# Vehicle Localization

Hannah Rae Kerner

21 April 2015

# Spotted in Mtn View: Google Car

# Why precision localization?

- in order for a robot to follow a road, it needs to know where the road is
- to stay in a particular lane, it needs to know where the lane is
  - for an autonomous robot to stay in a lane, localization must be accurate to decimeters at least

# Vehicle Localization Problem

- Autonomous driving and ADAS applications can be significantly improved by more accurate (cm-level) vehicle localization
  - important for safety in urban environments
  - narrow passages, turns, etc
  - GPS-denied areas e.g. parking garages, in between buildings, etc
- GPS-IMU-odometry based methods are not adequate for this positioning accuracy

# Techniques for Improvement

- Many techniques for increasing location accuracy for urban driving
  - Extended Kalman Filters, Belief Theory, multi-vehicle cooperation, and more…
- We’ll look at the one published by the group that led the development of the Google driverless car

# Map-Based Precision Vehicle Localization in Urban Environments

**Jesse Levinson, Michael Montemerlo, Sebastian Thrun**
**Stanford Artificial Intelligence Laboratory (2008)**

Augment inertial navigation (GPS + odometry) by:

1. learning a detailed map of the environment
2. using the vehicle’s LIDAR sensor to localize relative to that map

http://www.roboticsproceedings.org/rss03/p16.pdf

# 1. **Learning a detailed map**

Map contains:

- 2-D overhead views of the road surface
- infrared spectrum
- captures lane markings, tire marks, pavement, vegetation (grass), etc

# Acquiring the map

multiple SICK laser range finders pointing downward at the road, mounted on vehicle

**Figure 2.** Visualization of the scanning process: the LIDAR scanner acquires range data *and* infrared ground reflectivity. The resulting maps therefore are 3-D infrared images of the ground reflectivity. Notice that lane markings have much higher reflectivity than pavement.

- return range to sampling of points on the ground
- return measure of infrared reflectivity
- result: 3-D infrared images of ground reflectivity

# Eliminating Dynamic Objects

fits a ground plane to each laser scan and removes objects above the plane

- other cars, buildings, lamp posts, etc along the road are not included in the map

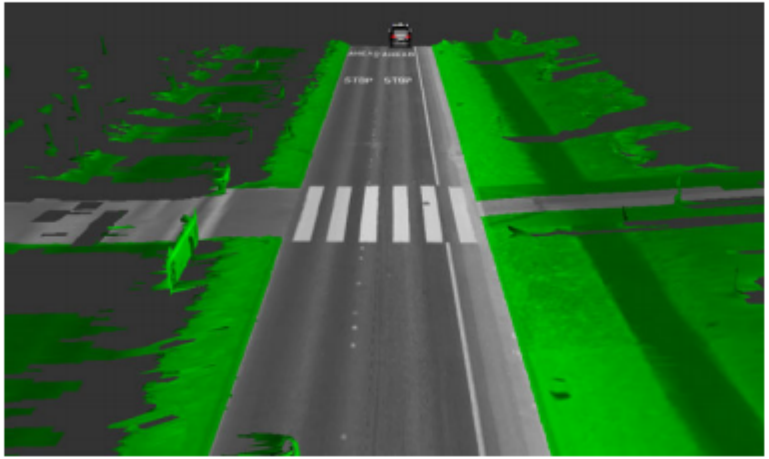

**Figure 3.** Example of ground plane extraction. Only measurements that coincide with the ground plane are retained; all others are discarded (shown in green here). As a result, moving objects such as car (and even parked cars) are not included in the map. This makes our approach robust in dynamic environments.

# Map Storage

- rectangular area acquired by range scan decomposed into square grid
- saves only squares for which there is data
- after lossless compression, grid images require ~10MB per mile of road at 5cm res.
- thus a 200GB hard drive can hold 20,000 miles of data
- particle filter maintains cache of image squares near the vehicle, thus requiring constant amount of memory

# 2. Localizing relative to map in RT

1. Particle filter analyzes range data to determine the ground plane the vehicle is on
   (also combines GPS data when available)
2. Correlates measured infrared reflectivity with the map (using the Pearson product-moment correlation)
3. Tracks location by projecting particles forward through time via the velocity outputs from inertial navigation system

# Localization

- uses hardware-accelerated OpenGL to render map for localization (faster than real-time even with low-end graphics card)
- localization computed with 200 Hz motion update
  - measurements arrive from each laser at 75 Hz
- uses a particle filter (Monte-Carlo localizer)
  - maintains 3-D pose vector: *x, y, yaw*

# Weather Complications

- wet surfaces tend to reflect less IR light than dry ones, so maps in the same loc. differ slightly
- particle filter normalizes brightness and standard deviation for each range scan as well as corresponding map stripes

(a) Map acquired on a sunny day.

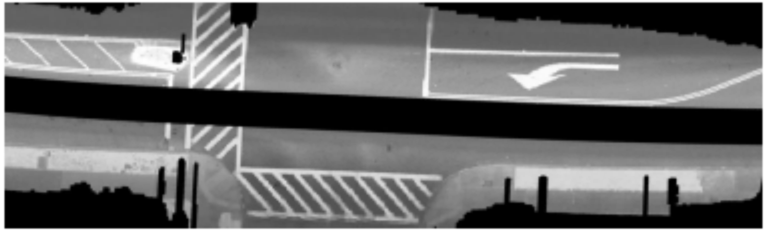

(b) Same road segment on a rainy day.

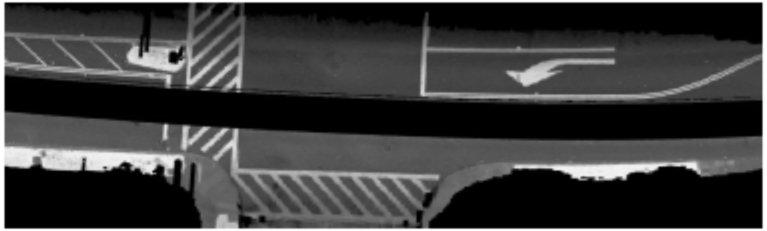

**Figure 4.** Patch of the map acquired in bright sunlight on a sunny day (top), and at night in heavy rain (middle). By correlating scans with the map, instead of taking absolute differences, the weather-related brightness variation has almost no effect on localization.

# Experimental Results

- state-of-the-art inertial nav system
- three down-facing laser range finders: left, right, and rear
- 5-cm pixel resolution

**Figure 1.** The acquisition vehicle is equipped with a tightly integrated inertial navigation system which uses GPS, IMU, and wheel odometry for localization. It also possesses laser range finders for road mapping and localization.

# Experimental Results

tested mapping algorithm successfully on variety of urban roads, e.g. this map acquired in Burlingame in 32 loops:

# “Ghosting” removal

(a) GPS leads to ghosting (b) Our method: No ghosting

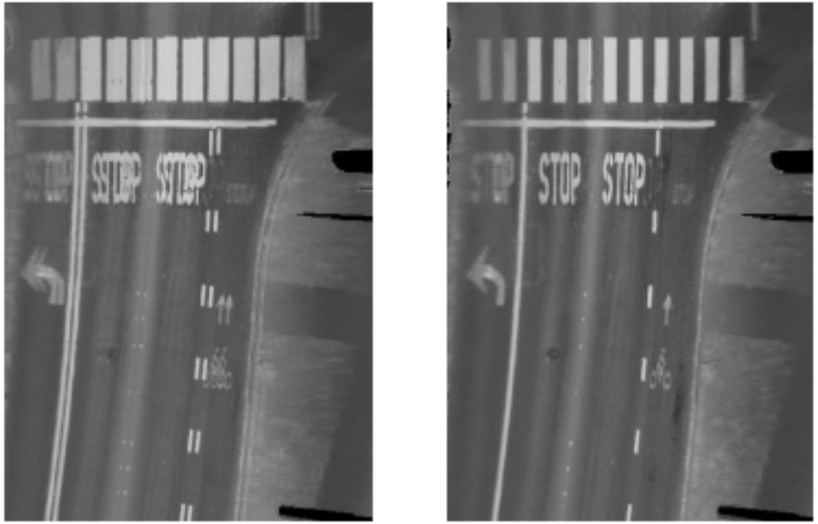


**Figure 8**. Infrared reflectivity ground map before and after SLAM optimization. Residual GPS drift can be seen in the ghost images of the road markings (left). After optimization, all ghost images have been removed (right).

(a) Map with hole (b) Ghosting (c) Our approach

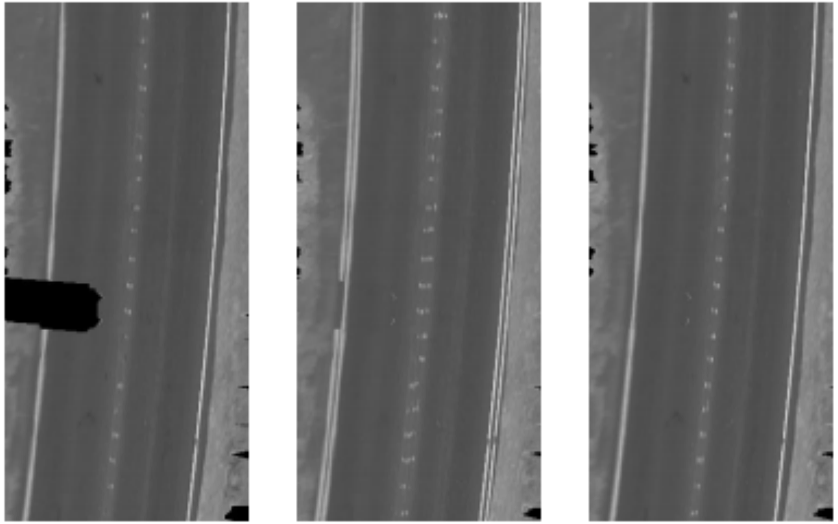

**Figure 9**. Filtering dynamic objects from the map leaves holes (left). These holes are often filled if a second pass is made over the road, but ghost images remain (center). After SLAM, the hole is filled and the ghost image is removed.

# Empirical Results

- very reliably tracks location of vehicle with relative accuracy of ~10cm
  - used 200 to 300 particles
- both mapping and localization processes robust to dynamic and hilly environments
  - so long as the road surface remains approx. laterally planar in the neighborhood of the vehicle

# Localization without GPS

successfully localizes even with GPS turned off (using only odometry and steering angle)

| Stanford Ave. | | |
|---|---|---|
| Distance Traveled (m) | Our Error (cm) | Odometry Error (cm) |
| 50 | 7 | 98 |
| 100 | 3 | 149 |
| 150 | 35 | 0 |
| 200 | 13 | 8 |
| 250 | 4 | 133 |
| 300 | 22 | 272 |
| 350 | 8 | 428 |
| 400 | 23 | 589 |
| 450 | 13 | 783 |
| 499 | 10 | 924 |

**Figure 12.** This table compares the accuracy of pose estimation in the absence of GPS or IMU data. The right column is obtained by odometry only; the center by particle filter localization relative to the map. Clearly, odometry alone accumulates error. Our approach localizes reliably without any GPS or IMU.

# Localization using only LIDAR

- GPS, IMU, and odometry were all ignored
- particle state vector: x, y, yaw, steering angle, velocity, and acc.
  - initialized near true position
  - assumed reasonable rates of change
- reasonably successfully tracked pos. and velocity

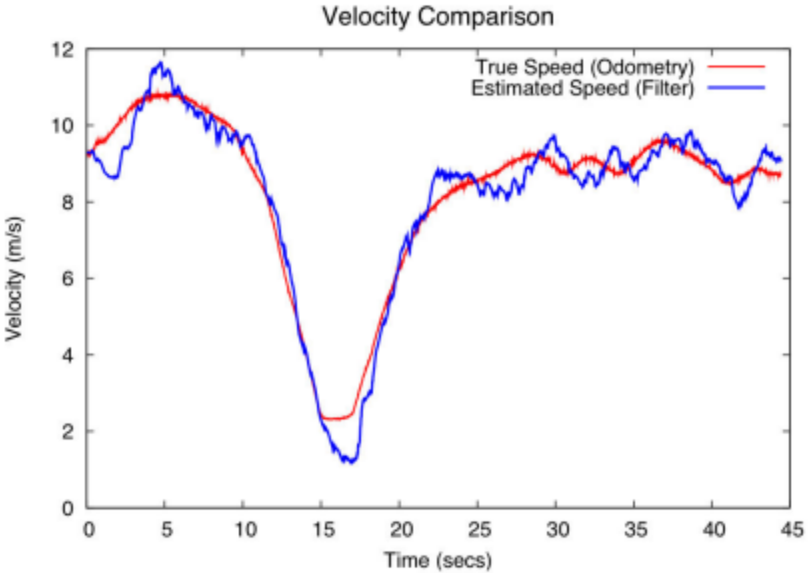

# Empirical Results

- localization results after 20 minutes of driving on top of acquired map
- lateral error almost always within 10cm but on turns sometimes as much as 30cm

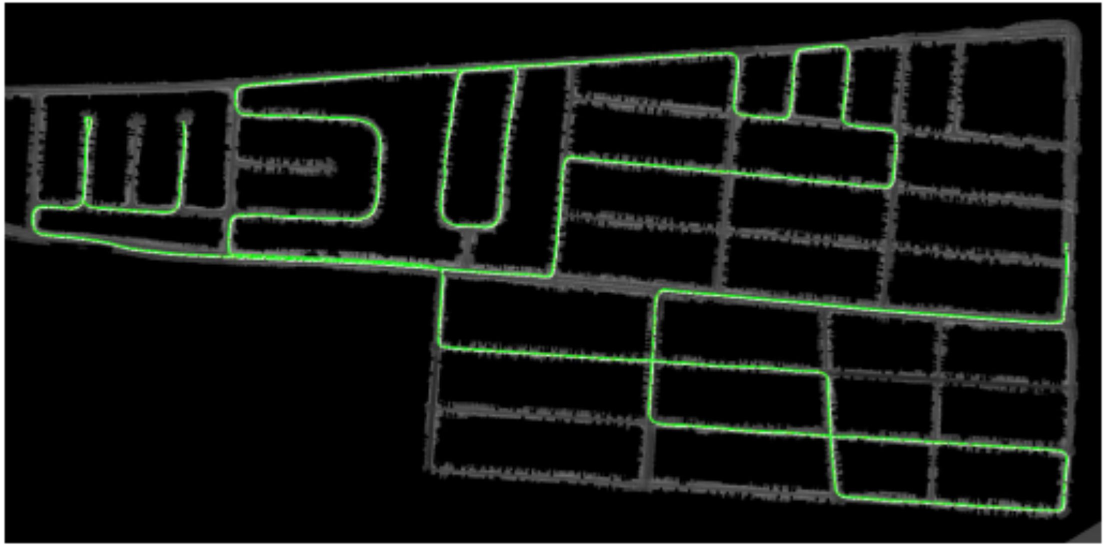

**Figure 10.** (Best viewed in color) Typical driving path during localization shown in green, and overlayed on a previously built map, acquired during 20 minutes of driving. For this and other paths, we find that the particle filter reliably localizes the vehicle.

# Importance of Localization Techniques

average disagreement between real-time GPS pose and their localization method was 66-cm

(a) GPS localization induces ≥1 meter of error.

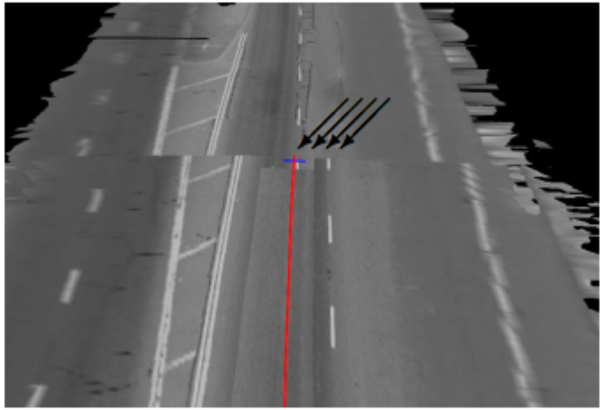

(b) No noticeable error in particle filter localization.

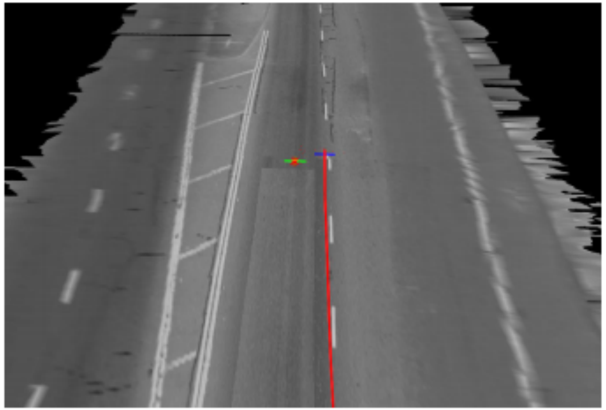

**Figure 11**. (a) GPS localization is prone to error, even (as shown here) with a high-end integrated inertial system and differential GPS using a nearby stationary antenna. (b) The particle filter result shows no noticeable error.

# Autonomous Driving Experiments

- ten attempts to drive autonomously through an urban area
    - gas and brakes operated mostly manually, but all steering done by computer
- followed fixed reference trajectory through Stanford campus without error 10/10 times
- often the lane width not occupied by vehicle was less than 2 meters, yet GPS-only consistently failed within meters: GPS localization **not** sufficient

# Conclusions

- accurate localization enables autonomous cars to perform accurate lane keeping and obey traffic laws
- GPS is not sufficient for autonomous vehicle localization, yet almost all outdoor localization work is GPS-based
- this method is better for both accuracy and availability
- disadvantage of approach: reliance on maps