RoadX: An Automated Road Inspection System

Mei Chung, Arnav Jagasia, Adele Li, Mark Lewis, Jeffrey Zhou
Mark Yim, Department of Mechanical Engineering and Applied Mechanics (MEAM)
Siddharth Deliwala, Department of Electrical and Systems Engineering (ESE)
Richard Montanez, Deputy Commissioner of Streets, City of Philadelphia
Team-1

Abstract

Road inspection is a crucial responsibility for city governments, and efficient identification and repair of road defects can improve city transportation systems, reduce accidents, and better utilize the municipal road repair budget. RoadX is an efficient road inspection system that consists of data collection devices and an enterprise platform to automatically identify and monitor defects on city roads. Mounted on city vehicles, RoadX devices collect image, GPS, and time data of road defects. The RoadX platform applies an image recognition model to classify road defects and allows government personnel to analyze and use the data for city planning. To develop RoadX, we interviewed the Philadelphia Department of Streets, and we designed RoadX based on the pain points of their current inspection workflow.

1 Motivation and Functionality

1.1 Overview of Problem and Need

Maintaining civil infrastructure and road systems is a key challenge for countries across the world. Many developed countries are struggling to address the threat of old, wearing infrastructure. In developing countries, the development of new infrastructure outpaces their ability to inspect and monitor it (Feng et al., 2017). Inefficient road inspection processes can prevent the timely detection of infrastructure defects and can lead to increased civilian accidents. Moreover, delayed identification of damaged infrastructure may require larger investments to repair, which places a strain on already-constrained infrastructure budgets. Efficiently monitoring civil infrastructure and road systems is crucial for both safety and economic development.

Many cities perform road inspection processes manually, which is both time-consuming and costly (Varadharajan et al., 2014). In fact, the City of Philadelphia relies on seven road inspectors to manually examine all of the city's roads, and the city pays about \$50,000–\$60,000 per month for this manual inspection (Montanez, 2019). A more automated road inspection process that does not rely on manual inspection could greatly improve the efficiency and accuracy of city government's road inspection processes (International Labour Organization, 2017). Through an interview with Deputy Commissioner Montanez and on-site meeting at Department of Streets, we conducted primary research to better understand the current problem. The key pain points are identified below:

Manual Inspection: We find that in the city of Philadelphia, inefficient inspection and maintenance processes for road defects area indeed an areas that need attention, especially given the old infrastructure in Philadelphia. Currently, in order to inspect the roads in the city, road inspectors manually walk all of the city's roads, visually inspecting the quality of the road and scoring the road based on their experience and judgement. It takes a year to manually inspect every road in the city.

Disparate Data Sources: On top of manual inspections conducted by inspectors in the Department of Streets, Philadelphia also relies on self-reported road defects from city residents. The city gathers this data through a form on their website, facilitated by Philly311, the city's information portal. However, these disparate sources are manually processed across inconsistent intervals, and the adhoc data reporting is not sufficient to identify all road defects in the city.

Financial Challenges: The biggest challenge to currently improving road inspection processes is the limited budget for the Department of Streets. Repairing defects costs at least \$60 million a year, not including the cost of the road inspectors themselves. As such, automating the inspection process to better use the constrained financial resources would be extremely beneficial, and the Department could reallocate some of the staff that previously performed manual inspection elsewhere for greater value.

Data Requirements: The automated inspection system needs to provide certain crucial pieces of data in order to allow users to inspect the defects. This includes (1) the type of defect, (2) an image of the defect, (3) date and time of image, and (4) the geolocation of the defect.

1.2 RoadX Value Proposition and Functionality

RoadX is an automated road inspection system for city governments which consists of vehicle-mounted image collection devices and a software platform to both classify and analyze road defect data. Specifically, our platform identifies longitudinal cracks, alligator cracks, lateral cracks, wheel mark cracks, and potholes. City governments can use RoadX to automate their annual road inspection inspection process and analyze road defect data in order to repair plans.

The RoadX data collection devices can be mounted on municipal route-based vehicles, such as garbage or sanitation trucks. These vehicles already traverse every road in the city, and with the RoadX device mounted, they can collect images of the roads on their routes. For Philadelphia specifically, the Department of Streets controls the routes of sanitation and garbage trucks, making these vehicles the ideal mount for the RoadX system. The RoadX Device would collect images of the roads at a frequency that is a function of the speed of the vehicle. In addition to the images, the devices collect GPS and timestamp data as the vehicles go on their daily routes. Renderings of the RoadX device can be seen in Appendix A, and an image of the device prototype can be seen in Appendix B.

The RoadX devices store all of the image, GPS, and timestamp data on an external USB drive. In our interviews with the Department of Streets, we learned that they strongly prefer using USB drives to transfer the data rather than any wireless upload process. This was because (1) there was no need for real-time upload capabilities and (2) they did not want to bear the costs of a 3G or 4G LTE data connection for uploading data. The USB drive can easily be removed from the device and connected to a user's computer, where the RoadX Platform's Data Upload Portal lets the user navigate their filesystem to upload the data from the USB drive.

The RoadX platform applies an image recognition model to data uploaded from the RoadX devices to detect and classify road defects and presents users with an analysis dashboard to investigate the model's findings. This dashboard allows road inspectors and city government officials to interact with all of the data from the RoadX device and use filters and map-based tools to better understand the city's road defects. The city can use the image, GPS, and timestamp data together to carry out their defect detection and repair workflow. A list of user stories that we implemented for the RoadX Platform is presented in Appendix C.

1.3 Demo

A demo video for RoadX can be found at `www.tinyurl.com/roadxvideo`. This video does give an overview of the entire RoadX system, but it specifically walks the viewer through the core features RoadX Platform and shows the image classification model in action on a newly uploaded image from a RoadX Device. As mentioned earlier, an image of the RoadX device can be found in Appendix B.

2 RoadX Business Model

2.1 Customer Segment and Stakeholders

The primary customers for RoadX are local government authorities that have jurisdiction over road maintenance. Within the United States, such responsibilities usually fall to city-level streets departments. We believe that RoadX would be an attractive solution for city-level streets departments as automating the road inspection process could directly reduce maintenance and labor costs.

In order to deploy RoadX, city information technology departments will be a crucial stakeholder. While city streets departments will be the ultimate end-user of the product, city IT departments will handle the procurement and deployment process for a new technology like RoadX.

2.2 Market Opportunity and Growth

We use the Philadelphia Department of Streets as an illustrative example of a city streets department to determine the total market size for road inspections as there is no literature on the size of this market in the United States. The Philadelphia Department of Streets spends up to $600,000 a year on road inspections. This figure is based on the lower

Components	Units	Unit Cost	Quantity	Total Cost
GPS Module	unit	2.30	1	2.30
Camera	unit	3.60	1	3.60
Microchip	unit	1.20	1	1.20
Cigarette Lighter Plug Adapter	unit	0.48	1	0.48
Polypropylene (Casing)	lb	1.72	0.0684	0.12

Table 1: Cost breakdown for manufacturing a RoadX Device at scale.

bound of the $50,000-$60,000 estimated cost of road inspections that the city bears each month (Montanez, 2019). This translates to a $106 million market in the United States (further details are presented in Appendix E). We expect the market size to have gradual, consistent growth, as our main customers are city authorities, and these cities have stable yet low growth rates.

Beyond the US market, we see potential opportunities for expanding into other markets, such as developed countries with aging infrastructure that demands inspection, maintenance, and repair (Canada, UK, Germany, Australia, Japan, etc).

2.3 Cost Analysis

With the fixed cost of the manufacturing plant aside, the cost structure only includes the cost of manufacturing our device. The device will include a microchip, camera, GPS module, casing and a mount. At scale, we estimate our product to cost about $8 per device, and the details are presented in Table 1 above. The server costs for hosting the platform depend on whether the users want to host on-premise, in which case we do not bear the server costs, or if they want a cloud solution, in which case these server costs are additionally factored into the monthly platform subscription fee.

2.4 Revenue Model

Our revenue model includes a fixed price for each RoadX device, and an enterprise SaaS model for the platform, where we incur recurring revenue from a monthly subscription fee to keep the device and software workflow running.

We price our device at $50 each. The devices are meant to be affordable and durable, so governments are encouraged to deploy multiple devices. Furthermore, our devices cannot be easily used without the platform. Thus, the affordable devices act as a gateway product for our entire solution.

The monthly subscription fee is $300/mo, as benchmarked from our platform-only competitor, DroneDeploy (DroneDeploy, 2020). Our software platform can support a fleet of devices, and it serves as a single source of truth for defect mapping for a city. If any government entities wish to host their data on the cloud, we will include the overhead for server costs into the subscription fee.

3 Related Work and Competition

3.1 Related Literature

3.1.1 Carnegie Mellon University: Computer Vision for Road Inspection

Road surface inspection is very subjective and labor intensive (Varadharajan et al., 2014). Varadharajan et. al. proposed using images collected from cameras mounted on the windshield of vehicles and computer vision to detect distress on the road surface. They segmented the ground plane and used texture, color and location information to detect road distress in the image. Varadharajan et. al are working with the City of Pittsburgh, and they hope to integrate the data into the city's existing Asset Management system. The equipment they developed costs around $1,000, which becomes prohibitively expensive for a city like Philadelphia if they want to scale to more than 50 vehicles as this approaches the current cost of manual road inspection. Moreover, this solution does not provide a system to analyze the images and leaves it to the city to process and analyze the defect data.

3.1.2 Seki Lab: Road Damage Using Deep Neural Networks

Seki Lab from the Institute of Industrial Science at The University of Tokyo published one of the original papers within road defect literature that accurately classifies the type of road damage rather than just focusing on the detection of the presence or absence of damage (Maeda et al., 2018). Researchers used an object detection method using convolutional neural networks to train the damage detection model with an original dataset. They classified eight types of road damage: longitudinal crack from wheel mark part, longitudinal crack from construction joint part, lateral crack, lateral crack from construction joint part, alligator crack, pothole, crosswalk blur, and white line blur.

In addition, the lab published a robust, open-source dataset of images for road damage. The dataset consists of 9,053 road damage images captured using a smartphone on a car. Within the dataset, there are over 15,435 instances of road surface damage examples including those captured in different weather and illuminance conditions. This dataset and model was the starting point for RoadX.

While this image detection model boasts very high accuracy, the model itself cannot be commercialized. Without a tool for analysis, this model would leave the city with a wealth of data that they would need to maintain and process independently.

3.2 Competition

According to the US Department of Transportation, there exist specialized vehicles with cameras to get distress data on roads. The cost of these vehicles, however, is prohibitively high for city government use, especially given constrained city infrastructure budgets (Miller et al., 2003).

Commercial solutions for defect inspection include products like DroneDeploy and Scopito. Both of these companies use drones to carry out the detection. Scopito is a cloud software designed to make inspection analysis faster and easier by quickly tagging and annotating images and creating personalized solutions for each individual municipal government. DroneDeploy is a platform driven by a drone mapping program to create 3D maps and analyze data in real-time. DroneDeploy has expanded to multiple applications for industries such as mining and construction and has raised over $56 million in venture capital. Both firms' products are limited to aerial image collection, and the evaluation and identification of road issues are still done manually by civil engineers. We have identified that the bottleneck of this system is not only data collection, but analysis and verification. Therefore, we designed a system that automates the entire process, from image collection to defect identification and evaluation.

Moreover, current city authorities use outdated scoring software in order to process road defects and suggest methods to repair road defects. These methods are not enough to provide a comprehensive mapping and geographic overlay of current and past defects. Additionally, none of the current products employed by city government allows for near real time analysis and a workflow to automate systematic data collection and processing.

4 RoadX Technical Approach

4.1 RoadX Device

The RoadX Device is the data collection component of the product that can be mounted on municipal vehicles. The hardware in the RoadX Device prototype consists of a camera, GPS, an external flash drive to store the image and geolocation data, and a RaspberryPi. The components are encapsulated in a custom casing that we designed, which is made of acrylic and 3D-printed PLA. The design of the device was determined by the ability to collect the necessary data, ease of assembly, and aesthetic appearance. The casing is attached to two strong magnets that can mount to any magnetic surface of the municipal vehicles.

The software for the RoadX device consists of Python scripts to automate the camera and collect GPS data on the RaspberryPi. When the RaspberryPi first turns on and the native operating system runs the boot process, we create a new process and run scripts to collect image and geolocation data from the camera and GPS module, respectively. Both the GPS and image data are saved to an external flash drive. Since the scripts begin running when the device turns on, a user only needs power on the RoadX device in order to begin data collection.

For ease of testing and debugging, we also configured our RaspberryPi to connect to a local Wi-Fi hotspot and automatically set up an SSH tunnel. This allowed us to more easily extract the data off the RoadX device from a nearby laptop connected to the same hotspot while debugging.

4.2 RoadX Platform

4.2.1 Image Recognition Model

We use the Seki lab's pre-trained road defect detection model as the core of the RoadX platform's defect detection framework (Maeda et al., 2018). This model and the training dataset are available under a Creative Commons Attribution-ShareAlike 4.0 international license, which allows others to share, modify, and use the model, even for commercial use. The model takes as input a 300x300 pixel image of a road and outputs the classifications of any defects in the image. Based on discussion with the Philadelphia Department of Streets, we only consider longitudinal crack, alligator crack, lateral crack, wheel mark crack, and pothole classifications from this model.

4.2.2 Infrastructure Analysis Platform

The RoadX Platform's goals are two-fold: first, to enable users to easily transfer data from the RoadX device to the platform, and second, to allow users directly interact with the data to perform analysis. The current version of the RoadX Platform currently uses MongoDB Atlas as the remote database, a Flask server written in Python, and a Typescript front-end built with the React framework.

The data upload portal allows users to send a compressed file of images and CSV of GPS data to the server and then queue this data for analysis. The road defect data can be displayed either as a list of entries or overlayed on a map, created with MapBox tiles. In order to display the original images from the device, if a user clicks on an entry in the platform, we fetch the analyzed image from our database and dynamically assign it a URL to view on the platform or in a new window or tab. Further details about the platform's implementation can be found in our Github repository at `www.github.com/arnavjagasia/RoadX`, and the demo video linked both in Section 1.3 and in the repository's ReadMe demonstrates these features in action.

5 Evaluation

5.1 RoadX Device Evaluation

To evaluate the RoadX Device, we outlined a set of technical requirements that the device had to meet for it to be a minimum viable product and then evaluated our current device's capabilities against these specifications. We present four evaluation metrics and our results below:

1. Physical Constraints – the device is small and light so that it can easily mount to a vehicle

2. Image Collection – the device can take and store images of sufficient quality for the image recognition model

3. Metadata Collection – the device can collect geolocation and timestamp metadata in addition to the images

4. Automation and Accessibility – the device can collect data as long as it is turned on (no other enablement needed) and the data is easily extractable from the device

After each iteration of prototyping our device, we evaluated the device on the specified criteria. From this testing, we find that our latest prototype of the device meets the evaluation specifications initially outlined.

5.1.1 Physical Constraint Evaluation

The current device measures 195mm x 75mm x 50mm and weighs 170g. Its size and weight are sufficiently small to mount to vehicles without burden.

The device is mounted with two magnets that support 5kg of vertical force each and have a dry static coefficient of friction of 0.35 against painted aluminum. The magnetic force and friction is sufficient to keep the device robustly in place even through harsh vehicle movements or vibrations and high air speed. The surface of the device that directly contacts the vehicle is coated with a thin layer of polystyrene in order to prevent any scratching of the vehicle. The device can be readily mounted, adjusted, and removed forcefully by hand. Renderings and an image of the final device are shown in Appendix A and Appendix B, respectively.

5.1.2 Image Collection Evaluation

The current version of the device uses a PiCam in order to take 300x300 pixel images every second, and we found that the quality of the PiCam is sufficient for the image recognition model.

In previous prototypes, we used a USB webcam to take higher resolution images with a wider angle of capture. We successfully set up this camera with the RaspberryPi and tested images from this camera with our image recognition model. We found that there were several critical limitations to using a USB webcam. Since the webcam connects to the RaspberryPi through a USB port, the frame rate of the camera was too slow to be mounted on a moving vehicle. Moreover, the RaspberryPi OS does not support programmatically controlling the frame rate of a device connected by USB. The PiCam is a camera that connects to the RaspberryPi's designated camera module, and the RaspberryPi OS allows a user to better configure the camera's settings. With the PiCam, we adjusted the frame rate to 1 per second in order to capture images in a moving vehicle without blur.

5.1.3 Metadata Collection Evaluation

The current version of the device uses a GPS to collect geolocation and timestamp data as the device collects images. We wrote a Python script to parse the geolocation into a human-readable format and configured the RaspberryPi to run this script with the image collection script once the RasbperryPi boots.

In earlier prototypes, we were not able to collect GPS readings indoors, and we determined this would not meet our evaluation criteria in case the municipal vehicle, on to which the RoadX device is mounted, went under tunnels or bridges and could not log metadata. To ensure that the GPS reading is collected regardless of the environment, we attached a GPS antenna and then confirmed through testing that the GPS reading is collected even when the device is covered by casing or is indoors. Our latest prototype of the device has a hole in the casing for this GPS antenna, as seen at the right side of the device in the image in Appendix B, and this revision allowed us to satisfy the metadata collection evaluation criterion.

5.1.4 Automation and Accessibility Evaluation

By executing the Python scripts in a separate process as the RaspberryPi operating system boots, the device automatically starts collecting data once it is turned on. This allows the device to collect data without needing someone to start the data collection scripts, remotely or through physical interaction with the device. The collected data gets saved on an external USB drive so that the images are easily accessible after the device is turned off. Automated inspection is a key aspect of our product – as described in our product's name (RoadX: An *Automated* Road Inspection System), and we confirmed that the data collection can be fully automated by simply powering the device on.

5.1.5 Proposed Evaluation Plans and Impact of Covid-19

In addition to these evaluation criteria, we planned to test the device on a running vehicle. Due to COVID-19, however, we could not execute the evaluation plans we had initially proposed in order to test and refine our hardware. Below is the original evaluation plan that otherwise would have been used:

- *Robustness of Hardware* – The robustness of the hardware would have been tested by mounting the device on a moving vehicle. The speed of the vehicle would be increased incrementally to the maximum speed of garbage trucks (65 mph) to test whether the device remains securely mounted.

- *Vibration* – After taking footage with the device for various speeds of the car, the images would be fed into the model to determine if the image quality is good enough for image recognition. If yes, there is no need for the vibrations to be dampened – i.e. the vibration of the device is negligible. If the image quality due to the vibrations is severely hampered, then the vibration of the device would be measured using an accelerometer. Using the vibration data, a damper would be designed inside the device to stabilize the camera and counter the vibrations.

- *Image frequency, image angle, and mount location* – We also wanted to test the impact of varying combinations of image size, frequency, and angle of the camera on the image recognition model's performance while on a moving vehicle. We would select the combination that resulted in the highest confidence levels for accurate defect classifications.

5.2 RoadX Platform Evaluation

5.2.1 Image Recognition Model

Since the image recognition model that the RoadX Platform uses is a pre-trained neural network, we wanted to evaluate its generalizability to Philadelphia streets. Specifically, we wanted to collect images of Philadelphia streets, label the images, and then assess the false negative and false positive rates of the model.

Our team took over 300 images of Philadelphia streets, primarily of roads in University City and Center City District. We made sure to take images in varying weather and lighting conditions and from a variety of angles to test the generalizability of the model. As mentioned in our proposed hardware evaluation plan, we had hoped to specifically test the angle and height of the device in a separate evaluation process. Since we needed ground truth labels for this data, we inspected each image and manually provided labels for the five types of defects our model considers.

The labeling process took between one to three minutes per image, so this was a very time intensive task for our team. Since the input to the image recognition model is a 300 x 300 pixel image, we also wrote a script in Python to resize the images we had collected to the appropriate size for the model.

In evaluating our model for the baseline false negative and false positive rates on Philadelphia data, we wrote a test loop to determine if the model could accurately identify the presence of a defect in the image. We found that it could identify correctly whether or not there was a road defect with 78.2% accuracy, and the false positive rate was 1.3% and the false negative rate was 19.2%.

We had planned to collect further images of Philadelphia roads to forward train the image recognition model and then test it again on the same test set to determine whether forward training the model could improve the false negative and false positive rates. Unfortunately, due to Covid-19, we were unable to perform this second round of image collection to carry out this task.

5.2.2 Infrastructure Analysis Platform and Dashboard

Since we interviewed the Department of Streets at the outset of this project to understand the problems and needs in road inspection, we had hoped to perform user studies with them to evaluate whether our platform solved some of the pain points they had addressed. The questionnaire that we had hoped to use for the user studies can be found in Appendix D. Due to the COVID-19 disruption, we were unable to perform these user studies. Instead, we focused on building out further features of the software platform to create an easier filtering mechanism for defect analysis. The full list of platform features can be found in Appendix E, where they are described as the user stories we used for our Agile development process.

We hoped that automating the road inspection process with RoadX would make defect discovery faster. Currently, it takes Philadelphia a full year to inspect all of the city's roads. Instead, our solution proposes mounting the device on vehicles to capture data. The Department of Streets had offered to let us test our device on a sanitation vehicle, but as we mentioned in the hardware evaluation section, we were unable to perform this any live testing due to the disruption of Covid-19. Since we evaluated that the image collection automation was successful, we are confident that a city that uses RoadX will be able to significantly speed up the amount of time it takes to perform manual inspection. Ideally, we would have also liked to confirm this as well during our live tests. Given that product was based on the needs and pain points of the Philadelphia Department of Streets, Philadelphia would be the most appropriate candidate city for an actual deployment of RoadX, and going forward, the Philadelphia Department of Streets could advocate for the solution to other cities as well.

6 Societal Impact

6.1 Value Proposition and Impact

RoadX not only streamlines and consolidates city-level road defect information into a single source of data, but it also provides city governments with an easy-to-use workflow application, which will ultimately save them time, money, and resources in the defect detection process. The price of the solution ($50 per device, $300 monthly subscription) is significantly cheaper than the current manual inspection process, which is $50,000-$60,000 per month. More so, by using an automated data collection process, RoadX effectively reduces the cost to identify each new defect.

6.2 Societal Concerns

6.2.1 Data Privacy Risks

RoadX devices can collect three types of data: images, GPS readings, and timestamps. The RoadX device does not intend to collect any personally identifiable information (PII), but we recognize that two types of PII could manifest in RoadX Devices' image data. First, license plates may be present in some of the pictures if the mount on the vehicle is sufficiently high. Some jurisdictions consider license plates as PII, and we encourage the city's data officers to apply the same retention policies to RoadX image data as they would for any Automated License Plate Reader (ALPR) platform's data. ALPR regulation is an important subdomain of data privacy legislation at the city level, so cities should treat RoadX image data with the same data privacy, security, and retention standards (American Civil Liberties Union, 2013).

Second, we understand that an individual could appear in the picture, though this would occur with very low probability due to the angle of the device's camera. Since road inspectors still have read and write access to all of the data, we recommend that the city give them full permission to remove such images from the database. If not, the RoadX image data should be held to the same data privacy, encryption, and retention standards as CCTV footage.

Ultimately, RoadX is designed to be flexible for any data privacy standards a city chooses to enforce. We choose to err on the side of giving road inspectors more control of the image data because they have full control over the current inspection process in the status quo.

6.2.2 Manipulation Risks

In order to reduce risks of manipulation or unauthorized access, the Philadelphia Department of Streets recommended that the device store data on a USB drive for transfer to an authorized computer with access to the RoadX Platform. Since an end-user conducting the data transfer can use an encrypted flash drive, we see low risk for data manipulation, even in the event that the data is lost during the transfer.

Furthermore, RoadX always includes a human road inspector in the decision making loop, and the automated process only serves to augment and not replace their abilities. With a road inspector reviewing images, we minimize the risk of wrongly identifying to existence of potholes; rather, our platform serves to augment the accuracy of the classifications.

6.2.3 Cyber Security Risks

We believe the RoadX system has low potential for security risks. We minimize any internet-based attack vectors for the RoadX Device as the device is not connected to the Internet and will remain offline. The only security risk for the RoadX Device would necessarily have to come from a hostile actor interacting directly with the physical device through its USB port, which users use to export data to the platform. We remain confident, however, that the security risk for the RoadX Device is low because the devices are mounted to city vehicles (garbage trucks and patrol cars), and tampering with any of these vehicles will have a significant legal consequence as they are city property.

The RoadX Platform has an authentication portal to limit access to registered users, and the city government could coordinate this with their current authentication procedure for other third-party procured software. Moreover, the data is property of the city, so the database that the platform connects to can be housed on-premises if the city chooses. The platform does not need an external Internet connection (the only connection would be for rendering the map in the analysis pane, but the map tiles for the city can be downloaded once and then saved on disk), so the entire application can be served locally within the city's network firewall. Given that our system does not depend on internet access, we feel confident that the platform has a low security risk.

6.2.4 Risks to Vulnerable Groups

Road defect detection processes may inherently marginalize vulnerable groups, but the RoadX platform does not introduce new risks to vulnerable groups. In road defect detection, the city can disadvantage a certain community by under-classifying defects in their community and thus conducting fewer road repairs. We hope that by using the RoadX Platform, the city will be able to more efficiently and exhaustively detect defects in all communities, possibly improving defect classification in previously under-classified areas of the city.

No matter how accurate the model, we recognize the risk that false negatives or positives might pose. For example, if one community has cobblestone roads and another community has asphalt roads, even a well-trained model might produce different distributions of classifications for the two types of roads. While the two types of roads may exhibit different structural qualities to merit these different classification distributions, the variance in the composition of roads in most major cities may lead to the risk of a high false positive rate. To address this problem, the user of the RoadX Platform – typically, an experienced road inspector from the city government – always has the option to override the RoadX Platform's automated defect classification and replace it with his or her own classification (or no classification at all, signifying the model predicted a false positive). This allows for human oversight on the defect detection process, and we believe this will help mitigate any potential risk to vulnerable groups.

6.2.5 Unintended Consequences and Mitigation

We identify one key long-term consequence from adoption of RoadX in place of the current road defect detection process: the fear that RoadX will eliminate the need for road inspectors. We do recognize that our platform will replace the manual inspection process that road inspectors regularly perform. We, however, do not believe that cities will no longer need road inspectors. Instead, we argue that road inspectors will remain an invaluable part of the road inspection process once RoadX is deployed. Road inspectors can use their subject matter expertise to operate on the RoadX Platform, where they can carry out analysis and oversight. We believe that any city that adopts our device should reallocate road inspectors from manually inspecting every road to a more analysis-focused role, where they can more efficiently apply their domain knowledge to understanding the state of the city's roads.

7 Discussion

Despite the disruption in our prototyping and evaluation process posed by Covid-19, our final prototype of the RoadX system takes a strong first step in tackling some of the inefficiencies of manual road inspection. Based on the evaluation of the RoadX device that we were able to complete, we found that we were able to construct a lightweight, small device that could easily automate image and metadata collection. Mounted to a fleet of vehicles that collectively traverse the city's roads each week, a system of these devices would collect image data each week that would have otherwise taken a full year of manual inspection. While we would have liked to carry out further benchmarking and evaluation, this system collects data in a faster and more cost-effective approach than the current manual inspection process.

Our baseline evaluation indicates that the current image recognition model may be able to generalize to Philadelphia with forward training on Philadelphia data. This, however, would require hundreds of tagged image samples. Throughout this project, we learned that given the novel, automated approach for road inspection that RoadX employs, no one was able to provide us with the quality or quantity of labeled image data that we needed improve the accuracy of our identification model. This led us to spend countless hours manually tagging photos. As increasingly more cities adopt automated approaches to road maintenance, cities can start to compile a robust set of training data. The first few cities to embark on this undertaking, however, will have to spend time collecting enough data to forward train the automated system to U.S. road infrastructure standards and even further to better handle any idiosyncrasies of their city's roads. Despite these challenges, we were able to show that the image recognition model can quickly classify road defects in images and that users can easily interact with this classified data for analysis. In the future, we could also work with city governments to identify the defects by priority or cost, thus providing officials with a simpler, outlined repair workflow.

If we were to continue this project, or if any city were interested in continuing such work, the city would need to commit time and resources to setting up automated data collection. Once the evaluations that got disrupted due to Covid-19 could be carried out to fine tune the mounted devices, the RoadX Devices would be ready to use. The compute and storage capabilities of the RoadX Platform would necessarily need to be scaled up to meet the demands of the city's proposed road inspection plan. Even despite the obstacles we faced in fully completing our proposed evaluation, RoadX proves that road inspection processes need not be manual, and cities can achieve for faster and more-cost effective inspection processes through an automated road inspection system.

Acknowledgments

We would like to thank Professor Ani Nenkova, Dr. Sangeeta Vohra, and all of the students of CIS400/CIS401 for their feedback and guidance over this past year. We would like to especially thank Professor Mark Yim and Mr. Sid Deliwala, who both consulted with us on this project and our approach from our first brainstorming sessions in the summer of 2019. We would also like to express our gratitude and thanks to Mr. Richard Montanez, Deputy Commissioner of Streets, and the entire Department of Streets for discussing with us their work, the road inspection process in the city of Philadelphia, and how we could best focus our work to improve this process.

References

American Civil Liberties Union. 2013. You are being tracked: How license plate readers are being used to record americans' movements. Technical report, American Civil Liberties Union.

DroneDeploy. 2020. DroneDeploy Pricing Plans. https://www.dronedeploy.com/pricing.html. Accessed: 2020-04-26.

Chen Feng, Ming-Yu Liu, Chieh-Chi Kao, and Teng-Yok Lee. 2017. Deep active learning for civil infrastructure defect detection and classification. In *Computing in civil engineering 2017*, pages 298–306.

International Labour Organization. 2017. *Managing Maintenance of Rural Roads in India*. National Rural Roads Development Agency, Ministry of Rural Development, Government of India, New Delhi, India.

Hiroya Maeda, Yoshihide Sekimoto, Toshikazu Seto, Takehiro Kashiyama, and Hiroshi Omata. 2018. Road damage detection and classification using deep neural networks with smartphone images. *Computer-Aided Civil and Infrastructure Engineering*.

John S Miller, William Y Bellinger, et al. 2003. Distress identification manual for the long-term pavement performance program. Technical report, United States. Federal Highway Administration. Office of Infrastructure.

Richard Montanez. 2019. Interview with Richard Montanez on Philadelphia Road Inspection Process.

Srivatsan Varadharajan, Sobhagya Jose, Karan Sharma, Lars Wander, and Christoph Mertz. 2014. Vision for road inspection. In *IEEE Winter Conference on Applications of Computer Vision*, pages 115–122. IEEE.

Appendix

A RoadX Device: Renderings

Figure 1: Rendering of RoadX Device

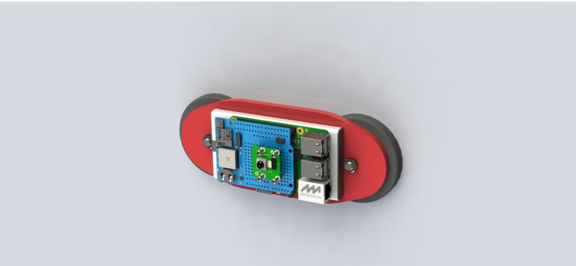

Figure 2: Rendering of RoadX Device Without Casing

B RoadX Device: Image

Figure 3: Image of RoadX Device Prototype

C List of RoadX Platform User Stories

Authentication User Stories

1. A user can log in and log-out of the platform.

Data Upload Workflow User Stories

1. A user can register a new device with the platform.

2. A user can upload image and GPS files from file-system.

3. A user can queue new data for automated defect discovery.

Analysis User Stories

1. A user can view the geographic distribution of road defects on a map.

2. A user user can view the predicted classification(s) for each entry.

3. In order to better understand the model's predictions, a user user can view the confidence with which the model predicted that there was a defect.

4. In order to better understand the model's predictions, a user can view the original image that the RoadX Device took of the road, with a bounding box, label, and confidence over the defect.

5. A user can zoom into the image in sufficient detail or in a full-screen window to assess the model's prediction.

6. A user can see the time the data was uploaded and the time the image was taken to understand the history of the defect based on the image collection data.

7. A user can override the model's findings with a manual override. Future users should see that the updated classification was because of a manual override but still see what the model had predicted, if they want.

8. A user can filter data on the map and the list by the mode's confidence in the classification.

9. A user can filter the data on the map and the list to any subset of defect categories.

Non-User Story Technical Features

1. Batch upload files as a zip file and then unzip them on the server to limit network traffic

2. If the automated discovery process found no defects, remove the image from the main datastore to conserve storage in the prototype platform.

D User Study Evaluation Form

We wanted to conduct user studies with staff members at the Department of Streets, who we had interviewed in the fall. We had intended to use our device to collect data from Philadelphia and then let staff members at the department upload data and conduct basic analysis. We would have used the following questionnaire to understand their views about RoadX. The questionnaire contains a mix of quantitative scoring and qualitative questions.

Questionnaire

1. On a scale of 1-5, how easy was it to upload data from the RoadX device to the platform (1 = Difficult, 5 = Easy)?

2. On a scale of 1-5, how easy is it to use the Map analytical view (1 = Difficult, 5 = Easy)?

3. On a scale of 1-5, how easy is it to use the List analytical view (1 = Difficult, 5 = Easy)?

4. On a scale of 1-5, how useful are the toolbar and filter features(1 = Not useful, 5 = Useful)?

5. On a scale of 1-5, do you feel that the RoadX platform improves your current workflow? (1 = Strongly Disagree, 5 = Strongly Agree)

6. What do you like about the RoadX platform?

7. What do you feel is the most important feature of RoadX?

8. What are the biggest pain points you face conducting analysis of defects, even using RoadX?

9. What would you improve about RoadX?

10. Was there anything missing from this product that you expected?

E Market Size in United States

We estimate that the US Market Size is $100 million. Our target market is every US metro area with population over 100,000, and the total population in this target market is 279,330,000. The cost of inspection in Philadelphia is $600,000 each year, and Philadelphia has a population of 1,581,000.

We assume that the cost of road inspection scales linearly with the number of roads and that the population density of Philadelphia is similar to those of other cities our target market across the U.S. We excluded U.S. cities with population less than 100,000 in our market sizing as we want to first focus deployment of RoadX to developed urban centers.

We see that Philadelphia represents $1,581,000/279,330,000 = 0.566\%$ of the entire target market in the U.S. Since Philadelphia spends $600,000 each year on road inspections, we predict the total U.S. market spends $\frac{\$600,000}{0.00566} \approx \$106,000,000$ on the administration behind repairing road defects.