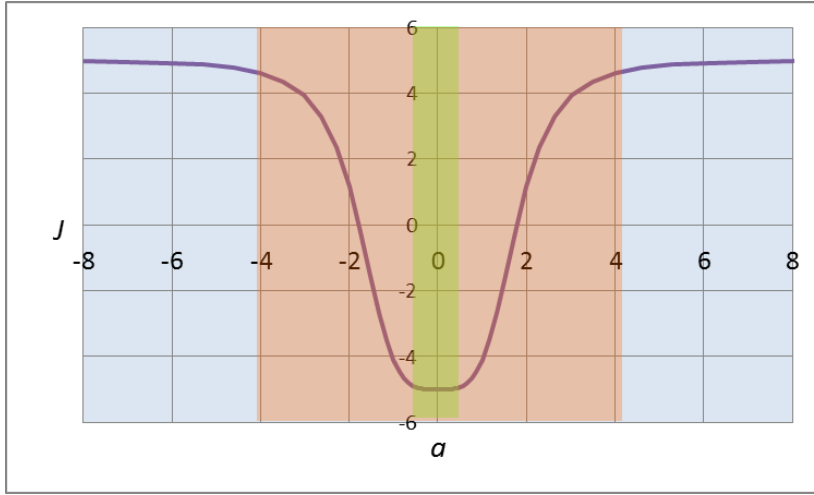


Figure 33: Inverse quartic plot with $b_1 = 10$, $b_2 = 0.1$, $b_3 = 5$

This cost function is like the inverse quadratic function, but the plot is flatter near the bottom, and has steeper sides, more like the improved function shown in Figure 31. The three coefficients play the same roles as for the inverse quadratic function.

Shifted error function

The error function is defined as:

$$\text{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt \quad (28)$$

It can be incorporated into a cost function as shown in equation (29).

The absolute value of a is taken to make the graph symmetric.

$$J = \sum_{k=k^*-K}^{k^*} \sum_{m=1}^M [b_1 \cdot \text{erf}(b_2[|a_m(k)| - b_3]) + b_4] \quad (29)$$

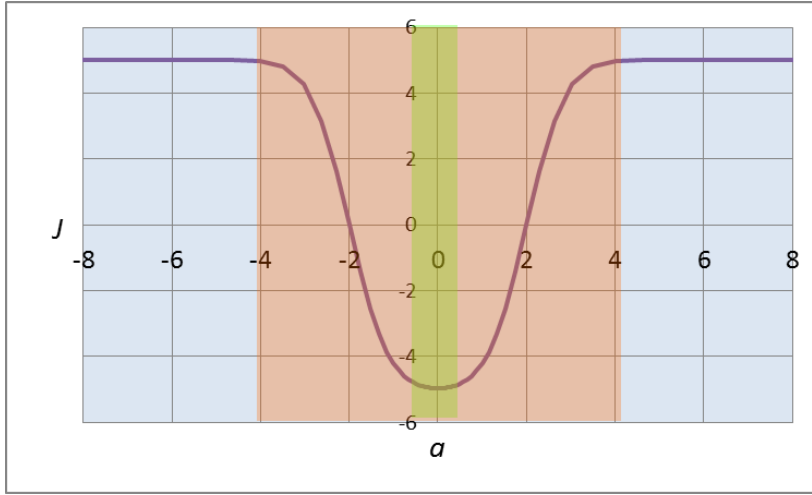


Figure 34: Shifted error function plot with $b_1 = 5$, $b_2 = 1$, $b_3 = 2$, $b_4 = 0$

The graph is very similar to the inverse quartic graph, but the four coefficients in equation (29) allow for more precise control of the function. Rather than one coefficient controlling both the width of the central (green) region and the steepness of the curve in the red region, b_3 and b_2 respectively control these features. The coefficient b_3 adjusts the horizontal spacing of the curve, with a larger b_3 producing a wider flat region at the bottom of the curve and x-intercepts more widely spaced. The slope in the red region is proportional to b_2 . The coefficient b_1 again controls the distance between the minimum and the maxima, and b_4 shifts the curve between the negative and positive regions.

In order to decide which symmetric cost function is best suited to the RAFL application (at least for the data set considered here), we devised an objective measure of cost function performance. The algorithm using the best function would ideally calculate a much lower function value for the correct location than for any other location. If there are a large number of suggested locations, C is the cost function value at the correct location, m is the mean of all the function values for all the locations, and

s is the standard deviation of all the values, a Z score for the value at the correct location is given in Equation (30). The lower the Z score, the better the function sets apart the correct location from the others.

$$Z = \frac{C - m}{\sigma} \quad (30)$$

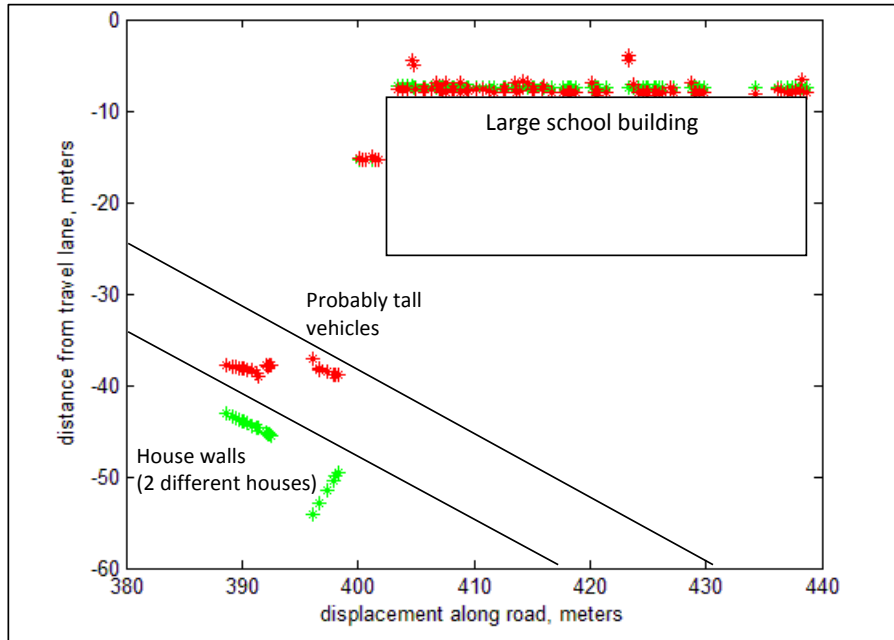


Figure 35: Sample measurements (red) and reference map (green) at 55 seconds

Figure 35 shows the laser measurement at 55 seconds, and all the measurements for 50 meters in back of that. They are aligned using the DRS, with the location of the last measurement set at a longitudinal position of 438.7 meters, and a lateral position of 0.71 meter. The green points in the plot show the reference map points found by the ray tracking function. The point (438.7, 0.71) was chosen to maximize the overlap between the two plots. We are certain that this is the location

the algorithm should give, so we can use the cost function value at (438.7, 0.71) as C in equation (30).

We ran several trials of the algorithm with this data. Each time we used the exact same swath, filtering, reference map, and reference map confidence values, but a different cost function, or different coefficients. The table below summarizes the best Z scores we obtained with each cost function. For each function, we used several different sets of coefficients, and the lowest value of any of them is shown here.

Table 3: Z score values for different cost functions

Cost function used	Lowest Z -score obtained at $t=55$ seconds
Quadratic	Could not correctly locate vehicle
Inverse quadratic	-1.83
Inverse quartic	-2.01
Shifted error function	-2.17

As expected, the shifted error function, which has the most flexibility, is most able to differentiate the correct solution from an average solution. However, its advantage over the inverse quadratic and inverse quartic functions is small. More research will be needed to verify the cost function quality on a larger data set.

Asymmetric cost functions for error minimization

All of the cost functions presented so far use the absolute value or square of the difference between the reference swath range and the sample swath range. This subsection explores whether a function that includes the sign of the difference can make any improvement in predicting vehicle location.

From observing Figure 35 and other superimposed comparisons of the reference swath and the sample swath, one can see that the laser measurements often fall short of the range values predicted by the reference map. In both cases, tall vehicles to the side of a road or in a driveway probably account for the discrepancy. If the building on the reference map is far from the RAFL vehicle, then the discrepancy can be large. These large discrepancies can occur even when the hypothesized position is correct.

On the other hand, an obstruction will never result in a higher than predicted laser reading. These are caused by 1) reference map mistakes, such as mistaking a roof for a parking lot, as discussed before; 2) aerial map feature that are not at the height of the laser rangefinder, but either above or below; 3) laser rangefinder misalignment, 4) multipath errors in the laser beam, 5) laser rangefinder inaccuracy, or 6) a laser beam that is completely absorbed or deflected. In case (6) the equipment returns a maximum value, and the data is discarded in our filtering step.

Our hypothesis is that laser rangefinder inaccuracy is within a meter, and the other 4 causes are much rarer than obstructions. Therefore, a swath that shows laser measurements much smaller than reference map locations is more likely the correct one than a swath that shows laser ranges exceeding reference map values by a large amount.

Any symmetric cost function has the same value for both of these situations, meaning they are equally likely to be selected as the best estimate. Increasing the penalty for a range past the map target, relative to the penalty for a range short of the map target, could result in better estimates. We propose an asymmetrical correction to

the cost function, so that it varies with the difference between reference range and sample range as shown in Figure 36.

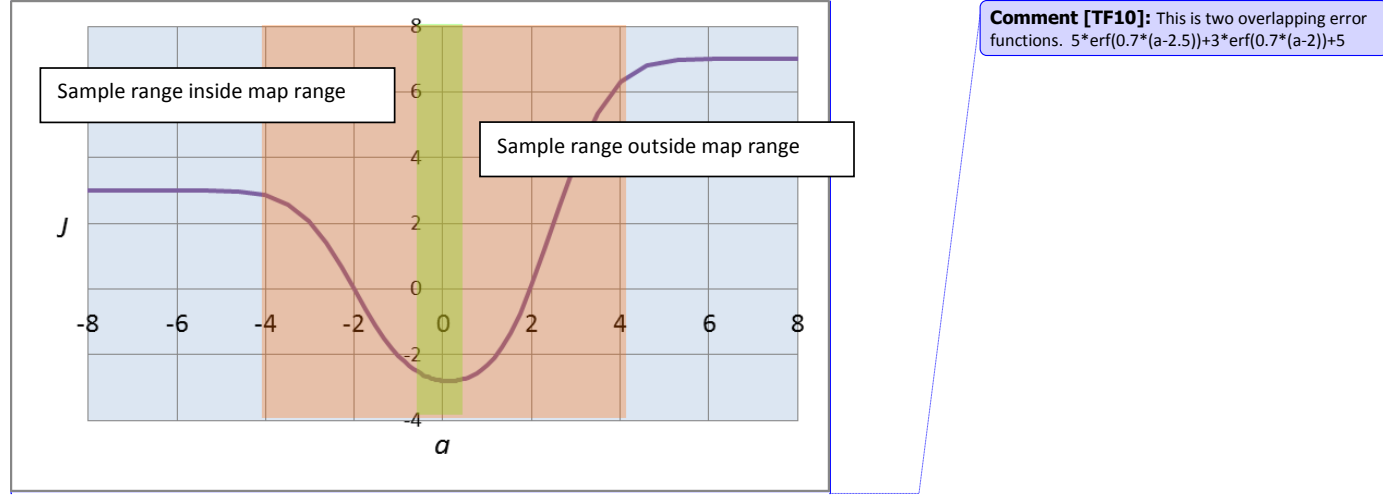


Figure 36: Asymmetrical error function cost function

One possible formula uses two overlapping sigmoidal curves. Again, the error function with multiple coefficients provides flexibility. This equation is derived from Equation (29), with \mathbf{a} indicating the vector between a map point and the point at the laser rangefinder target. For asymmetry, the equation is split into two halves, one using $a_m(j)$ as an argument, and one using $-a_m(j)$.

$$J = \sum_{k=k^*-K}^{k^*} \sum_{m=1}^M f(a_m(k)) \quad (31)$$

where

$$f(a) = [b_1 \cdot \text{erf}(b_2[a - b_3]) + b_4 \cdot \text{erf}(b_5[-a - b_6]) + b_7] \quad (32)$$

Recall that r values are all negative, so a positive value of $a_m(j)$ means the reference value of R taken from the map is closer to the vehicle than the measured value of r . Accordingly, positive values of a are given higher penalties.

In order to evaluate the potential benefit of the proposed asymmetric cost function, it is instructive to compare its performance to that of the symmetric cost functions described above (at least for the case of the single trial under consideration)..

To compare the performance of several cost functions, we deliberately ran our algorithm with some non-optimal settings, i.e. minimal filtering of the laser input, no shape matching or obstacle detection, and a small sample swath. This allowed for greater divergence between the best and worst cost functions, since optimal filtering can often compensate for shortcomings in the swath comparison algorithm, and vice versa.

First, we watched the video to verify that the GPS readings, which are completely independent from our estimates, are reasonably accurate. Using the video, we could verify that the GPS matches location along the road within 15 meters, and that the true longitudinal velocity was relatively constant, with no abrupt stops or accelerations, and definitely no backtracking. This gave us two metrics against which to measure our estimates: first, that the positions are within 10-15 meters of the GPS positions, and second, that the longitudinal position estimates keep increasing at a relatively constant rate.

After using the video to establish these performance metrics, we ran the algorithm 3 times, each time with the same set of laser and DRS data, the same filtering, and the same reference map. The only difference between these three trials was the cost function used. Each trial generated a set of position estimates for a series of chosen time steps. In order to compare positions to the GPS data, we plotted the longitudinal positions on the vertical axis against the time steps on the horizontal axis.

As can be seen in Figure 37, the asymmetric error function cost function performs better than the inverse quadratic cost function with no asymmetrical term. Both offer significant improvement over the quadratic cost function.

We repeated these steps many times, each times with a different data set. Most often, all the cost functions performed similarly, except for quadratic cost function, which generally was worse. A minority of data sets showed significant differences between the improved functions, and Figure 37 is representative of that group.

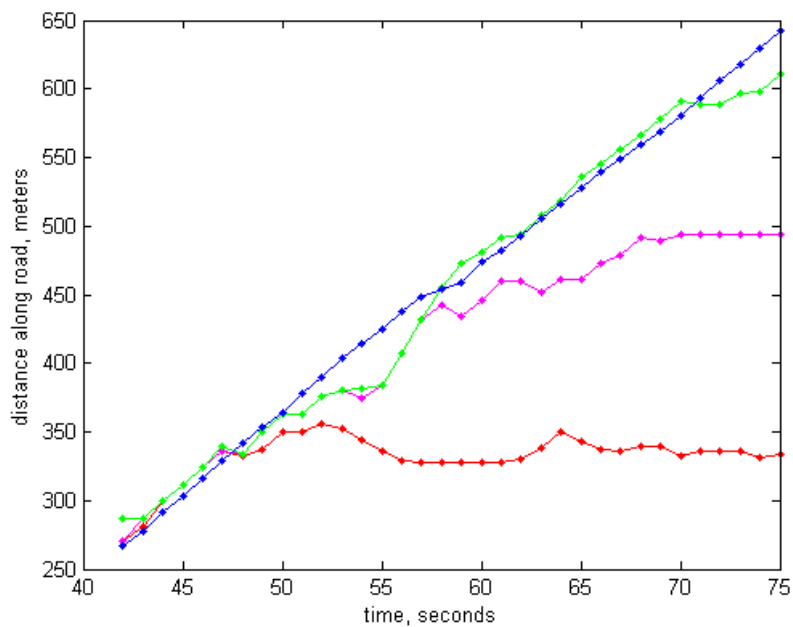


Figure 37: Performance comparison of cost functions. BLUE-GPS. RED-estimates using the quadratic cost function. MAGENTA-estimates using the inverse quadratic cost function. GREEN-estimates using the asymmetrical error function cost function.

Section 4: Effects of combined improvements

We have presented a number of improvements in this chapter, as summarized in Table 4 (a copy of Table 1 presented earlier). The effects of the improvements were highly dependent on the data set used, as well as the other parameters. Sometimes an improvement had a large effect, and sometimes it had no effect. None of the improvements presented in this thesis had a noticeable negative effect in any data set we studied.

Table 4: Improvements to our method presented in Chapter 2

Section 1: Map Improvements	Section 2: Filtering the sensor input	Section 3: Improving the cost function
Adjusting the map using ground-based measurements	Median filtering of the laser input	Replacing the quadratic cost function with a function less affected by outliers
Adding confidence to reference map segments	Shape matching algorithm	Asymmetrical adjustments
	Obstacle detection algorithm	

As seen in Chapter 1, even the quadratic cost function with no filtering, on the original map can sometimes produce feasible estimates. And when it does not, very often any one of the improvements in Table 4 is sufficient to improve the accuracy beyond the level that can be verified. The effect of the combined improvements is best viewed by looking at the end product: the trajectory of the car as described by the algorithm with all improvements made.

The following three graphs show these trajectories. According to the video data for this trial, the car stayed in the right lane (the lane farther away from the railroad tracks in the images) and maintained a fairly constant velocity for the entire time

sequence we used. Each point represents a position estimate from the RAFL method.

The estimates are 0.5 second apart.

As can be seen, replacing the quadratic cost function with the inverse quadratic cost function yields a large improvement in performance. Replacing that function with a cost function utilizing the error function with asymmetrical adjustments, as well as adding filtering, shape matching, obstacle detection, map corrections, and map confidence; produces a modest improvement over the inverse quadratic cost function.



Figure 38: The estimated trajectory using the quadratic cost function.



Figure 39: The estimated trajectory using the inverse quadratic cost function.



Figure 40: The estimated trajectory using the full suite of improvements.

Table 4 lists 7 improvements to our method, and all have significant effects on the accuracy. I have left out minor tweaks whose effect was difficult to ascertain.

However, the effects were not all equal. Table 5 shows the relative size of the effects.

Some comments on the effects:

Replacing the quadratic cost function with any of the other cost functions had a large effect on the performance of our method, in nearly every case except the few where the quadratic cost function performed adequately. Improving the cost function further provided only marginal gains that generally overlapped with the gains from other improvements. The combined effects of all improvements is only marginally better than the effect of the inverse quadratic cost function with some filtering.

Table 5: Effects of concepts for improvement

	Concept utilized	Effect on algorithm performance
1	Adjusting the map using ground-based measurements	Large near map points affected. No effect elsewhere.
2	Adding confidence to reference map segments	Moderate impact
3	Median filtering of the laser input	Moderate impact
4	Shape matching algorithm	Has limited impact. Probably not worth the computation needed.
5	Obstacle detection algorithm	Has impact only around obstacles, and when the sample swath is small. Probably not worth the computation needed.
6	Replacing the quadratic cost function	Anything other than the quadratic cost function is a very large improvement. Beyond that, the effects of improvement are moderate.
7	Asymmetrical adjustments	Moderate impact

Concepts 2-7 interact with each other, so the magnitudes of these effects depend on which other concepts are in place. Changes 2, 3, 4, 5, and 7 have the most effect when the other concepts are not introduced, that is, when there is the most room for improvement. Layering the improvements onto each other provides little benefit in our research, though it is possible that the capability to assess accuracy within centimeters may refine this hypothesis. The exceptions are the shape matching and

obstacle detection algorithms, which perform better when the laser input, which they both use, is filtered. Similarly, moving from the inverse quadratic cost function to the asymmetric error function cost function provides noticeable benefits only when other concepts (2,3,4,5,7) are not enacted. Concept 1 acts independently of the others. Replacing an erroneous part of the map with a corrected part provides similar benefits whether the calculations and filtering are optimized or not.

Implementing any or all of these concepts involves tradeoffs. The impact of inaccurate sections of the map can be reduced either by concept 1, which involves significant human effort beforehand, or by increasing the size of the sample swath, which uses more computing power each time our method is used. Changing the quadratic cost function to one of the improved symmetric or asymmetric functions presented is a must, but a designer concerned about computational power may decide to pick only one of the several other improvements suggested in this chapter.

Even with only the cost function change and one other concept introduced, the accuracy of our method increases substantially. The improvement is not in accuracy in any one trial, which is difficult to quantify, but in the number of trials for which our method achieves moderate to high fidelity to the video. Chapter 2 showed a trial in which even the quadratic cost function produced good results. Implementing the changes described in this chapter gives better performance in non-ideal situations.

Chapter 4: Conclusion

This conclusion discusses our contributions to the field of ADAS, and possible further work based on our research. Our major contributions include developing the RAFL method, designing a field test to verify our method, and adapting the inputs and mathematics to best suit our method. Some areas for future research include calculating the error with respect to ground truth, testing our method in different areas and with different equipment, and incorporating our estimates into an ADAS with sensor fusion methods.

Contributions

Our first contribution is the Registration with Automotive-Fixed Laser (RAFL) method of position estimation. This allows an estimate of vehicle position in both the lateral and longitudinal directions along mapped roads. We have shown a way to create a reference map from aerial data, and how to create a swath to compare to it. We also presented ways to identify and classify objects in both the reference map and the sample swath, in order to make our method more robust with respect to obstacles.

Our second contribution is the demonstration of our method with a road test. We used equipment costing less than three hundred dollars, along with a laptop computer and a few mounting brackets, and showed that our method functions with only the information from a DRS and a single laser rangefinder. We created a map of the area, and ran our algorithm after collecting a swath. We were able to demonstrate that our method has potential for calculating vehicle positions in real-world

applications. Additionally, we were able to adapt the cost functions and filtering to our method to improve its performance with more difficult data sets.

Future Work

Some follow-up work could increase the impact of what we have learned so far. We suggest five areas for further research: better estimation of the error of the RAFL method, testing in other streets and neighborhoods, sensor fusion, reducing the computation time of the algorithms, and sharing information between vehicles.

In order to improve the system, we intend to better calculate our error. The GPS estimate is not accurate within our target range, so we will need a better estimate of ground truth. A vision-based system coupled with surveyed points, similar to the method used by Bajikar and Gorjestani [6], could give us accuracies within a few centimeters. We could then decouple the accuracy of our method from the accuracy of our GPS receiver.

Testing our method with different equipment or different data sets could help us better understand our error. As of now, we do not know whether our accuracy is limited more by the reference map resolution, or more by the accuracy of our laser rangefinder and our DRS. We could also understand better whether our algorithm is tuned to our particular situation, or whether it could be used more widely. We have used our method to come up with accurate estimates only for a 700-meter stretch of Boston Avenue in Medford, Massachusetts, and this is the only place we have developed a reference map. Since our method's foundation is a comparison between a reference map and a sample swath, changing the reference map would have a significant impact on our method's performance. Areas with buildings of different sizes,

shapes and distances from the road could present new challenges. We would likely need significant revisions to our algorithm before it became useful in areas with winding roads or significant changes in elevation. A wide body of data could also reduce guesswork in tuning our algorithm, since we could test proposed improvements on statistically significant samples.

A third improvement could come from integrating our estimates with estimates from other sources, such as a GNSS receiver or an Inertial Navigation System. To assist with sensor fusion, we have been able to develop a method for assessing the quality of our estimates, with the possibility of weighting high-quality estimates more than low-quality estimates. Our efforts were preliminary, and further development could lead to better integration with other position estimates.

We did not focus on optimizing our computer algorithms for speed. If they are to be incorporated into an onboard computer, they will need to be efficient, since the computer will likely have several other ADAS and non-ADAS tasks to perform simultaneously. Before the RAFL method is incorporated in a non-research setting, work will need to be done to reduce our method's computational demands.

Lastly, there is a potential for collaboration between multiple vehicles. Our reference map was made by hand, point by point, and revision based on laser data was an uncertain and time-consuming process. It is possible that a number of vehicles equipped with rangefinders could streamline this process. With the proper safeguards, vehicles could give inputs to map coordinates and confidence levels, either directly with one another or through a central server. Multi-vehicle collaboration could increase both the efficiency and accuracy of our method.

Impact

The fields of ADAS and autonomous navigation can benefit society through improving the ease and safety of transportation. Our major contributions to these fields include developing the RAFL method, designing a field test to verify our method, and optimizing the inputs and mathematics used for our method. Future research, such as verifying the method against accurate ground truth, extending the method to more geographic locations, and incorporating our estimates into an ADAS with sensor fusion methods, are needed. With this research, the RAFL method may help bring safety improvements at a reasonable cost.

¹ Murray CJL et al. "Global and regional mortality from 235 causes of death for 20 age groups in 1990 and 2010: a systematic analysis for the Global Burden of Disease Study." *Lancet*, 2012, 380: p 2095-2128.

² Peden, M. et al. "World Report on Road Traffic Injury Prevention." World Health Organization. Geneva, 2004.

³ Staubach, M. "Factors correlated with traffic accidents as a basis for evaluating Advanced Driver Assistance Systems." *Accident Analysis & Prevention*, Volume 41, Issue 5, September 2009, p 1025-1033.

⁴ Kuehn, M., Hummel, T., and Bende, J. "Benefit estimation of advanced driver assistance systems for cars derived from real-life accidents." *German Insurers Accident Research*, Paper Number 09-0317; 2009. Note: Other accidents include collisions at intersections, rear-end collisions, collisions with cyclists, pedestrians, or obstacles. The 25% figure does not include accidents caused by intentional (though ill-advised) lane changes, though the distinction can be blurry at times. Another ADAS, blind spot detection, could help reduce these accidents.

⁵ Lu, M., Wevers, K., and Van Der Heijden, R. "Technical feasibility of advanced driver assistance systems (ADAS) for road traffic safety." *Transportation Planning and Technology*, 28:3, p 167-187, 2005.

⁶ Bajikar, S., Gorjestani, A., Simpkins, P., and Donath, M. "Evaluation of in-vehicle GPS-based lane position sensing for preventing road departure." *IEEE Conference on Intelligent Transportation Systems, Proceedings, ITSC*, p 397-402, 1997.

⁷ Obst, M et al. "Multipath Detection with 3D Digital Maps for Robust Multi-Constellation GNSS/INS Vehicle Localization in Urban Areas." *2012 Intelligent Vehicles Symposium*. 2012 June 3-7.

⁸ Chan, C., Bougler, B. Nelson, D. Kretz, P., Tan, H., and Zhang, W. "Characterization of magnetic tape and magnetic markers as a position sensing system for vehicle guidance and control." *Proceedings of the American Control Conference*, v 1, p 95-99, 2000.

⁹ Enkelmann, W. "Video-based driver assistance—from basic functions to applications." *International Journal of Computer Vision*, v 45, n 3, p 201-221, December 2001.

¹⁰ Chen, W. , Leon, J. , and Kuo, S. "Video-Based On-Road Driving Safety System with Lane Detection and Vehicle Detection." *2012 12th International Conference on ITS Telecommunications (ITST)*. 2012 Nov 5-8: p 537-541.

¹¹ McCall, J., and Trivedi, M. "Performance evaluation of a vision based lane tracker designed for driver assistance systems." *2005 IEEE Intelligent Vehicles Symposium Proceedings (IEEE Cat. No. 05TH8792)*, p 153-8, 2005.

¹² Zhao, H. and Shibasaki, R. "Reconstructing Textured CAD Model of Urban Environment Using Vehicle-Borne Laser Range Scanners and Line Cameras." *Computer Vision Systems: Lecture Notes in Computer Science*. 2001: Volume 2095, p 284-297.

¹³ Xu, Z. et al. "Map building and localization using 2D range scanner." *Proceedings 2003 IEEE International Symposium on Computational Intelligence in Robotics and Automation. Computational Intelligence in Robotics and Automation for the New Millennium (Cat. No.03EX694)*, p 848-53 vol.2, 2003.

¹⁴ Soloviev, A. Bates, D., and Van Graas, F. "Tight coupling of a laser scanner and inertial measurements for a fully autonomous relative navigation solution." *Journal of the Institute of Navigation*, Vol. 54, No. 3, p 189-205, 2007.

¹⁵ Madhavan, R. et al. Map-building and map-based localization in an underground-mine by statistical pattern matching. *Proceedings. Fourteenth International Conference on Pattern Recognition (Cat. No.98EX170)*, p 1744-6 vol.2, 1998.

¹⁶ Joerges, M. and Pervan, B. "Autonomous ground vehicle navigation using integrated GPS and

laser-scanner measurements.” *IEEE PLANS, Position Location and Navigation Symposium*, v 2006, p 988-997, 2006.

¹⁷ Levinson, J. et. al. “Towards fully autonomous driving: systems and algorithms.” *2011 IEEE Intelligent Vehicles Symposium*, p. 163-168, June 2011.