SpotholeAI - An Artificial Intelligence (AI) assistant to fix Potholes.

Vishal Bhalla

April 23, 2019

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. 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. Often times, such systems also struggle to identify fake complaints, and some time is spent in discarding such cases. On a whole, all of these existing systems require a lot of human discretion that too at multiple stages in the process of fixing potholes - be it reporting, identifying an action plan to actually fixing the issue. Given the scale of complaints daily, there is a pressing need of an intelligent and efficient way to automate the process which SpotholeAI addresses. 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, bringing accountability and transparency in the civic sector.

- **Topics**: SpotholeAI targets the below UN Sustainable Development Goals:
  - Goal 9: Build resilient infrastructure, promote sustainable industrialization and foster innovation
  - Goal 11: Make cities inclusive, safe, resilient and sustainable
- **Problem**:

  Currently, pothole fixing approaches adopt manual and time consuming efforts with lot of human intervention at multiple levels. A person has to register a complaint with detailed description of the issue along with its location [1]. 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 [2, 3, 4, 5, 6] but that is for a driver to avoid them. Other apps [7, 8, 9, 10] which help in fixing local civic issues, do not incorporate any Artificial Intelligence (AI) to automate the process, or cover broad road defects[11, 12]. Often times, such systems also struggle to identify fake complaints, and some time is spent in discarding such 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 reporting, identifying 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.

- **Proposal**:

  A user clicks an image of the civic issue (say pothole) with his mobile camera and uploads a complaint i.e. 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 help automate the plethora of complaints each day at the municipal corporation and also optimize its route plan to fix them [13, 14, 15]. 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 [16].

  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.

- **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 [17, 18].

  With SpotholeAI, we tend 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.

  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 [16] as well as responsibility from civic authorities to act upon complaints and fix them in a timely fashion. This citizen engagement will help build smart cities [19].

- **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 issue.

- **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 betters its predictions.

  Another risk is that the government or municipal authorities do not buy this product as it brings in an additional transparency factor or 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.

  The model should be robust to handle image blur. Mounting the smartphone on the dashboard, a video upload option can be used while driving [12], with the model still using some frames in it to detect the potholes.

  Also, need to blur any user identity - a person's face or a car license plate to account for privacy [12].

- **Data**:

  At present, there is no existing dataset [20] of potholes images and hence other datasets [21, 22, 23] that can be used in a transfer learning setting to train the models must be explored. However, collecting data for this use-case is a must and can be worked out after a tie-up with the local civic body wherein civil contractors need to be trained to click pictures of their current pothole fixing work that will be stored on a central repository or cloud. 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.

Moreover, this initial collected data can be scaled with the help of standalone platform independent libraries for Image Augmentation like Augmentor [1] to train the underlying Deep Learning models to be more robust.

The vision is to create a pothole dataset somewhat similar to the Kitti Road dataset [24] which will also be open-sourced and available online for incremental research on it. There should not be any privacy concerns with it, only the approval of the local body would be needed.

- **Labels**:

  In order for the app to train its underlying Deep Learning 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 can be to do Semantic Segmentation [2] wherein one labels the pixels of a pothole in images using a Fully Convolutional Network (FCN). 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 local civic body as well as a road contractor.

- **Technical System**:

  Data Augmentation libraries help scale the few pothole pictures to begin with but will primarily involve Transfer Learning on Image Segmentation models from Autonomous Driving to the pothole scenario. In particular, some related research work (papers and open source implementation) on Deep Learning models for Semantic Image Segmentation that can be applied in this scenario are DeepLab [25, 26], Fully Convolutional Network-Based approaches [27, 28, 29], Weakly Supervised Semantic Segmentation [30, 31] and Region-Based Semantic Segmentation [32, 33].

  From an infrastructure point of view, one of the primary needs to train the Deep Learning based models would be compute i.e. GPU/TPU - wherein one also leverage Deep Learning on cloud for fast training. The system will also require storage for the image-complaints for any future analytics as well as fast inference for real-time complaints on mobile.

  I found a mobile app developer from one of my Machine Learning talks at a Meetup event and he helped me design a hard-coded first prototype that works on 3 image complaints (based on severity of potholes) on my mobile. A pilot study was done on 10 users who gave feedback to publicize the app and include incentives for reporting issues.

[1] lhttps://github.com/mdbloice/Augmentor

[2] https://github.com/udacity/CarND-Semantic-Segmentation/

- **Extensions**: 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 instead of an user's mobile camera. Moreover, with near-about the same technical setup, resolution of other civic issues like garbage disposal, tree fall, mosquito menace, water pipe burst, etc. can then be addressed.

References

[1] "Halifax street - report a pothole or sinkhole." https://www.halifax.ca/transportation/streets-sidewalks/paving-repair/report-a-pothole, 2016.

[2] H. Song, K. Baek, and Y. Byun, "Pothole detection using machine learning," *Advanced Science and Technology*, pp. 151–155, 2018.

[3] S. Agarwal, N. Mittal, and A. Sureka, "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.

[4] "Pothole pal – interactive infrastructure tracking app." http://tiny.cc/qomk5y, 2016.

[5] R. Madli, S. Hebbar, P. Pattar, and V. Golla, "Automatic detection and notification of potholes and humps on roads to aid drivers," *IEEE sensors journal*, vol. 15, no. 8, pp. 4313–4318, 2015.

[6] J. Lin and Y. Liu, "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.

[7] "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.

[8] M. Foth, R. Schroeter, and I. Anastasiu, "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.

[9] A. Mednis, G. Strazdins, R. Zviedris, G. Kanonirs, and L. Selavo, "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.

[10] J. Eriksson, L. Girod, B. Hull, R. Newton, S. Madden, and H. Balakrishnan, "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.

[11] D. B. A. Chacra and J. S. Zelek, "Fully automated road defect detection using street view images," in *2017 14th Conference on Computer and Robot Vision (CRV)*, pp. 353–360, IEEE, 2017.

[12] H. Maeda, Y. Sekimoto, T. Seto, T. Kashiyama, and H. Omata, "Road damage detection using deep neural networks with images captured through a smartphone," *arXiv preprint arXiv:1801.09454*, 2018.

[13] M. J. Dueker and A. M. Fischer, "Fixing swiss potholes: The importance and cyclical nature of improvements," *Economics Letters*, vol. 79, no. 3, pp. 409–415, 2003.

[14] M. Marasteanu, D. Ghosh, M. Turos, M. Hartman, R. Milavitz, J.-L. Le, *et al.*, "Pothole prevention and innovative repair," tech. rep., Minnesota. Dept. of Transportation. Research Services & Library, 2018.

[15] S. Srivastava, A. Sharma, and H. Balot, "Analysis and improvements on current pothole detection techniques," in *2018 International Conference on Smart Computing and Electronic Enterprise (ICSCEE)*, pp. 1–4, IEEE, 2018.

[16] "Waze - connected citizens program." https://www.waze.com/ccp, 2019.

[17] "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.

[18] "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, 2016.

[19] E. M. Thompson, "What makes a city 'smart'?," *International Journal of Architectural Computing*, vol. 14, no. 4, pp. 358–371, 2016.

[20] "Road rehabilitation by mechanized pothole filling." http://donnees.ville.montreal.qc.ca/dataset/refection-de-chaussee-par-remplissage-mecanise-de-nid-de-poule, 2018.

[21] M. Everingham, S. A. Eslami, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes challenge: A retrospective," *International journal of computer vision*, vol. 111, no. 1, pp. 98–136, 2015.

[22] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The cityscapes dataset for semantic urban scene understanding," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3213–3223, 2016.

[23] B. Zhou, H. Zhao, X. Puig, S. Fidler, A. Barriuso, and A. Torralba, "Scene parsing through ade20k dataset," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 633–641, 2017.

[24] J. Fritsch, T. Kuehnl, and A. Geiger, "A new performance measure and evaluation benchmark for road detection algorithms," in *International Conference on Intelligent Transportation Systems (ITSC)*, 2013.

[25] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs," *IEEE transactions on pattern analysis and machine intelligence*, vol. 40, no. 4, pp. 834–848, 2018.

[26] C. Liu, L.-C. Chen, F. Schroff, H. Adam, W. Hua, A. Yuille, and L. Fei-Fei, "Auto-deeplab: Hierarchical neural architecture search for semantic image segmentation," *arXiv preprint arXiv:1901.02985*, 2019.

[27] R. P. Poudel, S. Liwicki, and R. Cipolla, "Fast-scnn: Fast semantic segmentation network," *arXiv preprint arXiv:1902.04502*, 2019.

[28] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3431–3440, 2015.

[29] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "Semantic image segmentation with deep convolutional nets and fully connected crfs," *arXiv preprint arXiv:1412.7062*, 2014.

[30] Q. Li, A. Arnab, and P. H. Torr, "Weakly-and semi-supervised panoptic segmentation," in *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 102–118, 2018.

[31] A. Khoreva, R. Benenson, J. Hosang, M. Hein, and B. Schiele, "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.

[32] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid scene parsing network," in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2881–2890, 2017.

[33] R. Girshick, I. Radosavovic, G. Gkioxari, P. Dollár, and K. He, "Detectron." https://github.com/facebookresearch/detectron, 2018.