Monitoring Roads and other Infrastructures with Smartphones in Pittsburgh – Stage I final report

Christoph Mertz
Robotics Institute, Carnegie Mellon University
cmertz@andrew.cmu.edu, 412-2687955

Table of Contents

The City of Pittsburgh and CMU completed the first stage of implementing and testing a smartphone based road monitoring system that was developed by CMU. In this project report we first want to describe the state of the system before the project start and then report on accomplishments achieved during the project.

Road monitoring system at project start

State and local maintenance departments are tasked with keeping roads in good repair which includes monitoring the roads to detect the presence of cracks, potholes, and other distress. Currently, this is done by inspectors, specialized vehicles [1], or citizen reports. These methods are respectively tedious and often inconsistent, generally expensive, or mainly only about severe and acute problems. None of them is able to monitor the road on a continuous basis at low cost, with consistent quality, and with minimal human intervention.

Our approach is to use images or videos collected by commodity devices such as smartphones. The devices are mounted in vehicle that travels the roads for other purposes and therefore no dedicated vehicles or drivers are needed. The captured images are analyzed automatically with machine vision algorithms and the resulting distress scores are passed to the asset management system.

Methodology

The **data collection system** is shown in Figure 1. A smartphone is mounted on the windshield and is powered by cigarette lighter. While the vehicle is driving the smartphone collects images or videos of the outside and tags them with time, GPS, and other selected information.

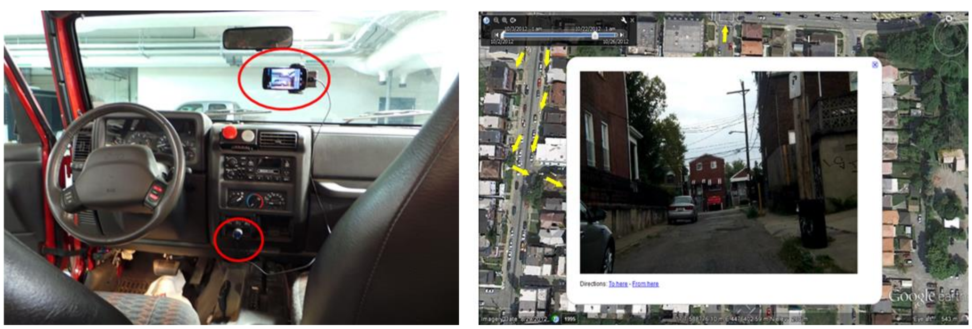

Figure 1 Left: Smartphone mounted inside a vehicle and powered by the cigarette lighter. Right: Example of road image displayed on Google Earth. The small yellow arrows on the street are markers pointing in the driving direction. When clicking on the marker the corresponding image appears.

One of the key ideas behind our data 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 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.

The details of the analysis of the images to determine the amount of damage of the road is described in the paper [2]. An example is shown in Figure 2. On the left is a typical image and on the right are patches with cracks where the intensity of red indicates the severity of the cracks. A distress score can now be calculated, it is the ratio of the area with and without cracks in front of the vehicle.

Figure 2 Left: A typical road image. Right: Classification result. The color indicates the severity of distress: blue = no cracks, light to dark red = light to severe cracks.

Results

We have mounted the system on a personal vehicle and collected data for almost 2 years. The routes included daily commutes and other typical day-to-day driving, together about 250 hours of driving on 600 miles of unique roads. We automatically selected a subset of the data with favorable weather conditions (daytime, overcast, no precipitation) by querying public weather sources. For this data we calculated the distress score, it is shown in Figure 3 overlaid on the GIS street center lines used by the city of Pittsburgh.

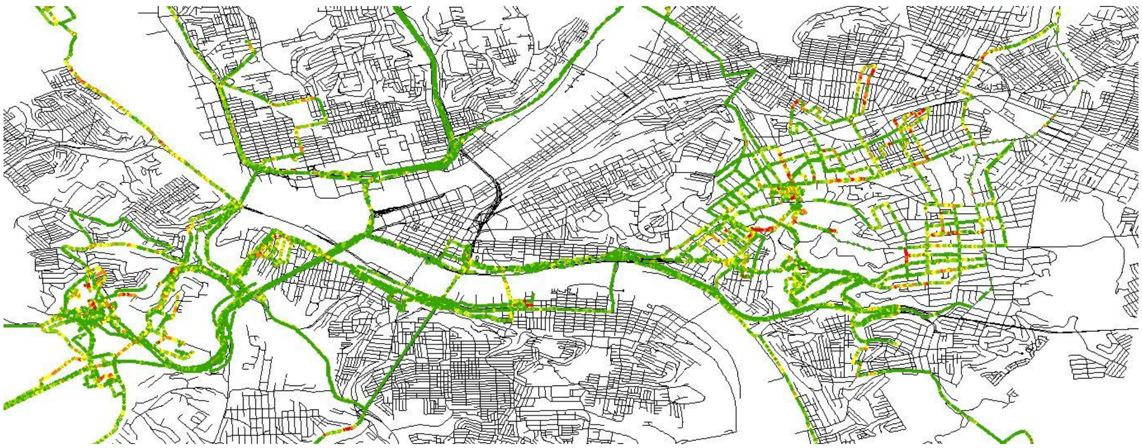

Figure 3 Distress score overlaid on a GSI street map of Pittsburgh. Green, yellow, and red are low, medium, and high amount of cracks, respectively.

We compared the performance of our system with that of human inspection; our system comes close to the performance of a human.

We also received the road ratings from the City of Pittsburgh. They contained the Overall Condition Index (OCI) and rehabilitation dates for each city road segment. We did find correlations between the OCI and our score, however, these two are difficult to compare. OCI contains measures like rutting and raveling whereas our score is only based on cracks. Instead, we show in Figure 4 the relationship of the time since the observed road has been rehabilitated and the average damage score. It is expected that the road deteriorates over time and the score increases accordingly. One can see that it is indeed the case.

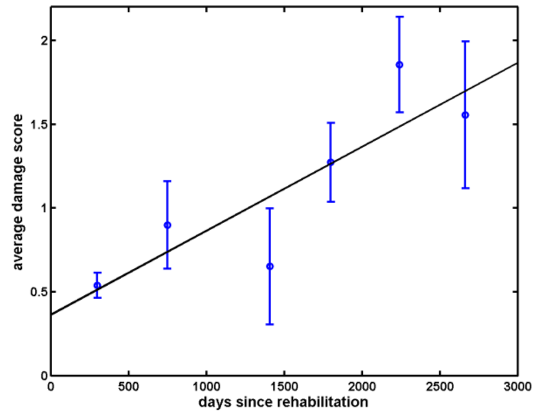

Figure 4 Average road cracks score vs. time since the last rehabilitation of the road. The error bars indicate the uncertainty of the averages.

The variations are large, as indicated by the error bars, because cracks are not evenly distributed along a road. One can find clusters of cracks and undamaged areas on the same road.

Project Activities

At the beginning of the project the City of Pittsburgh and CMU met together on several occasions to come up with a specific task that the City can do with the road monitoring system by the end of the pilot project. The task is to determine the 2016 paving schedule. This schedule has to be developed at the end of 2015. Currently the problem is that there are too many roads that need to be paved and there is insufficient data to determine which roads are in worst condition. The first goal of the pilot is therefore to collect data from all of Pittsburgh's streets. This data needs to be incorporated into the city's road management system in a way that it can be used to develop the schedule.

Data collection

We installed three data collection systems on city vehicles. A collection system consists of a smart-camera mounted on a windshield with a suction cup and powered from the cigarette lighter. These vehicles are driven during the work routine of the city employees or on routes purposefully selected to cover all the streets of Pittsburgh. Additionally, CMU has its own collection system that it continues to use. The Figure 5 shows the data that has been collected between the end of 2014 and April 2015. Large parts of Pittsburgh have already been covered and it is expected that all of Pittsburgh will be covered in a few months.

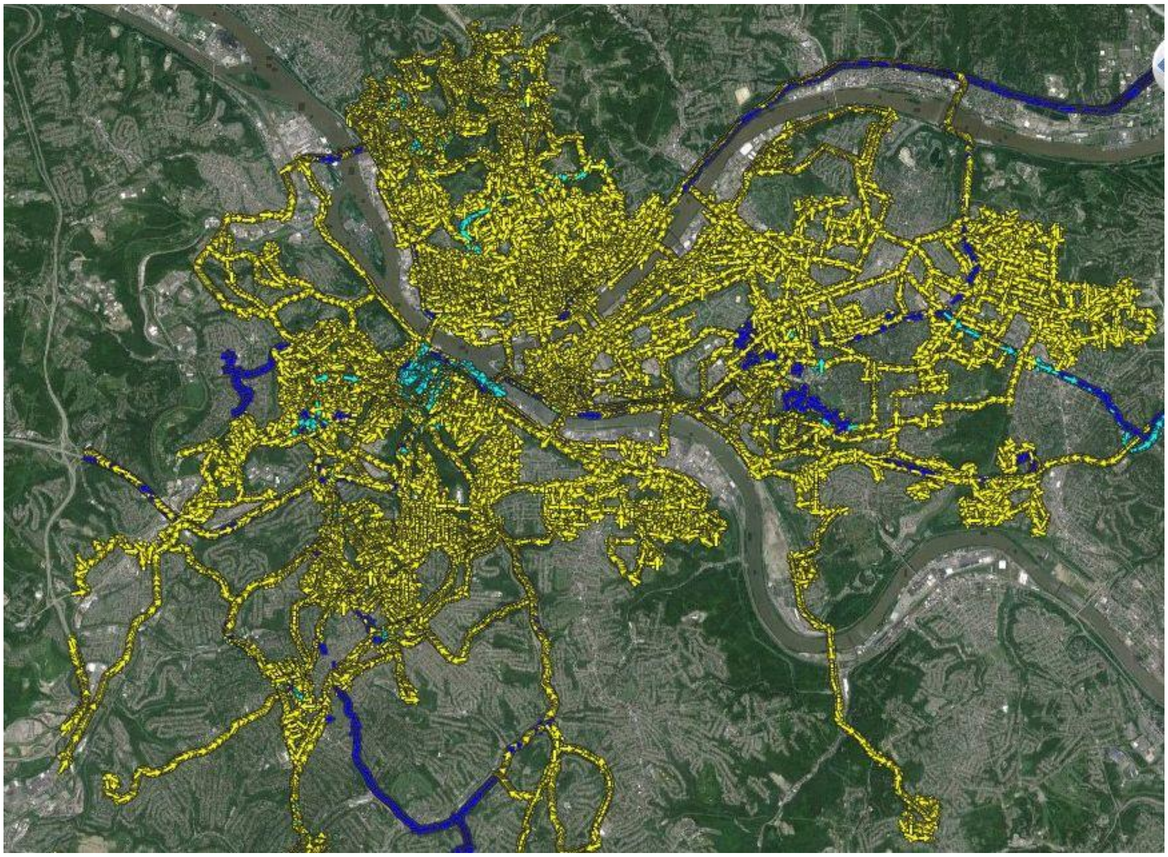

Figure 5 Data collected in Pittsburgh between end of 2014 and April 2015. Yellow, light blue and dark blue are day, dawn/dusk, and night collections respectively.

After a small amount of training the city employees were able to do the collection. The collection went smoothly for the most part, only a few times did the collection app crash. After about 2 weeks of collections the drivers bring the data on the microSD card to CMU where they exchange it for an empty one. CMU then processes the data.

We are in the process of re-writing the app to make it more stable and make it possible to download the data directly to cloud storage so that the drivers do not need to hand-deliver the data.

Data sharing between City of Pittsburgh and CMU

The current practice of the city is to manually inspect road distress condition. Inspectors estimate the score based on the percentage of distress area and severity for each road segment. Road scores are stored in Cartegraph databases. As mentioned above, CMU algorithm extracts images from the video and generates a score for each image. These scores are then averaged over a road segment to arrive at one score for each road segment.

We had several meeting with the city where we discussed the best ways to exchange data and develop common data formats. CMU data can be structured for display and other uses in three approaches:

1) The data is structured in a kml file. These are files that can be read by free software like GoogleEarth™. An example is shown in Figure 5. This data format can be used by anyone without the need to have access to proprietary software.

2) We can create shapefiles and geodatabases to display the score and image information with the ArcGIS platform. In the final map interface, users can view the road image, score, as well as other basic road information by clicking on the point in the map. An example is shown in Figure 3. This data format is used by many government agencies, it does require proprietary software to use and display it.
3) We can populate CMU data into Cartegraph, in the format that complies with city's database structure. Cartegraph is the proprietary software used by the city as their infrastructure management software (IMS).

After discussions with the city we jointly decided to use both the second and third approach. In the second approach, CMU generates a shapefile which can be displayed in city's ArcGIS system. We provide a geodatabase of photos that displays the score/image/road information upon clicking.
In the third approach, CMU generates scores that populate the Cartegraph database. An example of the data structure looks as follows:

	A	B	C	D	E	F	G	H	I	J
1	CarteID	Date	Inspected by	ID	Type	Length	ID	Distress	Severity	Extent
2	15889	10/4/2012	Computer	100420121	Random	213.5	100420121	AC cracking	Moderate	1.2908
3	15186	10/4/2012	Computer	100420122	Random	182	100420122	AC cracking	Moderate	2.1369
4	17480	10/4/2012	Computer	100420123	Random	455.5	100420123	AC cracking	Moderate	1.3112
5	18761	10/4/2012	Computer	100420124	Random	383.21	100420124	AC cracking	Moderate	0.14791
6	15989	10/5/2012	Computer	100420125	Random	560.21	100420125	AC cracking	Moderate	0

Inspection

Distress sample

Distress

We successfully tested incorporating this example into the Cartegraph IMS of the city.

Outreach

We had meetings with several governmental and university groups to present our work. Some of them expressed interest in our system and we started collaborations with them or agreed to consider collaborating later this year. Among them is Marshall Township. They borrowed one of our collection units and recorded data from their streets. An overview of the data is shown in Figure 6.

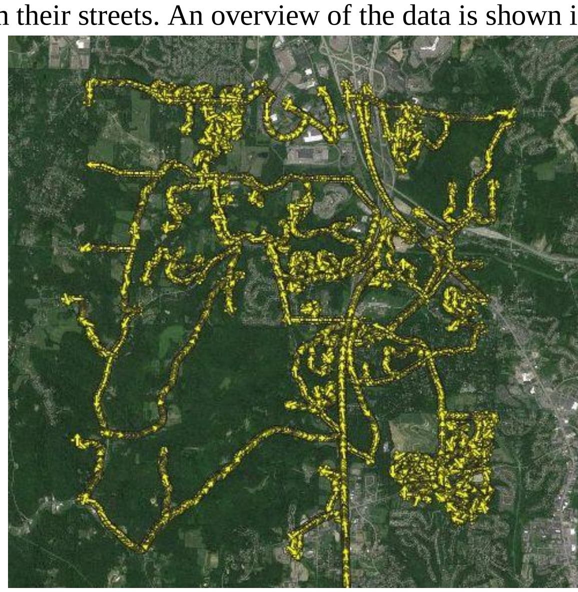

Figure 6 Data collected for Marshall Township.

System improvements

The software for the system was written as research code, i.e. it is used to test ideas. We have started to re-write the software to make it more user-friendly and efficient. The goal is also that everyone can use it without needing proprietary software tools or libraries.

We have also started the development of new capabilities. The images collected for road distress monitoring can also be used to inventory and assess many other infrastructure parts. Judging from the interest expressed by various governmental managers, we decided to look into stop-sign detection and lane marking assessment. Figure 7 shows the first results.

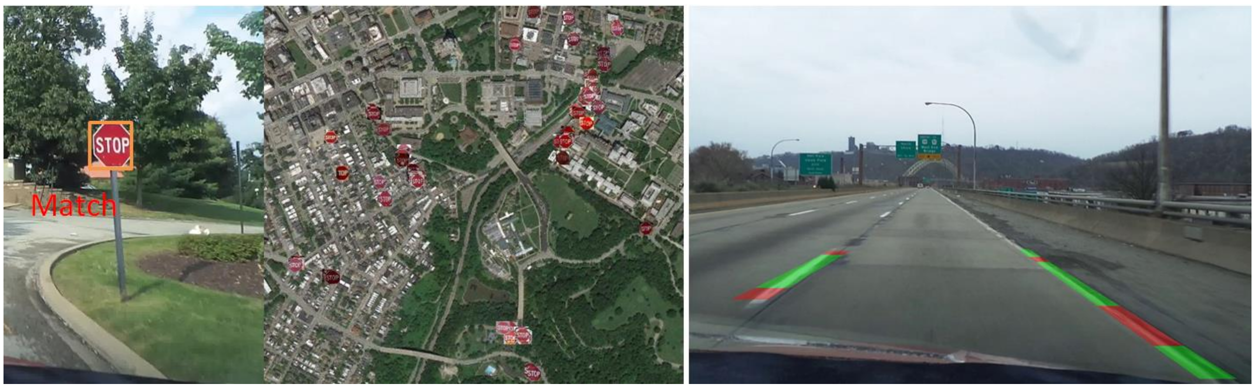

Figure 7 Left: Single detected stop sign and detected stop signs on a map. Right: Lane marking detection and assessment. Green indicates good lane marker and red indicates damaged lane marker.

Follow-on work

During this project we secured additional funding to complete the pilot project. We plan to have a first version of our system ready in the Fall to assist the city in developing their 2016 street pavement schedule. This will be followed by implementing the automatic scoring in their computing environment and tehn fully developing new capabilities like traffic sign detection.

References

[1] Examples: http://www.pavemetrics.com/, http://www.greenwood.dk/profiler.php, http://www.roadware.com/products/ARAN-Subsystems/, LCMS - Laser Crack Measurement System

[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.