

---

# SpotholeAI - An Artificial Intelligence Assistant to fix Potholes

---

Vishal Bhalla<sup>1</sup>

## Abstract

Currently, pothole fixing approaches adopt manual and time consuming efforts with lot of human intervention at multiple levels. There exists mobile apps which do detect potholes but that is for a driver to avoid them whereas other apps which help in fixing local civic issues, do not incorporate any Artificial Intelligence (AI) to automate the process and reap benefits of time and efficiency. Given the scale of complaints daily, there is a pressing need of an intelligent and efficient way to automate the process. SpotholeAI detects potholes from any user uploaded image-complaint to predict its severity, the raw materials required and the estimated time to fix it, by incorporating cutting edge AI research, primarily Deep Learning for Semantic Image Segmentation methods. With life saving applications that manifest the use of AI for social good, SpotholeAI intends to reduce the estimated time to report and fix not just potholes but also other civic issues by 10-fold, thereby bringing accountability and transparency in the civic sector.

## 1. Motivation

At present, the pothole fixing process is manual and time consuming and requires lot of human intervention at multiple levels. A person has to register a complaint with detailed description of the issue along with its location (Halifax, 2016). The complaint is then analyzed and assigned to a ward officer. Multiple follow-ups are needed that too in person and, at times, the issue is raised on social media or the help of a local corporator (elected representative) is used to resolve it.

There exists mobile apps which do detect potholes (Song et al., 2018; Agarwal et al., 2018; Gersztyn, 2016; Madli et al., 2015; Lin & Liu, 2010) but that is for a driver to avoid

them. Other apps (Foster, 2017; Foth et al., 2011; Mednis et al., 2011; Eriksson et al., 2008) which help in fixing local civic issues, do not incorporate any AI techniques to automate the process, or cover broad road defects (Chacra & Zelek, 2017; Maeda et al., 2018). Often times, such systems also struggle to identify fake complaints, and some time is wasted in discarding these cases. Holistically, all of these existing systems require a lot of human discretion that too at multiple stages in the process of fixing potholes; be it for reporting, identifying an action plan to actually fixing it. Given the scale of complaints daily, there is a pressing need of an intelligent and efficient way to automate the process which SpotholeAI addresses.

## 2. Goals

SpotholeAI targets the United Nations (UN) Sustainable Development Goals (SDG) of:

- Goal 9: Build resilient infrastructure, promote sustainable industrialization and foster innovation
- Goal 11: Make cities inclusive, safe, resilient and sustainable

## 3. Proposal

A user clicks an image of the civic issue (say pothole) with his mobile camera and uploads a complaint which is automatically geo-tagged on the SpotholeAI mobile app. SpotholeAI then detects potholes in the image, predicts its severity, the raw materials required and an estimated time to fix it by incorporating cutting edge AI research, primarily Deep Learning methods and sends the complaint to the municipal authority for fixing it. The end-user gets an in-app notification to track the progress of the complaint while the municipal authority can view all the complaints in a dashboard with an optimized route plan for the day to take their truck and fix these issues. Similarly, other civic issues can also be diagnosed from image complaints.

SpotholeAI will assist in the ability to identify potholes and prioritize fixing them with the appropriate raw materials in a speedy manner. This will help automate the plethora of complaints each day at the municipal corporation and also optimize its route plan to fix them (Dueker & Fischer,

---

<sup>1</sup>Department of Informatics, Technical University of Munich, Munich, Germany. Correspondence to: Vishal Bhalla <vishal.bhalla@tum.de>.

2003; Marasteanu et al., 2018; Srivastava et al., 2018) on priority. For the end-user, it would mean less traffic on roads, reduced damage to vehicles and overall happy drive and satisfaction with the work of the civic authorities (Waze, 2019). With the use of AI, SpotholeAI intends to reduce the estimated time to report and fix not just potholes but also other civic issues by 10-fold.

#### 4. Impact

With early reporting and fixing of potholes, this product has a life saving application. Pothole-free roads will lead to reduction in road accidents and less damage to both human and car (AAA-Oregon-Idaho, 2016; Dash). SpotholeAI tends to create (s)heroes - publishing their good and proactive work to help resolve civic issues at the earliest on social media, radio and print. Incentives are given to the supervising officer for fixing potholes in a ward with no recurring complaints and also to proactive users who report them.

The goal is to bring transparency in the whole process and churn out more accountability within the civic body and the various stakeholders involved in road construction and maintenance. This will in turn have a reverberating effect to instill more civic sense in people who proactively report (Waze, 2019) as well as responsibility from civic authorities to act upon complaints and fix them in a timely fashion. This citizen engagement will help build and sustain smart cities (Thompson, 2016).

#### 5. Evaluation

An opportunity to conduct a Proof Of Concept (POC) with the local municipal body would be an excellent way to test the proposal and fine tune the product before it is scaled further. The initial experiments are necessary to identify which civic issue needs priority and also important for example, to come up with the right severity levels and raw materials for the pothole fix. Once a standard is established, this template can then be used not just across local bodies within a region, but globally. All of this would only be possible in close co-operation with the civic authorities by seeking domain expertise from road contractors at a smaller-scale POC environment through direct interactions and site visits. Success can then be measured as the end to end turnaround time from reporting to fixing the civic issue at hand.

#### 6. Risks

The initial predictions from the model might not be accurate and hence the app will take in feedback from the visiting supervisor to correctly provide the actual on-ground severity, so that the model learns and better its predictions.

Another risk is that the government or municipal author-

ities do not buy this product as it brings in an additional transparency factor or it could be that they are actually overwhelmed with a lot of complaints but want something specific to another civic issue say for example garbage disposal on priority instead of pothole fixing.

Moreover, the model should be robust to handle image blur. Mounting the smart-phone on the dashboard, a video upload option can be used while driving (Maeda et al., 2018), with the model still using some frames in it to detect the potholes. Also, one needs to blur any user identity - a persons face or a car license plate to account for privacy (Maeda et al., 2018).

#### 7. Data

At present, there is no existing dataset (OpenDataPortal, 2018) of pothole images and hence other datasets (Everingham et al., 2015; Cordts et al., 2016; Zhou et al., 2017) that can be used in a transfer learning setting to train the models must be explored. It would be difficult for the city corporation to provide with the photos initially, so stock images of potholes <sup>1</sup> or creative commons <sup>2</sup> can help build a classifier to detect potholes. However, lacking any meta-data on the labels it can't be used in its entirety. Hence, collecting data for this use-case is a must and can be worked out after a tie-up with the local civic body. In order to keep the proposed data collection approach simple and easy to adapt for government bodies and teams, road contractors will provide with pictures of on-site repair, the raw materials used and time taken for fixing the pothole, to bootstrap the data labelling process through a mobile interface which has a photo upload option, drop-down button for catalogue of raw materials along with a corresponding field to enter the quantity, and, to note the time taken in hours, creating a database and annotation tool similar to LabelMe (Russell et al., 2008). This will be a one time activity and going forward more data will be accumulated as the app gets a widespread reach and more users upload image complaints.

Furthermore, this initial collected data can be scaled with the help of standalone platform independent libraries for Image Augmentation like Augmentor <sup>3</sup> to train the underlying Deep Learning models to be more robust. The vision is to also create a pothole dataset on the lines of SYNTHIA (Ros et al., 2016) - a large-scale synthetic dataset for semantic segmentation of urban scenes and Kitti Road dataset (Fritsch et al., 2013), which will also be open-sourced and made available online for incremental research. To avoid privacy concerns, prior approval from the local body would be taken.

<sup>1</sup><https://www.shutterstock.com/search/pothole>

<sup>2</sup><https://search.creativecommons.org/search?q=pothole>

<sup>3</sup><https://github.com/mdbloice/Augmentor>

## 8. Labels

At the crux of SpatholeAI, a pothole image is fed to a Deep Learning based Semantic Image Segmentation model that assigns a semantic label to every pixel in the image. For the app to train its underlying model and give accurate predictions, one has to collect annotated image data wherein each image is firstly segmented to highlight the area of interest (say potholes) along with multi-label information on its severity, the raw materials required as well as the estimated time to fix it.

A good starting point is to do Semantic Segmentation<sup>4</sup> wherein one labels the pixels of a pothole in images using a Fully Convolutional Network (FCN). Other latest work in Curve-GCN (Ling et al., 2019) that utilize semi-supervised approaches for efficient interactive annotation can also be used. Another option is to label the collected image data with the help of crowd-sourcing but unfortunately that would require domain expertise i.e. someone from the civic body as well as a road contractor.

In order to offset the annotation cost, labeling will constitute both coarse (only object bounding box labels) and fine-grained (having semantic segmentation labels and box labels). The underlying semantic image segmentation model can then use these coarse image-level annotations to train and infer the multi-label outputs in a weakly or semi-supervised setting. And, the resulting dataset will have a small set of fully labeled and a large set of weakly labeled images.

## 9. Social System

As a speaker at one of the Machine Learning talks at a Meetup event, the author found a mobile app developer who designed a hard-coded first prototype that works for 3 sample image complaints (based on severity of potholes) on an android phone. A pilot study was done on 10 users who gave feedback to publicize the app and include incentives for reporting issues.

To scale further, a computer vision partner to advise on research of Semantic Image Segmentation models and an investor to support this idea are absolutely essential. Likewise, collaborators with a varied skill set (mobile app developers, deep learning researchers, sales, etc.), annotated data collection, pitching and networking with stakeholders are equally important to build this together!

From an infrastructure point of view, one of the primary needs to train and experiment the underlying Deep Learning based models would be compute i.e. GPU/TPU - wherein one can also leverage Deep Learning on cloud for fast train-

ing. I will also require storage for the image-complaints for any future analytics as well as fast inference for real-time complaints on mobile.

## 10. Technical System

Although semantic image segmentation models have been extensively used to detect various objects in images, it has not been applied to a micro-level problem of detecting potholes or to resolve civic issues in government bodies. The research here will primarily involve Transfer Learning of Deep Learning based models for Image Segmentation from Autonomous Driving to the pothole scenario.

In particular, related research work on Deep Learning models for Semantic Image Segmentation that can be applied in this use-case include DeepLab (Chen et al., 2018; Liu et al., 2019), Fully Convolutional Network-Based approaches (Chen et al., 2014; Long et al., 2015; Poudel et al., 2019) and Region-Based Semantic Segmentation (Zhao et al., 2017; Girshick et al., 2018). Moreover, Data Augmentation libraries will help scale the few pothole pictures to begin with, in a Weakly and Semi-Supervised Semantic Segmentation (Papandreou et al., 2015; Khoreva et al., 2017; Ibrahim et al., 2018; Li et al., 2018; Song et al., 2019; Lee et al., 2019), as well as, Few-Shot learning (Shaban et al., 2017; Siam & Oreshkin, 2019; Hu et al., 2019; Wu et al., 2019) setting where small amount of coarse image-level annotations can be utilized to train models.

## 11. Extensions

The proposed AI-infused tool for managing urban potholes, evaluates their severity based on a user-uploaded photo, then proposes an optimized route in order to send municipal services to fix them. Moreover, with near-about the same technical setup, resolution of other civic issues like garbage disposal, tree fall, choked drains, water pipe burst, etc. can also be addressed. The pothole fixing scenario extends to highways, airport runways for safe landing and in defense for rescue and emergency operations wherein a drone sweeps the area to get an image scan, or can even leverage the availability of high-resolution satellite imagery (Bischke et al., 2017) instead of a user's mobile camera. A future modality is to have under-belly cameras installed already in a vehicle or public transport like buses, that will give live daily road status, promoting proactive repair before the potholes are formed.

## Acknowledgements

The author would like to wholeheartedly thank all the anonymous reviewers for their critical feedback which has helped streamline this research proposal and Mayur Rokade for the

<sup>4</sup><https://github.com/udacity/CarND-Semantic-Segmentation/>

development of the sample mobile app prototype.

## References

- AAA-Oregon-Idaho. Pothole damage costs u.s. drivers \$3 billion annually. <https://www.oregon.aaa.com/2016/02/pothole-damage-costs-u-s-drivers-3-billion-annually/>, 2016.
- Agarwal, S., Mittal, N., and Sureka, A. Potholes and bad road conditions: mining twitter to extract information on killer roads. In *Proceedings of the ACM India Joint International Conference on Data Science and Management of Data*, pp. 67–77. ACM, 2018.
- Bischke, B., Helber, P., Folz, J., Borth, D., and Dengel, A. Multi-task learning for segmentation of building footprints with deep neural networks. *arXiv preprint arXiv:1709.05932*, 2017.
- Chacra, D. B. A. and Zelek, J. S. Fully automated road defect detection using street view images. In *2017 14th Conference on Computer and Robot Vision (CRV)*, pp. 353–360. IEEE, 2017.
- Chen, L.-C., Papandreou, G., Kokkinos, I., Murphy, K., and Yuille, A. L. Semantic image segmentation with deep convolutional nets and fully connected crfs. *arXiv preprint arXiv:1412.7062*, 2014.
- Chen, L.-C., Papandreou, G., Kokkinos, I., Murphy, K., and Yuille, A. L. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2018.
- Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., and Schiele, B. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3213–3223, 2016.
- Dash, D. K. Bad roads killed over 10k people in 2015; 3,416 deaths due to potholes. <https://timesofindia.indiatimes.com/india/Bad-roads-killed-over-10k-people-in-2015-3416-deaths-due-to-potholes/articleshow/53482615.cms>.
- Dueker, M. J. and Fischer, A. M. Fixing swiss potholes: The importance and cyclical nature of improvements. *Economics Letters*, 79(3):409–415, 2003.
- Eriksson, J., Girod, L., Hull, B., Newton, R., Madden, S., and Balakrishnan, H. The pothole patrol: using a mobile sensor network for road surface monitoring. In *Proceedings of the 6th international conference on Mobile systems, applications, and services*, pp. 29–39. ACM, 2008.
- Everingham, M., Eslami, S. A., Van Gool, L., Williams, C. K., Winn, J., and Zisserman, A. The pascal visual object classes challenge: A retrospective. *International journal of computer vision*, 111(1):98–136, 2015.
- Foster, J. Listen: Montreal gives pot holes gps data you can track. <http://www.iheartradio.ca/cjad/news/listen-montreal-gives-pot-holes-gps-data-you-can-track-1.2536533>, 2017.
- Foth, M., Schroeter, R., and Anastasiu, I. Fixing the city one photo at a time: mobile logging of maintenance requests. In *Proceedings of the 23rd Australian computer-Human interaction conference*, pp. 126–129. ACM, 2011.
- Fritsch, J., Kuehnl, T., and Geiger, A. A new performance measure and evaluation benchmark for road detection algorithms. In *International Conference on Intelligent Transportation Systems (ITSC)*, 2013.
- Gersztyn, J. Pothole pal interactive infrastructure tracking app. <https://digitalcommons.wou.edu/cgi/viewcontent.cgi?referer=&httpsredir=1&referer=https://www.google.com/&httpsredir=1&article=1041&context=aes>, 2016.
- Girshick, R., Radosavovic, I., Gkioxari, G., Dollár, P., and He, K. Detectron. <https://github.com/facebookresearch/detectron>, 2018.
- Halifax. Street - report a pothole or sinkhole. <https://www.halifax.ca/transportation/streets-sidewalks/paving-repair/report-a-pothole>, 2016.
- Hu, T., Yang, P., Zhang, C., Yu, G., Mu, Y., and Snoek, C. G. Attention-based multi-context guiding for few-shot semantic segmentation. 2019.
- Ibrahim, M. S., Vahdat, A., and Macready, W. G. Weakly supervised semantic image segmentation with self-correcting networks. *arXiv preprint arXiv:1811.07073*, 2018.
- Khoreva, A., Benenson, R., Hosang, J., Hein, M., and Schiele, B. Simple does it: Weakly supervised instance and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 876–885, 2017.
- Lee, J., Kim, E., Lee, S., Lee, J., and Yoon, S. Ficklenet: Weakly and semi-supervised semantic image segmentation using stochastic inference. *arXiv preprint arXiv:1902.10421*, 2019.



- Li, Q., Arnab, A., and Torr, P. H. Weakly-and semi-supervised panoptic segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 102–118, 2018.
- Lin, J. and Liu, Y. Potholes detection based on svm in the pavement distress image. In *2010 Ninth International Symposium on Distributed Computing and Applications to Business, Engineering and Science*, pp. 544–547. IEEE, 2010.
- Ling, H., Gao, J., Kar, A., Chen, W., and Fidler, S. Fast interactive object annotation with curve-gcn. In *CVPR*, 2019.
- Liu, C., Chen, L.-C., Schroff, F., Adam, H., Hua, W., Yuille, A., and Fei-Fei, L. Auto-deeplab: Hierarchical neural architecture search for semantic image segmentation. *arXiv preprint arXiv:1901.02985*, 2019.
- Long, J., Shelhamer, E., and Darrell, T. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3431–3440, 2015.
- Madli, R., Hebbar, S., Pattar, P., and Golla, V. Automatic detection and notification of potholes and humps on roads to aid drivers. *IEEE sensors journal*, 15(8):4313–4318, 2015.
- Maeda, H., Sekimoto, Y., Seto, T., Kashiya, T., and Omata, H. Road damage detection using deep neural networks with images captured through a smartphone. *arXiv preprint arXiv:1801.09454*, 2018.
- Marasteanu, M., Ghosh, D., Turo, M., Hartman, M., Milavitz, R., Le, J.-L., et al. Pothole prevention and innovative repair. Technical report, Minnesota. Dept. of Transportation. Research Services & Library, 2018.
- Mednis, A., Strazdins, G., Zviedris, R., Kanonirs, G., and Selavo, L. Real time pothole detection using android smartphones with accelerometers. In *2011 International conference on distributed computing in sensor systems and workshops (DCOSS)*, pp. 1–6. IEEE, 2011.
- OpenDataPortal, M. Road rehabilitation by mechanized pothole filling. <http://donnees.ville.montreal.qc.ca/dataset/refection-de-chaussee-par-remplissage-mecanise-de-nid-de-poule>, 2018.
- Papandreou, G., Chen, L.-C., Murphy, K. P., and Yuille, A. L. Weakly-and semi-supervised learning of a deep convolutional network for semantic image segmentation. In *Proceedings of the IEEE international conference on computer vision*, pp. 1742–1750, 2015.
- Poudel, R. P., Liwicki, S., and Cipolla, R. Fast-scnn: Fast semantic segmentation network. *arXiv preprint arXiv:1902.04502*, 2019.
- Ros, G., Sellart, L., Materzynska, J., Vazquez, D., and Lopez, A. M. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3234–3243, 2016.
- Russell, B. C., Torralba, A., Murphy, K. P., and Freeman, W. T. Labelme: a database and web-based tool for image annotation. *International journal of computer vision*, 77(1-3):157–173, 2008.
- Shaban, A., Bansal, S., Liu, Z., Essa, I., and Boots, B. One-shot learning for semantic segmentation. *arXiv preprint arXiv:1709.03410*, 2017.
- Siam, M. and Oreshkin, B. Adaptive masked weight imprinting for few-shot segmentation. *arXiv preprint arXiv:1902.11123*, 2019.
- Song, C., Huang, Y., Ouyang, W., and Wang, L. Box-driven class-wise region masking and filling rate guided loss for weakly supervised semantic segmentation. *arXiv preprint arXiv:1904.11693*, 2019.
- Song, H., Baek, K., and Byun, Y. Pothole detection using machine learning. *Advanced Science and Technology*, pp. 151–155, 2018.
- Srivastava, S., Sharma, A., and Balot, H. Analysis and improvements on current pothole detection techniques. In *2018 International Conference on Smart Computing and Electronic Enterprise (ICSCEE)*, pp. 1–4. IEEE, 2018.
- Thompson, E. M. What makes a city smart? *International Journal of Architectural Computing*, 14(4):358–371, 2016.
- Waze. Connected citizens program (ccp). <https://www.waze.com/ccp>, 2019.
- Wu, Z., Wang, X., Gonzalez, J. E., Goldstein, T., and Davis, L. S. Ace: Adapting to changing environments for semantic segmentation. *arXiv preprint arXiv:1904.06268*, 2019.
- Zhao, H., Shi, J., Qi, X., Wang, X., and Jia, J. Pyramid scene parsing network. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2881–2890, 2017.
- Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., and Torralba, A. Scene parsing through ade20k dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 633–641, 2017.

## A. SpotholeAI Mobile App Screens

The workflow for any user uploaded complaint is depicted in the SpotholeAI mobile app screens of [fig. 1](#) and the corresponding flow from the municipal authority is shown in the app screens of [fig. 2](#)







Screen 1 (After login)	Screen 2 (After upload)	Screen 3 (After pothole detection on central repository)	Screen 4 (After approval from municipal authority)	Screen 5 (After actual fix by municipal truck)
 <div>Upload</div> <p><b>Note:</b> GPS co-ordinates are captured in the background.</p>	<div>GPS</div> <div>24.1403° N, 102.7112° E</div> <div>Image_Id</div> <div>IMG_239421</div> <div>Date of Report</div> <div>09/04/2018</div> <div>Complaint_Id</div> <div>PH_024301</div> <div>Status</div> <div>Image uploaded successfully </div>	<div>GPS</div> <div>24.1403° N, 102.7112° E</div> <div>Image_Id</div> <div>IMG_239421</div> <div>Date of Report</div> <div>09/04/2018</div> <div>Complaint_Id</div> <div>PH_024301</div> <div>Pothole Severity</div> <div>Level 2</div> <div>Materials</div> <div>Bricks, sand &amp; cement</div> <div>Estimated Time to Fix (ETF)</div> <div>3 days</div> <div>Status</div> <div>Sent to Municipal Authority  </div>	<div>GPS</div> <div>24.1403° N, 102.7112° E</div> <div>Image_Id</div> <div>IMG_239421</div> <div>Date of Report</div> <div>09/04/2018</div> <div>Complaint_Id</div> <div>PH_024301</div> <div>Pothole Severity</div> <div>Level 2</div> <div>Materials</div> <div>Bricks, sand &amp; cement</div> <div>Jurisdiction</div> <div>Ward 12</div> <div>Supervising Officer</div> <div>Mrs. Sulochana Bhosale</div> <div>Estimated Time to Fix (ETF)</div> <div>Scheduled tomorrow</div> <div>Status</div> <div>Road repair </div>	<div>GPS</div> <div>24.1403° N, 102.7112° E</div> <div>Image_Id</div> <div>IMG_239421</div> <div>Date of Report</div> <div>09/04/2018</div> <div>Complaint_Id</div> <div>PH_024301</div> <div>Pothole Severity</div> <div>Level 2</div> <div>Materials</div> <div>Bricks, sand &amp; cement</div> <div>Jurisdiction</div> <div>Ward 12</div> <div>Supervising Officer</div> <div>Mrs. Sulochana Bhosale</div> <div>Status</div> <div>Pothole fixed </div>

Figure 1. The user scenario of SpotholeAI app is depicted above. Firstly, a user clicks a pothole image as in the first screen to the left, and uploads this image-complaint on the app (with it's geo-location in the background) as shown in the second screen from the left. This complaint is then sent to the SpotholeAI central repository which returns the severity, raw materials required, and, the Estimated Time to Fix (ETF) using advanced AI based techniques shown in the third screen. This data is then sent to the municipal authority (fourth screen from the left) that assigns a supervising municipal officer, makes any corrections at its end and updates the progress of the repair work. In the last screen, the user is notified of any status update for the reported issue.

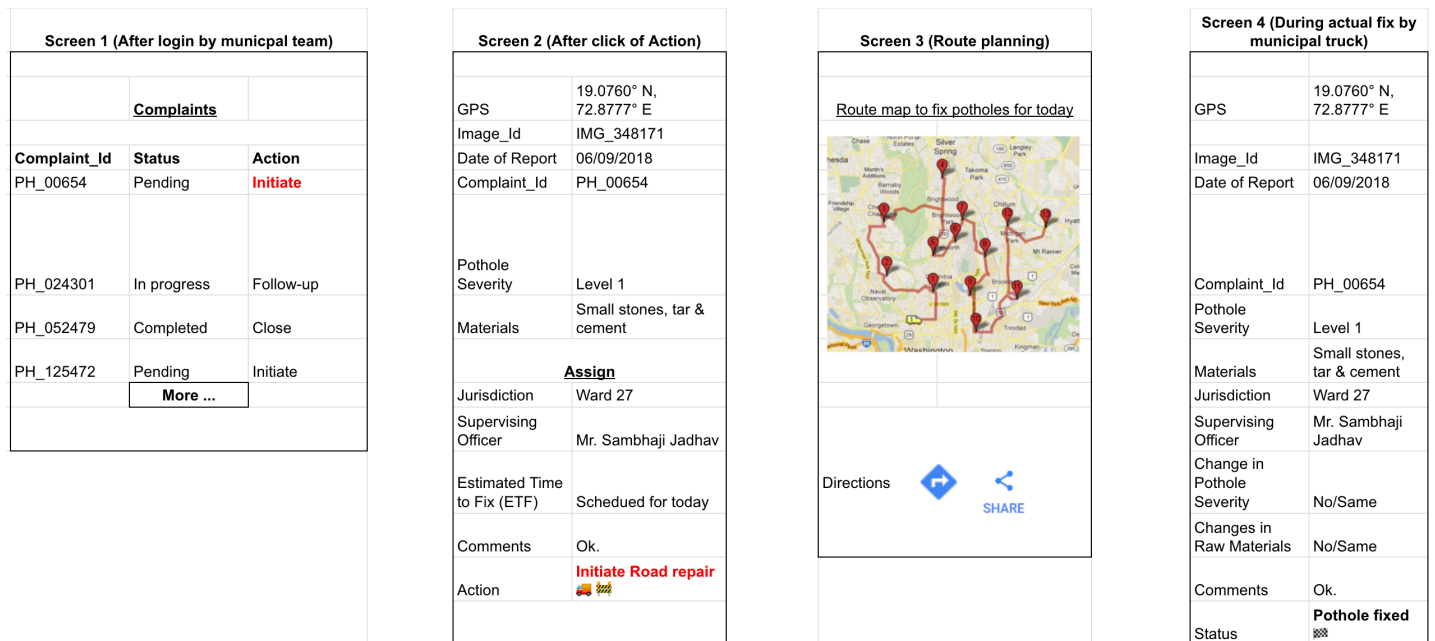


Figure 2. The figure above is representative mobile app screens for a municipal ward supervisor. A dashboard in screen 1, allows to track all the reported pothole image-complaints in an area. The next screen shows the status of one such complaint - it's severity and all other related details. Additionally, SpotholeAI provides an optimized daily route plan, based on priority, to send the raw-materials truck to repair potholes as can be seen in the third screen. Lastly, it also enables the supervising officer with an option to provide feedback on any discrepancies in the actual on-ground severity of the issue.