Pot-hole Detection and Reporting System Using Deep Learning

Wellington T. Mpofu, Belinda Ndlovu, Sibusisiwe Dube, Joseph Mutengeni
Department of Informatics and Analytics
National University of Science and Technology
Bulawayo, Zimbabwe
wellington.t.mpofu@gmail.com, belinda.ndlovu@nust.ac.zw, sibusisiwe.dube@nust.ac.zw, joseph.mutengeni@nust.ac.zw

Abstract

Pot-holes are a prevalent problem that affects the effective functioning of the road infrastructure. They cause a high number of road accidents, and vehicle damage, resulting in higher maintenance costs. Effective strategies are therefore required to report potholes to road maintenance authorities and thus facilitate constant upkeep and repair. This study focuses on the development of a deep-learning-based pothole detection and reporting system, to assist road maintenance personnel in making informed road repair decisions. A YOLOv5 model is developed and trained using a custom dataset containing pothole images obtained from a survey conducted by the researcher, as well as online sources. The model is integrated into the overall system, to facilitate the detection of potholes from uploaded images. Comprehensive data visualizations are created, with the location of detected potholes added to a map. Pothole maintenance procedures are implemented to ensure that the pothole information in the system is accurate and up to date. Overall, the system provides a clear and concise viewpoint from which road repair decisions can be made and justified.

Keywords
Pothole Detection, Reporting System, Deep Learning, YOLOv5 and Decision Making.

1. Introduction
Good road infrastructure plays a critical role in the transportation system. It enables the efficient movement of people and access to a variety of commercial and social activities (Ng et al. 2019). However, maintenance of this road infrastructure proves challenging, with potholes becoming a major hindrance to the process. Over time, potholes increase in number and size, leading to unpleasant impacts on different road stakeholders. Premanand et al. (2020) state that these bowl-shaped road cavities cause a significant proportion of automobile-related accidents, either directly or indirectly. J et al. (2020) further allude that road maintenance processes are prone to human errors and consume a significant amount of time as they are conducted manually in most places. These factors have resulted in the establishment of road maintenance projects in recent years to assess the performance of road infrastructure, support investment plans, and emphasize upkeep and repair (Rana et al. 2022).

Studies indicate that deep learning-based approaches have been used to detect potholes, with several researchers Srivastava et al. (2018), Rosli et al. (2022), and R et al. (2020) using map services to visually present these potholes to road maintenance authorities. However, identifying potholes and knowing their location does not provide sufficient information that road maintenance authorities need to make informed road repair decisions. Without further exploitation of other pothole information reporting and visualization strategies, challenges are faced in deciding which roads to prioritize fixing first and justifying the reason for that prioritization. This research, therefore, seeks to utilize deep-learning algorithms and data visualization techniques to add more versatility to the pothole reporting process, and thus provide a clearer standpoint from which road repair decisions can be made.

1.1 Objectives
To upload a pothole image and additional pothole information

To detect pothole presence in the uploaded image using deep learning algorithms.

To record pothole location on a map.

To generate visualization reports containing the pothole information.

To email all system users regarding new changes on the map and visualization dashboards.

2. Literature Review

Potholes predominantly form during the rainy season, when water seeps into the underlying soil of the road surface, causing a depression to form on the pavement surface (Premanand et al. 2020). They also form during dry seasons due to heavy traffic and the weight of vehicles using the road infrastructure. Potholes affect daily social and economic activities that depend on road infrastructure. They pose a threat to road users, causing a high number of road accidents. In Zimbabwe, according to police statistics released in 2016, an average of 2000 people die each year because of accidents caused by poor road infrastructure (Zikhali 2017). Potholes lead to higher vehicle maintenance costs, due to damages to wheels, shocks, and axles. Insurance companies suffer huge financial strain due to claims based on pothole damage are frequently made. These factors, therefore, have led to different attempts being undertaken to try and alleviate the problems caused by potholes.

Deep learning techniques have been used by many researchers to detect potholes. These techniques have vastly improved visual object recognition, with convolutional neural networks-based algorithms accurately processing images and videos (LeCun et al. 2015). The availability of high-quality datasets containing annotated training data has also fast-tracked the adoption of deep learning and the creation of high-performing object detection models. Overall, the object detection models are integrated within established pothole detection techniques to facilitate efficient pothole detection. The pothole detection techniques include vision-based, 3D reconstruction-based, and vibration-based techniques (Kim et al. 2022). Vision-based techniques process 2-dimensional (2D) image or video data captured with a digital camera using 2D image or video processing methods (Dhiman and Klette 2020). Several vision-based pothole detection studies have been conducted using cutting-edge deep-learning models, the YOLO family (Asad et al. 2022; Basher 2022; Dharneeshkar et al. 2020; Pereira et al. 2018; Yik et al. 2021). The models are trained using the Darknet and TensorFlow frameworks, to find the model best suited for pothole detection on 2D images. The results from their conducted experiments proved YOLOv5 as the most accurate pothole detection model. 3D reconstruction-based techniques rely on 3D point clouds generated by laser scanners or stereovision algorithms implemented with a pair of video cameras to detect potholes (Rana et al. 2022b). Studies conducted by (Asad et al. 2022) allude that 3D Reconstruction-based techniques measured the shape of potholes most accurately, but however, are less practical for large-scale pothole detection due to the complexity of the hardware required.

Vibration-based techniques capitalize on mobile device technology, with devices increasingly outfitted with sensors such as Gsensors, electronic compasses, gyroscopes, global positioning systems (GPS), microphones, and cameras which facilitate pothole detection (Wang et al. 2015). Pothole reporting and visualization strategies adopted in pothole detection systems play a critical role in avoiding the negative impacts of potholes. Several studies focusing on pothole reporting have been conducted with researchers focusing on alerting road maintenance authorities (Salcedo et al. 2022) to spark a quick response towards pothole fixing and some focusing on supporting both road users and road maintenance authorities (Kuthyar et al. 2021; Prakash Devrukhkar et al. 2022). These studies discuss the development of web and mobile applications integrated with map services. Map services are used as the main pothole data visualization strategy in these systems, whereby road maintenance authorities could view all detected potholes. Prakash Devrukhkar et al. (2022) suggested a community-based pothole data collection process where community members could upload pothole images and create a report using a mobile application.

The researchers presented a feature where all system users were granted the ability to manage the status of their report and make comments where they saw fit. Adding a more dynamic approach to pothole reporting and visualization strategies Rosli et al. (2022) presented PotAlert, a system that offered convenient and simple access to information to road maintenance authorities for pothole management. Road maintenance officials could view data visualizations based on various categories, such as weekly reports and pothole intensity predictions. These visualizations were intended to provide support to the road maintenance authorities in making informed decisions. This research expands on this idea and introduces cloud-based data visualization techniques to give a more transparent picture of the decision-making process. It presents a deep learning-based pothole detection and reporting system that assists experienced road maintenance personnel in conducting the pothole data collection processes. All detected potholes are geotagged, and their location is updated on google maps. Additionally, it introduces a questionnaire-based survey to gather more information about the general state of the road where a pothole would have been located. Visualization reports are then generated based on this information, acting as another variable to consider in the decision-making process. Overall, the system seeks to provide a holistic view of all necessary information for road repair decision-making.

3. Methods

To shape the development of the proposed system, Applied Research was chosen as the research methodology, while Rapid Application Development was chosen as the software development methodology. These two methodologies provided the framework to conduct work in a systematic and iterative manner.

3.1 Data Collection

The researcher identified two categories of data to be used in the research work, that is primary and secondary data. For primary data collection, the researcher travelled to multiple locations in low-density areas which included: Pumula South, Pumula North, Pumula East, Nkulumane 5, Tshabalala, and Lobengula West. With these areas proving to have low traffic during the daytime, it became easy for the researcher to capture pothole images on foot using the Samsung A13 smartphone mounted on a tripod stand from a video-making kit. Figure 1 demonstrates this setup.

Figure 1. On-foot pothole image capturing setup.

In high-traffic areas and roads, the researcher mounted the tripod stand on a Nissan X-Trail moving at low speeds to ensure a safe and efficient way of capturing the pothole images. This setup also enabled the researcher to obtain images from various angles, adding more diversity to the overall pool of images being created. This setup is shown in Figure 2 below.

Figure 2. Pothole image capturing setup mounted on a Nissan X-Trail

The last piece of primary data the researcher attempted to collect was information on the type of questions road maintenance personnel had to answer when conducting their manual road surveys. The researcher approached a road construction company based in Bulawayo named Bitumen World. However, due to a delayed response, the researcher had to find alternative means to acquire the information. Using secondary data collection sources on the internet, the researcher sifted through multiple websites and applications to find the frequently asked questions on road maintenance surveys. A key observation noted in this phase was that the questions were diverse, with some questions having structured and unstructured response types. Furthermore, the researcher also obtained a pothole image dataset from Kaggle to increase the number of pothole images to be used for training the deep learning model.

3.2 Data Analysis and Preparation

The initial step in this phase was visually analyzing the captured pothole images to ensure they did contain potholes. This led to a few images being discarded as they were unclear due to capturing errors. Moving on, the researcher manually annotated the pothole image dataset using the LabelImg Tool. This tool allowed the researcher to draw object bounding boxes on the pothole images using the tool's Python Tkinter User Interface as shown in Figure 3. Upon saving the images with bounding boxes, the tool created a folder with the image annotations converted into the YOLO format. The format is as follows: (class id, x, y, width and height), where the class id defined the object class (potholes), and the rest of the coordinates represented the drawn bounding box. The next step in the dataset preparation was creating a YAML file containing the class number (0), class name (pothole), and a path to the training and validation folders. Finally, the dataset was split into training and testing data.

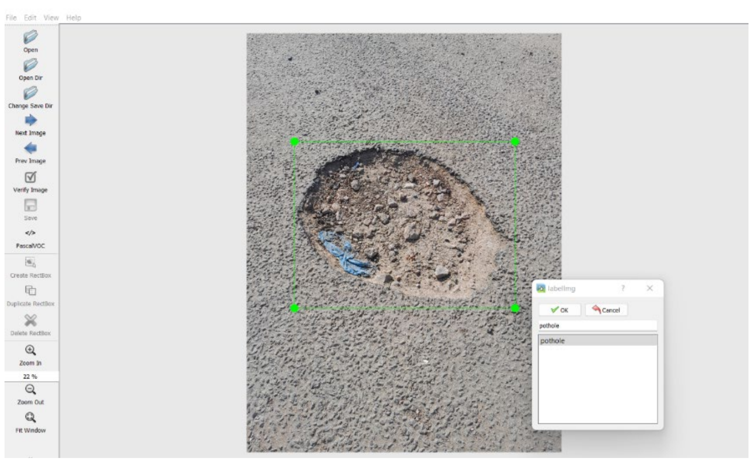

Figure 3. Annotating pothole images using the LabelImg Tool

3.3 System Design

With data analysis and preparation complete, the next step was to design and create the overall system components using the system architecture described in Figure 4.

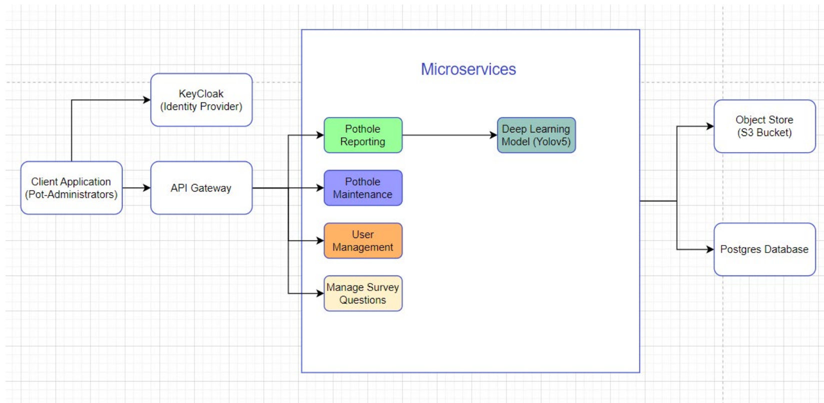

Figure 4. System Architecture

A web-based frontend user interface was created, to act as the main point of user interaction with the system. The user interface allowed the users to perform various tasks shown in Figure 5.

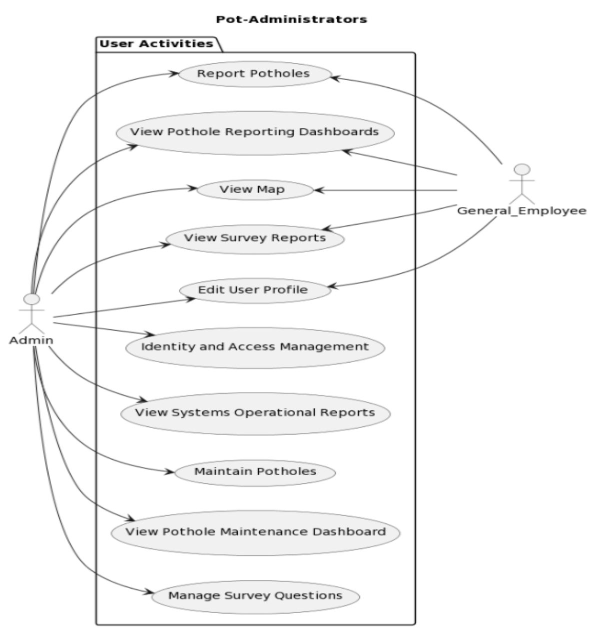

Figure 5. System Use Cases

Various serverless backend microservices were created and integrated to form the pothole reporting process. System users, admin, or general employees initiate the pothole reporting process by selecting an image from their internal storage using the client application. Upon selection, the user clicks the verify image button, which triggers the request to the backend application. The image is uploaded to an S3 bucket, which in turn triggers a lambda function that performs image verification processes. These include checking if the image contains location metadata of where the image was taken (longitude, latitude, and altitude) and checking if an image had previously been verified or not. The execution results are published to an SQS queue, with the client application acting as a consumer of the messages in the queue. Messages are processed and the verification status is returned to the user. If the verification process is successful, the user proceeds to answer the questionnaire-based survey questions for assessing the road condition where the pothole would have been taken.

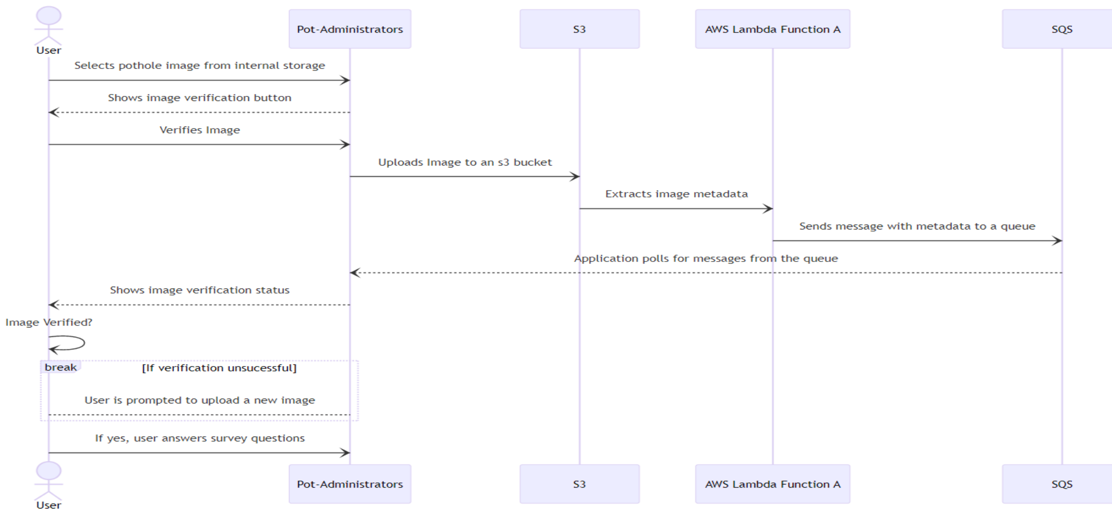

Figure 6. Pothole Reporting Sequence Diagram 1

Finally, the user clicks the submit button to trigger the pothole detection process. The request is published to a second queue, with a lambda function polling for messages from the queue. The lambda function processes the message from the queue and triggers a test instance to the trained deep learning model. The model detects potholes from the uploaded image and returns its execution results to the lambda function. If the results indicate that potholes were detected, the lambda function invokes a process for updating the image's geolocation data on a map. The overall execution results are persisted on the Postgres database deployed on an Amazon RDS instance. Furthermore, the results are published to a third queue, with the client application consuming these messages and displaying them to the user as the pothole detection results.

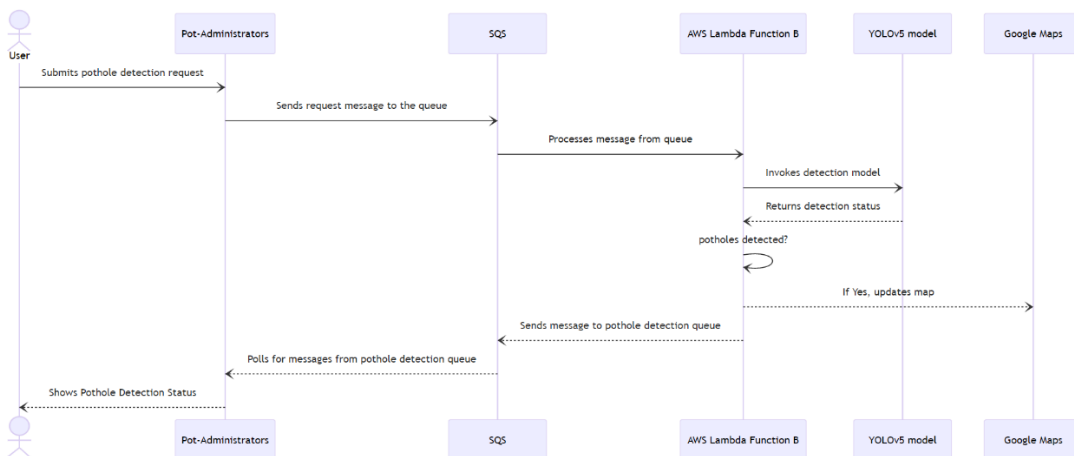

Figure 7. Pothole Reporting Sequence Diagram 2

Finally, a pothole maintenance feature was implemented, with its main goal being to simulate the physical road maintenance process on the ground. This was achieved using Camunda, a business process modelling and orchestration tool. Using Camunda Modeler, a business process model simulating the process of conducting road maintenance work on the field was created and deployed to a Camunda Engine. This process model is shown in Figure 8.

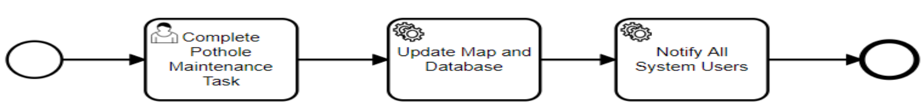

Figure 8. Pothole Maintenance Process Model

When road maintenance work commences on a specific road, admin users initiate a maintenance process instance from the front-end user interface. Upon completion of the on-field work and resolution of the potholes, the admin users complete the representative task of the business process model. A service task for updating the map and the database is triggered, with the resolved potholes being removed from the respective data stores. Lastly, all system users are notified via email of the completion of the maintenance work and updates on the map and the database, which leads to an update on the visualization dashboards.

4. Results and Discussion

The following section discusses the results of the conducted work.

4.1 YOLOv5 Model

The Python programming language was used on the Google Collab platform to train and test the YOLOv5 model. Using the pothole image dataset acquired in prior stages, the researcher trained the model using the YOLOv5s weights provided by the Ultralytics GitHub repository. Tensorboard was then used to visualize the results of the training process. Results are shown in Figure 9.

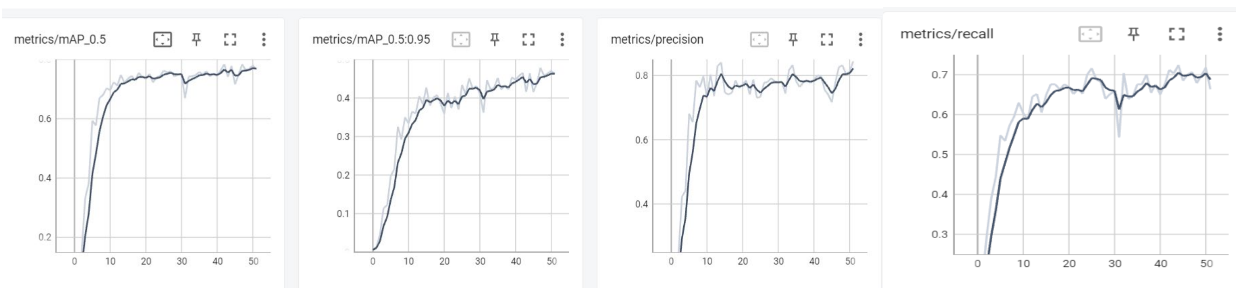

Figure 9. Model Training Results (mAP, mAP_0.5:0.95, Accuracy, and Precision)

Figure 9 indicates the key model evaluation metrics, which are precision, recall and mean average precision. Precision measures the number of correct bounding-box predictions out of all positive predictions, while Recall measures the number of true bounding-box predictions out of all predictions. Mean Average Precision measures the overall accuracy of the YOLOv5 model. Hence the model training process achieved a Precision of 89.4%, 67.5% Recall and 74.8% mAP. These results are further represented in Figure 10.

obj_loss	cls_loss	Instances	Size
0.0265	0	69	640:
Instances	P	R	mAP50
25	0.894	0.675	0.748

Figure 10. Best Training Results

Following the training process, the YOLOv5 model was embedded into a FASTAPI application and deployed to an EC2 instance on AWS. The deployed FASTAPI application acted as the inference point to the model, by exposing two REST API endpoints through which requests could be made to obtain pothole detection results. Figure 11 represents the swagger documentation of the cloud based FASTAPI application from which pothole detection queries can be made.

Figure 11. Swagger Documentation for the FASTAPI application

Making inferences to the model using the Swagger Documentations gives results presented in Figure 12.

Figure 12. Model Testing Results Using the Swagger Documentation

4.2 Pothole Reporting

From a user interface perspective, pothole reporting is done using the screen shown in Figure 13.

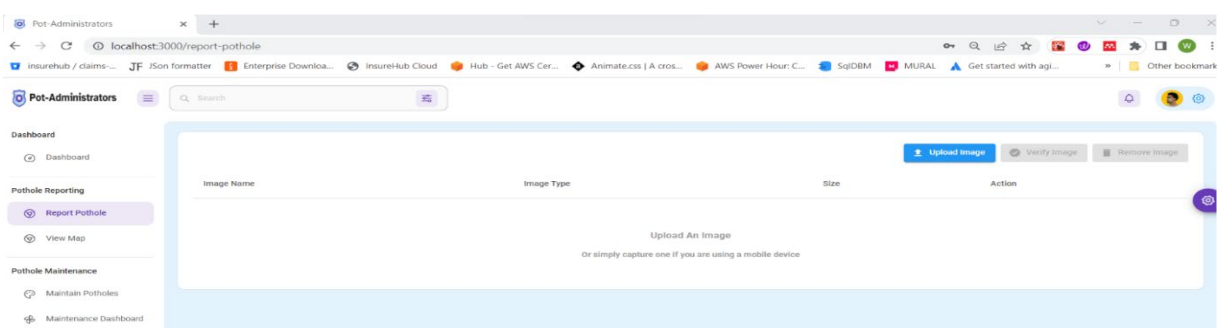

Figure 13. Pothole Reportingf

The screen facilitates users to upload or capture images with potential images. The images are verified to check if they contain geolocation data. If they do, the user can proceed to answer a questionnaire survey or simply make the pothole detection request. If the verification process is successful, the following screen is shown.

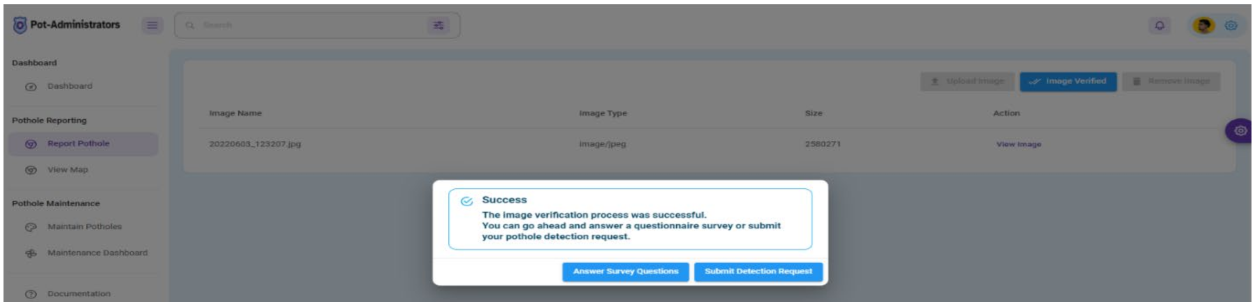

Figure 14. Image Verification Success

Answering the questionnaire survey is not compulsory, hence the users can go ahead and submit the detection request. Figure 15 highlights the output upon successful pothole detection. It indicates the successful detection of potholes in the uploaded image with an accuracy of 93.923%.

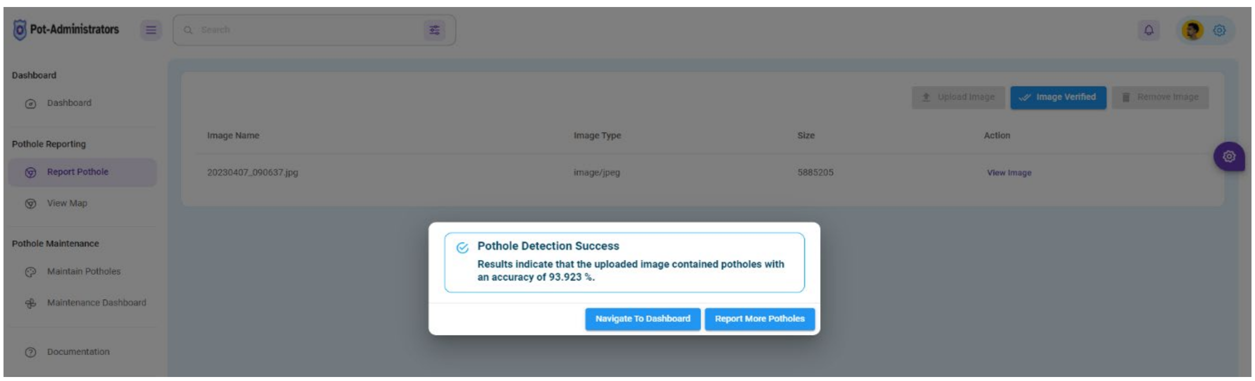

Figure 15. Pothole Detection Success

When potholes are successfully detected, an email notification is sent to all system users. This email notifies the user of the new changes in the visualization dashboards and the map. This email notification is shown below.

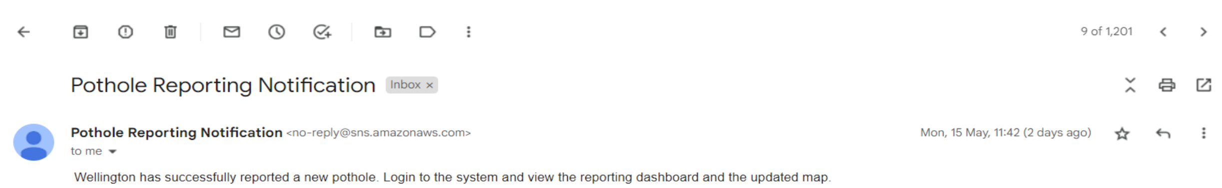

Figure 16. Pothole Reporting Email Notification

Figure 17 indicates the updated visualization dashboard. The dashboard indicates various metrics on reported pothole data. These include the total number of resolved and unresolved potholes annually, as well as the total number of survey responses obtained during the pothole reporting process.

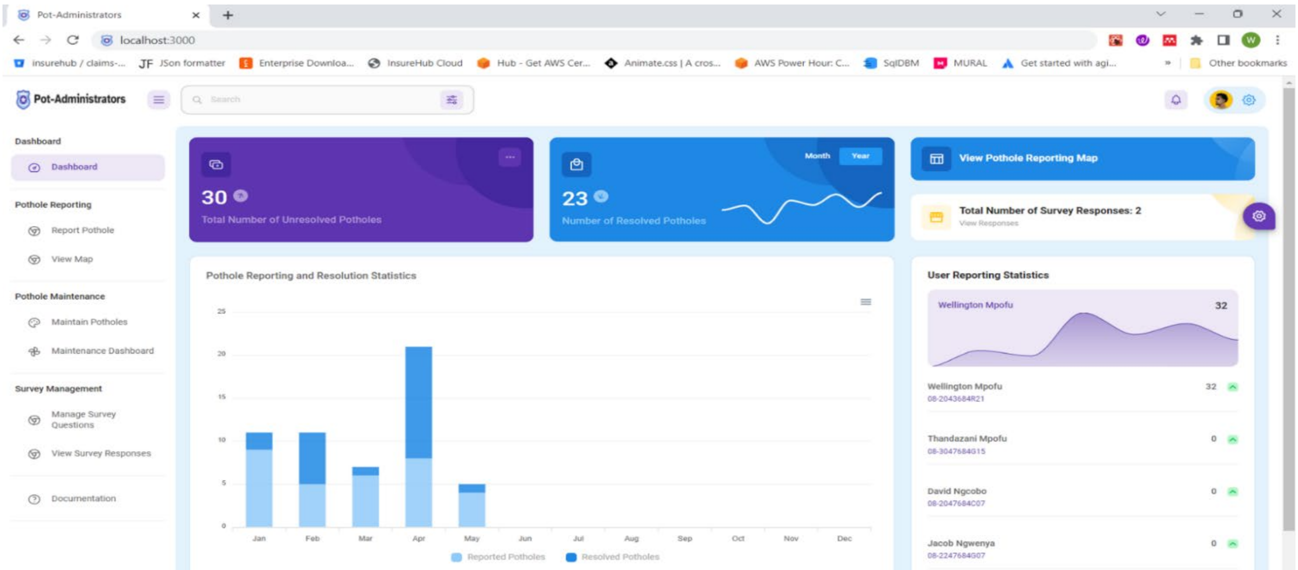

Figure 17. Main Dashboard

A map component shows the location of all detected potholes. The high-level map shows the number of potholes in a clustered manner, for a specific region. When zoomed into the region, the cluster reduces in size until the individual pothole markers are seen. This map is highlighted in Figure 18.

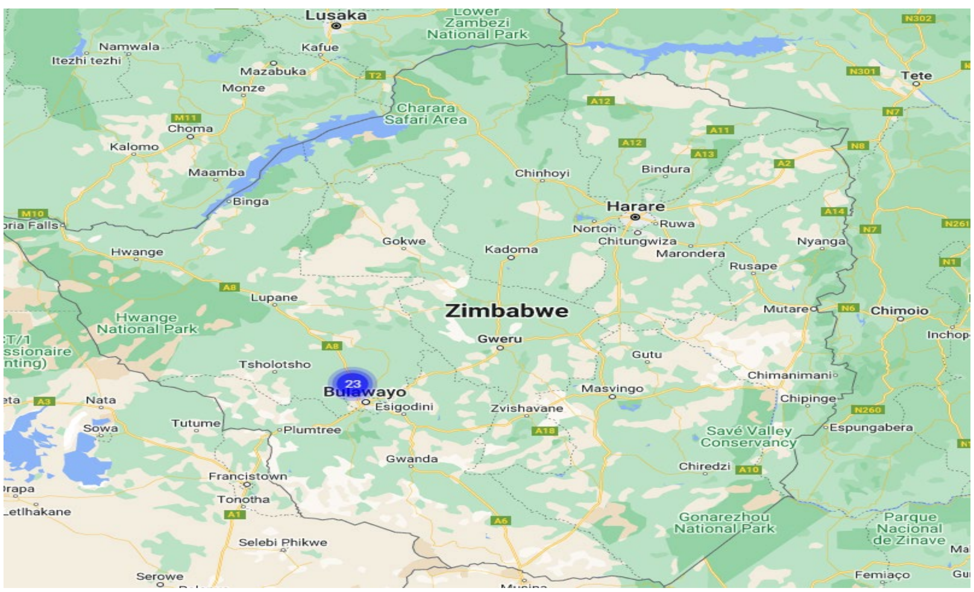

Figure 18. High-level Pothole Map View

4.3 Pothole Maintenance

When the maintenance process is started from the system's user interface, users are redirected to a third-party Camunda platform. This platform allows the users to view all pothole maintenance processes being undertaken. Figure 19 presents the business process model showing 2 running maintenance processes for different road regions.

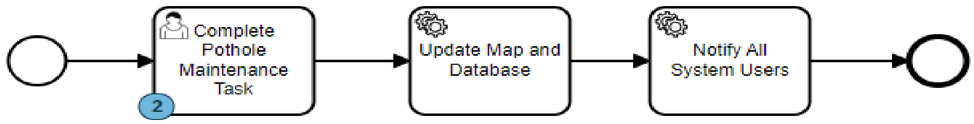

Figure 19. Pothole Maintenance Process

When the actual maintenance work is completed on the ground, the users can then complete the maintenance process. This ensures that all potholes within the maintained region are removed from the system, thus updating the database, dashboards, and map. Figure 19. shows the task completion screen. Finally, an email notification is sent to all users, updating them on the changes on the visualization dashboards and the map.

4.4 Proposed Improvements

The output of the study was a web-based pothole detection and reporting system, with mobile responsiveness at its core, to support smaller mobile devices. However, not all system user interface components were able to achieve the required level of mobile responsiveness required. Hence future developments would be to build a mobile application tailored to perform the same functionalities as the web application. This would improve user interactivity and experience on smaller mobile devices. In future work in the system, batch processing and pothole detection can be implemented. This feature would allow users to upload a folder containing all pothole images collected from a region of interest. These would then be detected in the batch process, with results returned to the user when all processing is completed.

5. Conclusion

This study presented a deep-learning-based pothole detection and reporting system that provides comprehensive data visualizations. The first objective of the study was achieved by creating a user interface feature to upload or capture an image. To capture additional information on potholes in the specified collection region, a questionnaire survey was implemented where system users could provide more context on the state of the road based on their perspective. The second objective was achieved by using a member of the YOLO family of algorithms (YOLOv5) to build a model to detect potholes. To make the trained model available to other system components, it was then embedded in a FASTAPI application and deployed to a cloud environment.

The third and fourth objectives were successfully met. Before pothole detection, an image verification step was implemented, to check if the contained location information of where it was captured. If the information was available, the user would then submit their pothole detection request. Upon successful detection, the location metadata captured was used to add the detected potholes on a map. Visualization reports within the defined dashboards would then be updated accordingly. The fifth objective was also met as all registered system users were able to receive email notifications upon successful pothole detection. A pothole maintenance process was also implemented to ensure that the pothole information in the system portrayed an accurate view of the state of the road infrastructure. All these features led to the successful implementation of a system that supports informed road repair decision-making.

References

Asad, M. H., Khaliq, S., Yousaf, M. H., Ullah, M. O. and Ahmad, A., Pothole Detection Using Deep Learning: A Real-Time and AI-on-the-Edge Perspective, *Advances in Civil Engineering*, 2022.

Basher, R, *Deep Learning-Based Real-Time Pothole Detection for Avoiding Road Accident*, 2022.

Dharneeshkar, J., Dhakshana, S., Aniruthan, S., Karthika, R. and Parameswaran, L., Deep Learning based Detection of potholes in Indian roads using YOLO, *Proceedings of the 5th International Conference on Inventive Computation Technologies,* 2020.

Dhiman, A. and Klette, R., Pothole Detection Using Computer Vision and Learning, *IEEE Transactions on Intelligent Transportation Systems*, 2020.

Kim, Y. M., Kim, Y. G., Son, S. Y., Lim, S. Y., Choi, B. Y. and Choi, D. H., Review of Recent Automated Pothole-Detection Methods, *Applied Sciences Switzerland,* 2022.

Kuthyar, S. R., Roopashree, S., Rasika, V., Manjesh, S., Girimaji, R., Arjun, R. D. and Priyadarshini, S., An Intelligent Pothole Detection and Alerting System using Mobile Sensors and Deep Learning, *Proceedings of the 2021 IEEE 18th India Council International Conference*, 2021.

LeCun, Y., Bengio, Y. and Hinton, G., Deep learning, *Nature*, 2015.

Ng, C. P., Law, T. H., Jakarni, F. M. and Kulanthayan, S., Road infrastructure development and economic growth, *IOP Conference Series: Materials Science and Engineering*, 2019.

Pereira, V., Tamura, S., Hayamizu, S. and Fukai, H., A Deep Learning-Based Approach for Road Pothole Detection in Timor Leste, *2018 IEEE International Conference on Service Operations and Logistics, and Informatics*, 2018.

Prakash Devrukhkar, A., Dethe, A. A., Patel, S., Londhe, S. F. and Sonawane, V., Artificial Intelligence Powered Pothole Detection, Reporting and Management Solution, *International Journal of Innovative Science and Research Technology*, 2022.

Premanand, T., Sridhar, S. and Durbha, A., Pothole Detection System, *International Journal of Engineering and Advanced Technology*, 2020.

R., G. B., C., C., Rao M., S. B., M., S., E., K. and J., S., Deep Learning Based Pothole Detection and Reporting System, *2020 7th International Conference on Smart Structures and Systems*, 2020.

Rana, G. A., Adnan, S. M., Nida, N., Ahmad, W. and Bilal, F., Asphalt Pavement Potholes Localization and Segmentation using Deep RetinaNet and Conditional Random Fields, *International Journal of Innovations in Science and Technology*, 2022.

Rosli, A. S., Roslan, M. and Nordin, A., Design and Development of a Pothole Tracking System (PotAlert) using Web Progressive Application (WPA). In *International Journal on Perceptive and Cognitive Computing*, 2022.

Salcedo, E., Jaber, M. and Requena Carrión, J., A Novel Road Maintenance Prioritisation System Based on Computer Vision and Crowdsourced Reporting, *Journal of Sensor and Actuator Networks*, 2022.

Srivastava, S., Sharma, A. and Balot, H., Analysis and Improvements on Current Pothole Detection Techniques; Analysis and Improvements on Current Pothole Detection Techniques, *International Conference on Smart Computing and Electronic Enterprise*, 2018.

Wang, H. W., Chen, C. H., Cheng, D. Y., Lin, C. H. and Lo, C. C., A Real-Time Pothole Detection Approach for Intelligent Transportation System, *Mathematical Problems in Engineering*, 2015.

Yik, Y. K., Alias, N. E., Yusof, Y. and Isaak, S., A real-time pothole detection based on deep learning approach, *Journal of Physics: Conference Series*, 2021.

Zikhali, W., POTHOLES IN ZIMBABWE: A HINDRANCE TO ECONOMIC DEVELOPMENT. *European Journal of Research in Social Sciences*, 2017.

Biography

Wellington T. Mpofu is a student currently studying Informatics at the National University of Science and Technology. Throughout his studies, Wellington has been actively involved in various projects which have earned him invitations to multiple innovation expos across Zimbabwe. In addition to his academic pursuits, he has gained practical experience through internships and part-time jobs in the tech industry.