



Technologies for Safe & Efficient Transportation

THE NATIONAL USDOT UNIVERSITY
TRANSPORTATION CENTER FOR SAFETY

Carnegie Mellon University

UNIVERSITY of PENNSYLVANIA

Smartphone Based Traffic Sign Inventory and Assessment

FINAL RESEARCH REPORT

Christoph Mertz (PI), John Kozar, Jinhang Wang,
Joseph Doyle, Christopher Kaffine, Anuja Kelkar,
Nikitha Poddatur, Lam Chan, Kedar Amladi

Contract No. DTRT12GUTG11

DISCLAIMER

The contents of this report reflect the views of the authors, who are responsible for the facts and the accuracy of the information presented herein. This document is disseminated under the sponsorship of the U.S. Department of Transportation's University Transportation Centers Program, in the interest of information exchange. The U.S. Government assumes no liability for the contents or use thereof.

Introduction

Road signs are an important part of the infrastructure and are needed to ensure smooth and safe traffic flow. Faded, occluded, damaged or vandalized signs can confuse or misinform drivers and lead to unsafe driving behavior. E.g. if a driver is not able to see a stop sign, he or she might drive into an intersection without stopping and cause an accident.

Government agencies are tasked with maintaining good signage and part of it is regular inspections to detect problems. Current methods involve manual inspections, specialized vehicles, or citizen reports. They are tedious, expensive, or not always reliable. In this project we developed a traffic sign inventory and assessment system that built on our smartphone based road inspection system.

Background and motivation

In previous years the Navlab group had developed a road inspection system that is based on a vehicle mounted smartphone. We have published a description of the system¹ and details of the computer vision techniques used to analyze the images². The smartphone is mounted on the windshield and is powered by cigarette lighter (Figure 1 left). While the vehicle is driving the smartphone collects images or videos of the outside and tags them with time, GPS, and other selected information.

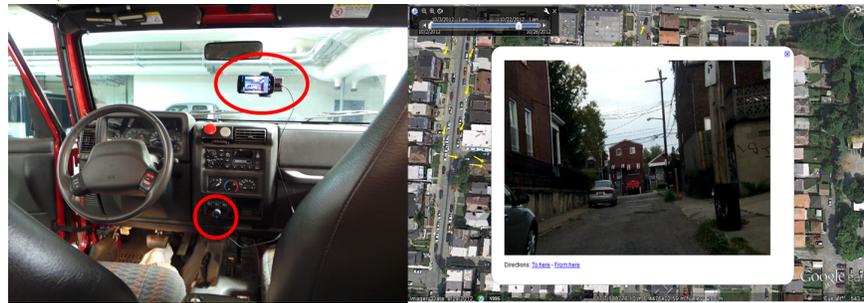


Figure 1 Left: Smartphone mounted inside a vehicle and powered by the cigarette lighter. Right: Example of road image displayed on Google Earth. The small yellow arrows on the street are markers pointing in the driving direction. When clicking on the marker the corresponding image appears.

One of the key ideas behind the collection system is that it can be easily mounted on any vehicle, especially those that drive on the roads on a regular basis, e.g. garbage trucks drive through every neighborhood once a week. It is therefore possible to collect data frequently without the need for a dedicated vehicle or a dedicated driver.

The images can be displayed in the asset management system of the department or with free software. An example is shown in Figure 1 (right) where the data is displayed on Google Earth. This will allow the user to inspect the road from a computer instead of physically going to the road.

In the first version of the system we automatically detected road distress by using computer vision algorithms to find road cracks in the images. In this second version we extended it to detect road signs, specifically stop signs. We choose this after receiving feedback from the city of Pittsburgh that traffic signs and graffiti are important problems. Traffic signs are

¹ C. Mertz, S. Varadharajan, S. Jose, K. Sharma, L. Wander, and J. Wang, "City-Wide Road Distress Monitoring with Smartphones", ITS World Congress, Detroit MI, September 2014.

² S. Varadharajan, S. Jose, K. Sharma, L. Wander, and C. Mertz, "Vision for Road Inspection" IEEE Winter Conference on Applications of Computer Vision, March 2014, accepted for publication.

significant parts of the road infrastructure and they need to be inspected and maintained. It is an important safety concern to detect degradation or obstruction as early as possible. Graffiti is for the most part a nuisance problem that is difficult for the City to address because the general public usually does not report it and regular inspections are too costly. The best method to defeat it is to remove it as soon as possible so that the intended exposure to the public is denied. Detecting graffiti in images is a very challenging machine vision problem. As a first step we want to detect graffiti or other vandalism on traffic signs. Those are not only a nuisance but also safety concerns and they are easier to detect because we know a-priori how a clean traffic sign should look like.

Sign detection methodology

The main software modules we developed are sign detection and sign assessment. Sign detection is an established vision application^{3,4}. Some of the methods are quite accurate; detection rates of 99% are not uncommon. We tested several methods to find the one best suited for our purposes. One important issue will be how easy it is to train the classifier. As there are many different traffic signs in each city, region, or country we want to make it straightforward to find any specified sign. For this we used an exemplar method. However, we found that using a method that utilizes HOG features and a SVM gives much more accurate results and it can also be extended to detect problems with a traffic sign. For both methods the first step is the same. We detect areas (“blobs”) in the image that are red (Figure 2, red blobs are indicated by boxes).



Figure 2 Stop sign detection: First, red areas are detected (boxes). Then the areas are analyzed if they look like a stop sign (orange box) or not (blue boxes).

These blobs are then tested if they appear to be stop signs or not. With the first method we only need one exemplar, i.e. a synthetic image of an ideal stop sign. A similarity distance between the blob and the exemplar is calculated and if the similarity distance is below a threshold, the blob is classified as a stop sign. In the second method we use a set of positive example blobs (about 1000 stop signs) and a set of negative example blobs (about 2000 red blobs not containing stop signs). For each set we calculate the HOG features of the images and then train a SVM classifier. We tested both methods and found that the

³ A. Mogelmoose, M. Trivedi and T. Moeslund, "Vision-Based Traffic Sign Detection and Analysis for Intelligent Driver Assistance Systems: Perspectives and Survey," *Intelligent Transportation Systems*, IEEE Transactions on, vol. 13, no. 4, pp. 1484-1497, 2012.

⁴ S. Houben, J. Stallkamp, J. Salmen, M. Schlipsing and C. Igel, "Detection of traffic signs in real-world images: The German Traffic Sign Detection Benchmark," in *International Joint Conference on Neural Networks* (submitted), 2013

similarity-distance method is much easier to use because one does not have to produce all the positive and negative examples, but the performance is significantly worse than the HOG-SVM method. However, the similarity-distance method can be used to make the selection of examples for the HOG-SVM training much easier. Figure 3 shows the efficient training pipeline.

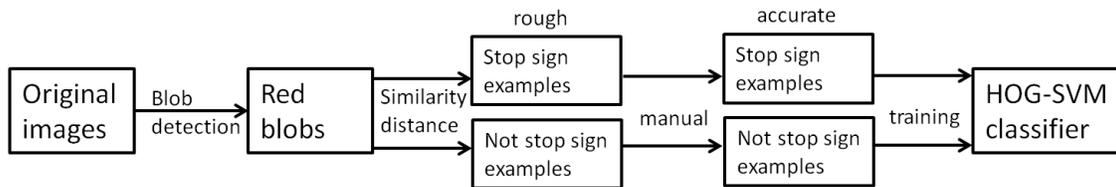


Figure 3 Training pipeline for stop sign detection.

After the blob detection the blobs are roughly sorted into stop-sign and not-stop-sign examples. These are then manually corrected to get the two sets of examples. These are then used to train the HOG-SVM classifier.

We further trained the HOG-SVM classifier to distinguish between various kinds of stop signs. These are regular stop signs, stop signs with “3 ways” etc. underneath it, other stop signs, defaced stop signs, stop signs occluded by vegetation and displaced stop signs (Figure 4).



Figure 4 Various kinds of intact stop signs and stop signs with problems that can be detected by the HOG-SVM classifier.

We analyzed our large data set of street images (more than 500 hours of recording, about 20 million images) to find the stop signs. First we sub-sampled the images to have only one every 5 meters. In these images we detected the stop signs and detected if there are any problems with them. Figure 5 shows all the stop signs we found in Pittsburgh and the Townships of Cranberry and Marshall. It also shows the stop signs with problems.

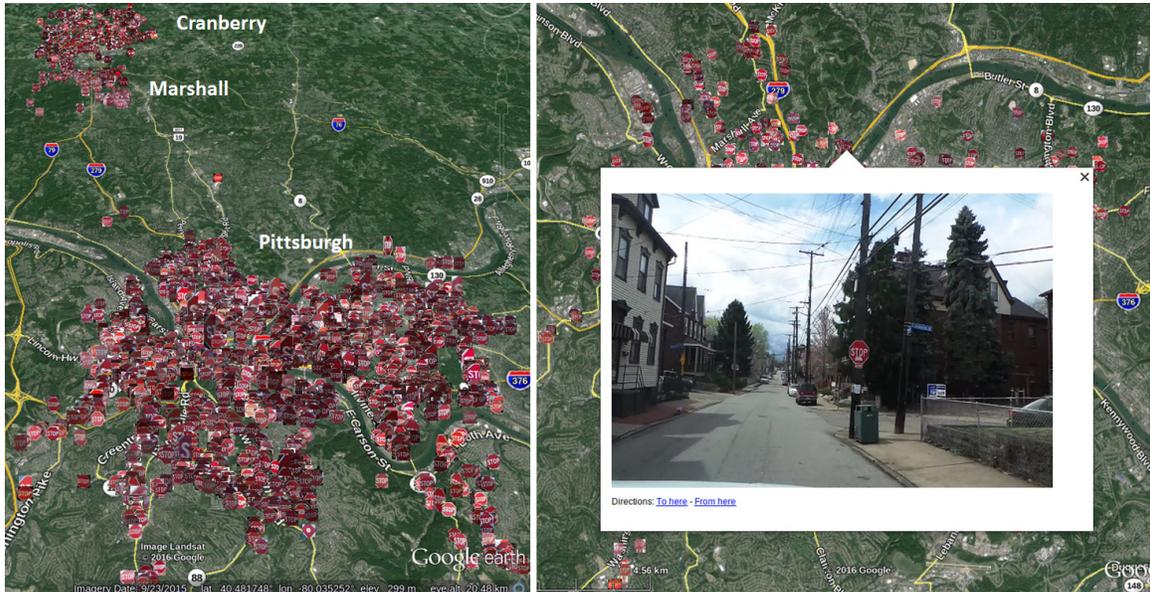
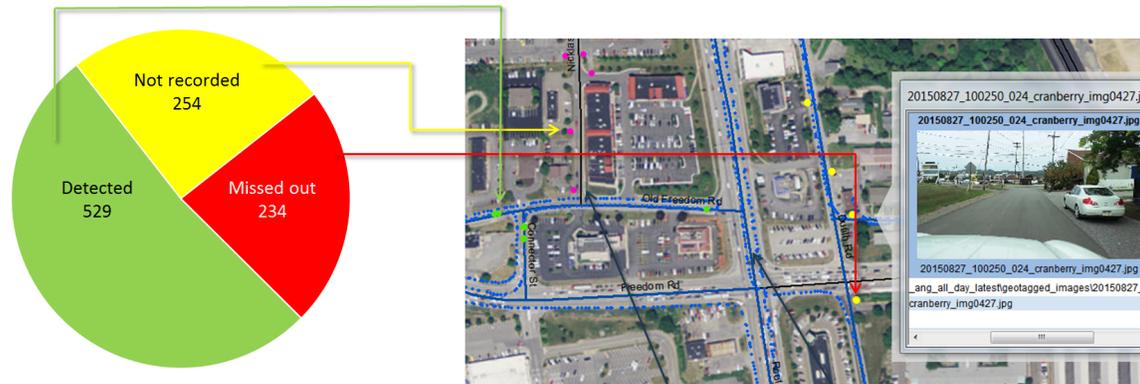


Figure 5 Left: All the stop signs we found in Pittsburgh, Cranberry Township, and Marshall Township displayed on Google Earth. Right: Stop signs in Pittsburgh with problems. By clicking on a stop sign icon one can see the corresponding full picture. Shown is a full picture with a stop sign defaced by a sticker.

The method to detect stop signs can easily be extended to detect any colored traffic sign. One would first have to define the desired color of the sign and find the color blobs in the images. Then one has to get one exemplar of the traffic sign and use the similarity distance method followed by manual checking to get the positive and negative examples. These in turn can then be used to train the HOG-SVM.

Real world test and evaluation

We tested the system in Cranberry Township with the help of their maintenance department. They collected the road data with our collection system. We analyzed the data to find all the stop signs and determine any problems with it. The result was compared with a stop-sign inventory they had established independently from our system. The database consists of accurate GPS locations (DGPS) of the traffic signs and street centerlines. Figure 6 shows the two data sets together.



Total 1017 stop signs in Cranberry dataset

Total 637 road segments, 441 segments are recorded

Figure 6 Comparison of Cranberry inventory with our data set. Cranberry: straight lines = street centerlines, large dots = stop signs. Our data set: Small dots = images recorded. Green large dot = our detection agrees with Cranberry inventory.

Figure 7 shows in more detail the reasons and percentages of missed detections. Of the approximately 1000 stop signs in the inventory we detected slightly over half. The main reason was that data was not taken on the corresponding streets. Of the remaining missed ones the great majority could not be seen by the system because the stop signs were facing away from the street (Figure 8 left), often these were stop signs at the exit of a parking lot. It is straightforward to fix all the aforementioned misses. One has to drive on all the streets in both directions and also cover all the exits of parking lots. There was also a case where the sign was not seen because it happened to be blocked by another vehicle.

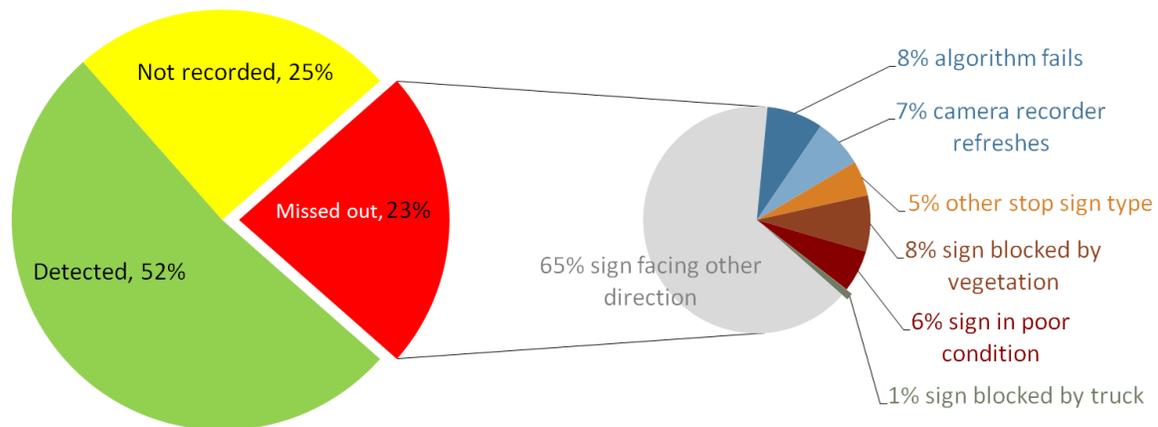


Figure 7 The reasons and percentages of the stop signs in the inventory that were not detected by our system.

In some cases it was our system that missed the signs. Either the algorithm did not detect it or the system missed a stretch of road because it was refreshing for a short time. We have already identified ways to reduce these misses by adjusting the parameters of the algorithm and reducing the refresh time.

Occasionally in the inventory traffic signs that are called “stop signs” are actually not the stop signs we were interested in but other kinds of stop signs, like warnings of upcoming stop sign or stop-here-for-red-light signs. These apparent misses are therefore only a matter of making the labels in the data bases consistent.

The remaining misses are the most interesting ones. They were missed because they were blocked by vegetation (Figure 8 left middle) or because they were in poor condition (Figure 8 right middle). These are cases which require maintenance.



Figure 8 Missed detections. Left: Stop sign facing away from the street. Left middle: Occluded by vegetation. Right middle: Faded. The picture on the right shows stop signs (light blue) that we detected by were not in the inventory.

There were also stop signs that we detected but were not in the inventory (Figure 8 right). These were mainly new developments where the township has not created the inventory, yet. They could either use our data as their inventory or go out and do their usual DGPS measurements.

Finally, we checked all the stop signs we captured for any problems. We detected defaced, faded, and partially occluded stop signs (examples shown in Figure 9)



Figure 9 Problems with stop signs: Defaced, faded, and partially occluded by vegetation.

This real world test helped us refine our collection and analysis methodology. Most importantly we were able to assess the stop signs of Cranberry Township and find problems: Signs completely occluded by vegetation, in poor condition, defaced, faded, or partially occluded by vegetation. Some signs were missing from their inventory.

Conclusion and Recommendation

The feedback from our partners in Cranberry Township and the City of Pittsburgh were very positive. In their opinion this is an efficient and cost effective way to inventory and assess street signs. There are still a few improvements to the full system that we can make, but it is basically ready to be pilot tested on a larger scale. The system should also be expanded to include other traffic signs.