

APPLIED RESEARCH

Road to Repair (R2R): An Afrocentric Sensor-Based Solution to Enhanced Road Maintenance

DARRYN ANTON JORDAN¹, STEPHEN PAINE²,
AMIT KUMAR MISHRA², (Senior Member, IEEE),
AND JAN PIDANIC³, (Senior Member, IEEE)

¹Sensorit (Pty) Ltd., Cape Town 8000, South Africa

²Department of Electrical Engineering, University of Cape Town, Cape Town 7700, South Africa

³Department of Electrical Engineering, University of Pardubice, 532 10 Pardubice, Czech Republic

Corresponding author: Jan Pidanic (jan.pidanic@upce.cz)

This work was supported by the INTER-EXCELLENCE (LTAIN19100) Programme of the Ministry of Education, Youth and Sports, Czech Republic, Artificial Intelligence Enabled Smart Contactless Technology Development for Smart Fencing, under Project LTAIN19100.

ABSTRACT Potholes are one of the most important issues in African road-networks. They pose a major threat to mobility and, with time, cause accelerated degradation of the underlying road infrastructure as well as extensive vehicle damage. To address the need for improved infrastructure management, an advanced data gathering solution is required. This paper presents one such solution. The pothole detection, classification and logging (PDCL) system is under active development by Sensorit (Pty) Ltd in collaboration with the University of Cape Town (UCT) Radar Remote Sensing Group (RRSG). This system represents the initial step in Sensorit's modular approach to producing fully autonomous vehicles for African markets. In this paper, an overview of the PDCL system is presented and early results are shown. The efficacy of the system's unique combination of active infrared stereo vision and mmWave frequency-modulated continuous-wave (FMCW) radar sensors is explored. Under various experimental conditions, range-Doppler maps (RDMs) produced by the radar were unable to provide meaningful pothole detections. In contrast, processed depth maps produced by the stereo vision system are demonstrated to successfully detect even shallow potholes.

INDEX TERMS Machine learning, radar, FMCW radar, road maintainance.

I. INTRODUCTION

Collaboration between Sensorit (Pty) Ltd, the University of Cape Town (UCT) Radar Remote Sensing Group (RRSG) and the University of Pardubice (UPA) is supporting development of the pothole detection, classification and logging (PDCL) system. The long term goal of this system is to enable active avoidance of potholes for autonomous vehicles and thus to reduce road accidents. Furthermore, the mapping of individual pothole locations and properties may be used to improve the rate and prioritisation of pothole repair.

The United States (US) department of transportation has defined a pothole as a bowl-shaped hole with a minimum plan

dimension of 150 mm [1]. They assign a pothole severity level based on the maximum depth of the hole from the surface of the road. Classification of potholes for the purpose of repair is, however, far more complicated than simply determining the potholes' dimensions. The specific repair action required for each pothole is dependant on several properties, such as the local road surfacing material, the state of the road base material, and whether or not road surface cracking has caused the pothole [2]. If these and other properties can be measured and mapped, then road maintenance services can manage their resources more effectively and prioritise repairs.

Current pothole detection methods are typically categorised according to their utilisation of either vibration sensing, 2D imaging or 3D reconstruction techniques [3].

The associate editor coordinating the review of this manuscript and approving it for publication was Marco Martalo¹.

Accelerometer-based vibration sensing [4] requires that at least one of a vehicle's wheels pass through or over a pothole for detection to occur. This is clearly undesirable, as tyre damage may result and potholes in the middle of a lane are unlikely to be detected. As such, vibration should ideally only be used to validate detections made using other sensors. Alternatively, traditional computer vision and contemporary machine learning algorithms have been applied to 2D images captured using hyperspectral [5] and thermal cameras [6] in addition to standard red, green, blue (RGB) colour sensors [7]. Finally, 3D reconstruction methods make use of more complex sensors such as lidar [8], structured light cameras [9], time-of-flight cameras, and stereo vision systems.

Absent from the preceding categories is any form of radar-based detection strategy. While uncommon, the ability of ground-penetrating radar (GPR) systems to successfully detect potholes has been demonstrated [10]. Recently, simulations have shown that the radar cross-section (RCS) of a pothole at automotive radar frequencies is only significant if the dielectric properties of the road surface are significantly different than those exposed within the hole [11]. Interestingly, these simulations also reveal that the potholes' wall curvature and overall dimensions only marginally effect its RCS. Given these findings, the initial development prototype of the PDCL system comprises an active infrared stereo vision system and a forward-looking 77 GHz frequency-modulated continuous-wave (FMCW) radar.

The two major elements of novelty presented in this work can thus be summarised as (1) the combination of optical and radar sensors for pothole detection and classification, and (2) the targeted development of such a system for use in developing countries with unique infrastructural challenges and tight budget constraints.

The remainder of the paper is structured as follows. Section II presents the PDCL system and provides an overview of its hardware and software components. The intricacies of capturing optical, radar and system metadata are explored in Section III. Aspects of the experimental configuration used to test the system are presented in Section IV. Section V reveals the results of early system trials and conclusions are drawn in Section VI.

II. SYSTEM OVERVIEW

A. HARDWARE

A block diagram of the PDCL system is presented in Fig. 1. This diagram reveals all of the system's major hardware components and their interconnections.

An Nvidia Jetson Nano developer kit [12] is at the heart of the system. Powered by the vehicle's internal 12 V battery through the use of a DC-DC converter, the Jetson Nano serves as the command centre from which all other components are controlled. The Jetson Nano board also ensures that all sensors interface and communicate with each other in a seamless manner.

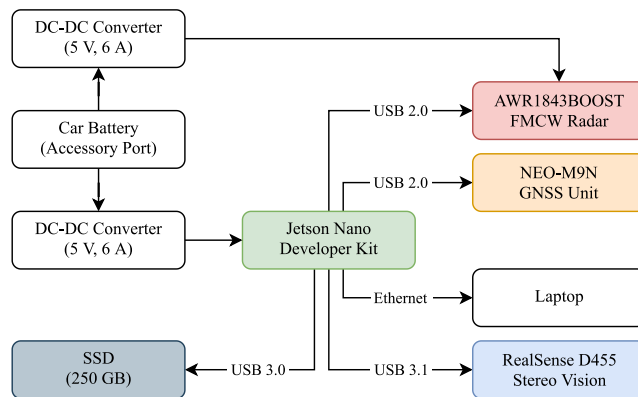


FIGURE 1. Block diagram of the PDCL system, illustrating all major hardware components and their interconnections.

The AWR1843BOOST evaluation module from Texas Instruments (TI) [13] is a complete 77 GHz FMCW radar solution. Its integrated patch antenna array consists of 3 transmit (TX) and 4 receive (RX) elements to enable object detection in three dimensions. It is powered by an independent DC-DC converter and communicates with the Jetson Nano over a universal serial bus (USB) 2.0 connection.

In contrast, the Intel RealSense D455 [14] is an active infrared stereo vision camera which is powered by and communicates over the same USB 3.1 port. It bundles an RGB colour sensor, a pair of infrared imagers, and an infrared projector into a single package. Additionally, the D455 contains a Bosch BMI055 inertial measurement unit (IMU), which is useful for vibration-based validation of pothole detections.

The radar and stereo vision sensors are supported by the NEO-M9N; a quad-band global navigation satellite system (GNSS) unit which boasts a hot start time of 2 s, up to 25 Hz update rate, and horizontal positional accuracy up to 1.5 m.

Finally, the Jetson Nano is connected to an external solid state drive (SSD) for storage of large datasets. Communication with a host laptop is achieved over Ethernet.

B. SOFTWARE

There are three distinct software components that make up the PDCL system: the control server, quick-look viewer and processing chain.

Firstly, the control server runs on the Jetson Nano. In addition to performing several real-time processing steps, the control server presents a web interface that allows the operator to control and monitor all aspects of the system using a laptop and a web browser. Examples of real-time processing steps include the colourisation algorithm which is described later in this paper, and the stacking and compression of sequential range-Doppler maps (RDMs) produced by the AWR1843BOOST. The web interface includes a section for live video feeds, where the D455's colour and depth frames, and the AWR1843BOOST's RDMs can optionally be displayed. Additionally, the interface includes a map section that displays the current location of the PDCL system and

the uncertainty thereof. Numerous system parameters are displayed on the interface to provide a real-time overview of system health. These parameters include system temperatures, memory utilisation, processor load, remaining storage space, sensor configuration details, and ambient weather conditions. Furthermore, buttons are available to control the initialisation and completion of dataset recordings.

As the name suggests, the quick-look viewer is a tool for rapid verification of dataset integrity. During measurement campaigns, each dataset can be immediately reviewed by the system operator following its capture using the quick-look viewer to check that it is readable, complete and usable. This tool is invaluable in the system prototyping and testing phases where datasets may unintentionally become unusable. With rapid integrity verification, experiments and data captures can simply be repeated where errors may have occurred, before leaving the test site.

Finally, the processing chain currently contains all computationally intensive and analysis-orientated processing steps that are presently not performed by the control server on the Jetson Nano. Examples of these processing steps include the *flattening* procedure which is described later in this paper and the time alignment of optical and radar data for analysis. In future iterations of the PDCL system, this processing chain will be tightly integrated into the control server for real-time operation and analysis.

III. DATA ACQUISITION

A. OPTICAL

The Python wrapper for Intel's C++ RealSense software development kit (SDK) was used to interface with, and extract data from the D455 sensor. In addition to control and configuration of the unit, the SDK exposes invaluable functionality such as real-time pixel alignment between colour and depth streams. When these streams are configured for the same resolution and frames per second (FPS), synchronisation is guaranteed between colour and depth frames. This, in addition to the sensor's global shutters, makes the D455 ideally suited for vehicle-borne pothole detection.

During system design it was decided to prioritise FPS over resolution, despite the fact that the D455 produces more accurate depth maps at higher resolutions. This decision was driven by the requirement to capture several images of each pothole at different ranges (and therefore scales) as the vehicle proceeds. In the future, this may enable temporal analysis for improved pothole detection, and supply as much data as possible to any future detection/classification algorithm that requires training data. A resolution of 640 by 480 pixels at a rate of 60 Hz was selected in an attempt to optimise the trade-off between FPS and resolution for a reasonable data rate. These parameters can, however, be adjusted in the control server to enable flexibility during experiment campaigns.

For bonnet-mounted sensing, the D455 was configured for depths between 1 m and 5.5 m, which falls within Intel's recommendation. Sensor accuracy is quoted as $\leq 2\%$ of

the depth [14], which means that its accuracy decreases as a function of depth.

Initially, the synchronised colour and depth frames were stored as matrices in separate datasets of an hierarchical data format (HDF) file. However, this approach proved to be impractical due to the associated high storage rate requirements and large file sizes. As such, it was decided to rather encode and store the sequence of frames as video files. To achieve this, each unsigned 16 bit depth map was transformed into an RGB image in a process known as *colourisation*. Once colourised, standard video encoding could be used to compress both the colour and depth streams, drastically reducing storage rates and file sizes.

The *colorize* method provided in the RealSense SDK was found to significantly stress the Jetson Nano's central processing unit (CPU). To avoid this, the colourisation algorithm published in [15] was implemented in compute unified device architecture (CUDA), to instead run on the Jetson Nano's graphics processing unit (GPU). A *decolourisation* algorithm for depth map recovery was also implemented in CUDA as part of the processing chain. Note that the inequality limits presented in [15] were found to contain errors that prohibit correct mapping to the *hue* colour map.

While this colourisation approach enables a significant data rate reduction, it carries other disadvantages. Firstly, the recommended *hue* colour mapping linearly varies hue while keeping saturation and value constant, which constrains the depth map to a maximum of 1529 unique values. As such, the dynamic range of the depth map is greatly reduced from 16 bit to $\log_2(1529) \approx 10.5$ bit. This quantisation introduces *banding* in the colourised image and prevents true recovery of the input depth map. Secondly, the hardware-accelerated video encoders supported by the Jetson Nano (NVENC) are lossy, which introduces additional error in the recovered depth map. Further issues associated with colourisation include *flying pixels* and *depth inversion*, which are both covered in detail in [15].

To quantify the aforementioned quantisation error, a synthetic depth map was passed through the colourisation and decolourisation pipeline without any form of video encoding. This synthetic depth map was generated by varying the depth value linearly along a single axis between 1 m and 5.5 m. Fig. 2 presents the quantisation error introduced during this process. The error signal observed in this figure is consistent with that of quantisation by truncation, with one exception; the inequality imposed during colourisation and decolourisation cause the sign of the non-zero mean value to alternate at the transition between limits. As such, the measured mean value alternates between plus and minus half of the least significant bit (LSB),

$$\pm \frac{1}{2} \text{LSB} = \pm \frac{1}{2} \frac{(5.5 - 1)}{1529} \approx \pm 0.147 \text{ cm.} \quad (1)$$

The measured root mean square error (RMSE) value in each region also agrees with that expected of quantisation by




FIGURE 2. Measured quantisation error introduced through colourisation, presented as a function of depth. The mean value of the positive and negative regions each agree with that expected of quantisation by truncation and are illustrated by dotted horizontal lines.

truncation,

$$\frac{1}{\sqrt{3}} \text{LSB} = \frac{1}{\sqrt{3}} \frac{(5.5 - 1)}{1529} \approx 0.17 \text{ cm.} \quad (2)$$

These results show that it is desirable to minimise the required range of depth values in order to minimise the effects of quantisation.

Next, the compounded error attributed to quantisation and encoding was assessed by passing the same synthetic depth map through the entire colourisation, encoding and decolourisation pipeline.

The effect of encoding the synthetic colourised depth map using Nvidia's implementation of VP8 is illustrated in Fig. 3, where the RGB channels of the encoded frame are decomposed in 3(a) and the compounded error of quantisation and encoding in the recovered depth map are presented in 3(b).

Significant levels of noise are clearly present in each of the RGB channels of Fig. 3(a), which should ideally be composed of perfectly linear segments. The hue colour mapping is seen to consist of six segments in which a single channel's value is varied linearly. The transition between these segments corresponds to the point at which the polarity of the quantisation noise's mean value changes.

The RMSE of the error signal in Fig. 3(b) was calculated to be 1.68 cm, approximately an order of magnitude greater than that of the quantisation error in Fig. 2. This RMSE is appreciable, since detection of potholes relies on searching for fluctuations in the measured depth map in the order of centimetres. Minimisation of this RMSE is therefore required to avoid false detections.

Table 1 presents the RMSE and the minimum and maximum difference between the recovered and original depth map for all hardware-accelerated video encoders supported by the Jetson Nano.

When applied to the colourised synthetic depth map, the H.264 and H.265 encoders from both vendors are seen to outperform VP8 in RMSE and extreme values. Based on its small extreme values and competitive RMSE, Nvidia's




FIGURE 3. (a) RGB channels of the encoded, colourised depth map, (b) compounded error of quantisation and encoding in the recovered depth map. Encoding was performed using Nvidia's implementation of VP8.

TABLE 1. Error analysis for depth map recovery.

Vendor	Encoder	RMSE (cm)	Min (cm)	Max (cm)
OpenMAX	VP8	1.49	-4.40	4.28
	H.264	0.80	-1.64	2.39
	H.265	0.80	-1.99	1.83
Nvidia	VP8	1.68	-4.45	4.28
	H.264	0.75	-1.82	2.20
	H.265	0.76	-1.70	1.70

implementation of H.265 was selected for use in the PDCL system.

It must be noted that while the synthetic depth map is representative of bonnet-mounted measurements in terms of spatial redundancy, the impact of temporal redundancy was not considered in these tests. As a result, each encoder's inter-frame prediction is not taken into account. An investigation into this is left for future work.

Finally, the importance of colourising disparity values (the reciprocal of depth), rather than raw depth values is illustrated in Fig. 4. The profiles presented in Fig. 4(a) and 4(b) are both one-dimensional vertical slices through colourised depth maps that were captured from the same stationary position, in a similar experimental configuration to that presented in Fig. 5. By colourising disparity, the number of quantisation levels at close depth are increased [15]. As a result, the depth




FIGURE 4. Comparison of RGB channels following colourisation using (a) depth and (b) disparity values. The profiles presented in (a) and (b) are one-dimensional slices through RGB images taken from the same stationary bonnet-mounted perspective.

values are more spatially dispersed in Fig. 4(b) compared to that of 4(a), and the dynamic range is better utilised.

B. RADAR

TI provides its mmWave SDK to enable configuration, control and data extraction from its suite of radars. As part of the SDK, TI has developed demos that enable users to begin working with their sensors rapidly. The *mmWave Demo Visualiser* is provided as a web-based interface that allows users to interact with these demos over a serial connection. This visualiser provided the foundation for interfacing with the AWR1843BOOST using the Jetson Nano.

The baud rate of the serial connection used to interface with the AWR1843BOOST is not high enough to support the transfer of raw analogue-to-digital converter (ADC) samples. As a result, the provided demos only output processed data products, owing to their reduced data rates, such as the magnitude of the RDM. Even so, the FPS of products such as the RDM magnitude is highly constrained [16].

Table 2 presents the parameters used to configure the AWR1843BOOST for data captures. These parameters were carefully tuned to sustain the achieved output of 8 RDM frames per second. At the maximum supported vehicle velocity of 5.5 m/s the system captures 4 RDM frames of each pothole over the 3.2m swath. While a maximum vehicle

TABLE 2. AWR1843BOOST parameters.

Parameter	Value	Unit
Antenna Configuration	4 TX, 1 RX	Elements
Initial Frequency	77	GHz
Sweep Bandwidth	1.54	GHz
Chirp Rate	100	MHz/ μ s
Range Resolution	0.1	m
Maximum Range	5	m
Velocity Resolution	0.1	m/s
Unambiguous Velocity	± 3	m/s

velocity of 5.5 m/s is not acceptable for the final version of the PDCL, it was sufficient for initial system tests.

Under the condition that detections were limited to only potholes that appear directly in front of the bonnet-mounted radar, it was hypothesised that the pothole returns could be separated from those of clutter based on radial velocity.

C. METADATA

In addition to optical and radar data, the PDCL system stores GNSS, weather and system configuration data in an HDF file. Camera parameters include the depth scaling factor, whether or not pixel alignment and disparity are enabled, image dimensions, and camera temperature. GNSS values include latitude, longitude, altitude, speed, and the error associated with each of these. System information includes processor load, remaining disk space, memory usage, and system temperature. Finally, weather data queried by the system includes ambient temperature, pressure, wind, and chance of rain.

This weather data may provide contextual information to a future detection/classification algorithm. For example, the sensitivity of such an algorithm might be dynamically adjusted based on rainy conditions, where potholes might be filled with water and present differently compared to dry conditions.

Furthermore, GNSS information may contribute to a solution that tracks the state of any particular pothole over time, provided each pothole can be uniquely identified on repeat passes.

IV. EXPERIMENTAL CONFIGURATIONS

The bonnet-mounted sensor configuration for the PDCL system is illustrated in Fig. 5. Both the D455 and AWR1843BOOST are seen to be attached to the vehicle using suction-cups and tilted for depression angles of 35° and 15° respectively. At a height of 0.9 m above the road surface, the lower and upper 3 dB points of the radar's antenna intersect the road at a range of 2.4 m and 6.4 m respectively. This ensured overlap with the depth camera's 58° vertical field of view (FOV). A pothole of interest is present near the rightmost edge of Fig. 5, but is obscured by the shade of overhead trees. This pothole was used for both stationary and moving tests. For clarity, the location of the pothole has been enclosed by a dashed circle.



FIGURE 5. Bonnet-mounted experimental configuration of the D455 and AWR1843BOOST sensors. The remainder of the PDCL system was positioned inside the vehicle. The location of a pothole of interest has been highlighted in front of the vehicle.

V. RESULTS

A. OPTICAL

Initial stationary measurements saw the vehicle parked 1.8 m away from the pothole, as seen in Fig. 5. Visualisation of the depth map proved challenging in this configuration due to the pothole's small (centimetre-level) deviation in the road surface within the depth map's dynamic range of several metres. To address this challenge the depth map was *flattened* by modelling the road's curvature in the width and height dimensions and then subtracting this model from the measured depth data. This was achieved by first computing the reciprocal of the depth map to form a disparity map, then fitting polynomials in the width and height dimensions through the centre of the map. The intended effect was to transform the perspective of the depth camera to that of top-down, such that the extent of values was reduced and any deviations would be easier to detect. The result of this procedure is presented in Fig. 6(b).

The RGB image and disparity map of Fig. 6(a) and 6(b) respectively are pixel-aligned to enable direct comparison. While the pothole is relatively difficult to identify in the RGB image, its presence is clear in the disparity map. Dashed circles have been used to highlight the location of the pothole in the RGB image and disparity map. Note that the dimensions of the disparity map have been reduced to aid in this visualisation. In addition to the pothole, the gutter and sidewalk on either side of the road present a significant deviation in the otherwise mostly flat road surface. This explains their prominence in the disparity map. The vertical dotted line at a width value of 320 illustrates the location at which a one-dimensional slice was taken through the disparity map.

Fig. 7 is the one-dimensional profile that corresponds to the vertical dashed line in Fig. 6(b). A dashed ellipse has been superimposed over the region of disparity values that corresponds to the pothole.

Even in stationary tests, both the measured disparity map and profile show significant variation between frames due to the previously analysed compound quantisation and encoding error. This compound error is visible in Fig. 7, where it appears as additive noise. It is vital to note that the OpenMAX VP8 encoder was used for all measurements in this

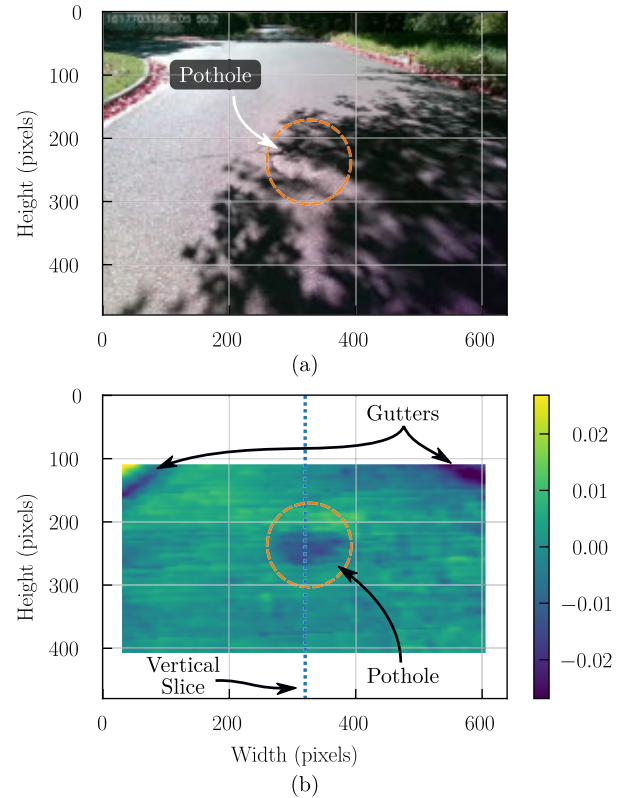


FIGURE 6. Pixel-aligned (a) RGB image and (b) relative disparity map of the same pothole seen in Fig. 5, both captured using the D455. Note that the disparity map in (b) presents relative depth values over the range of 20 cm.

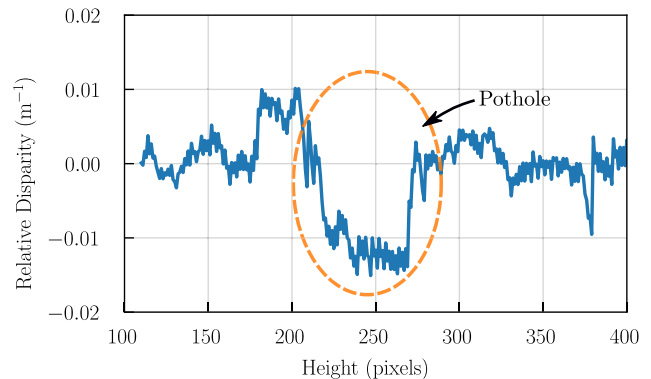


FIGURE 7. One-dimensional relative disparity profile which corresponds to the vertical dashed line in Fig. 6(b).

campaign, which took place before the detrimental impact of this encoder was known. Future work therefore includes more measurement campaigns using the superior Nvidia H.265 encoder.

Based on the presented results, it is hypothesised that a **convolutional neural network (CNN)** trained on flattened disparity maps should provide a suitable means of pothole detection and classification. This is reasonable, given the perceptible presence of the pothole in the flattened disparity map of Fig. 6(b) and the fact that CNNs are inspired by

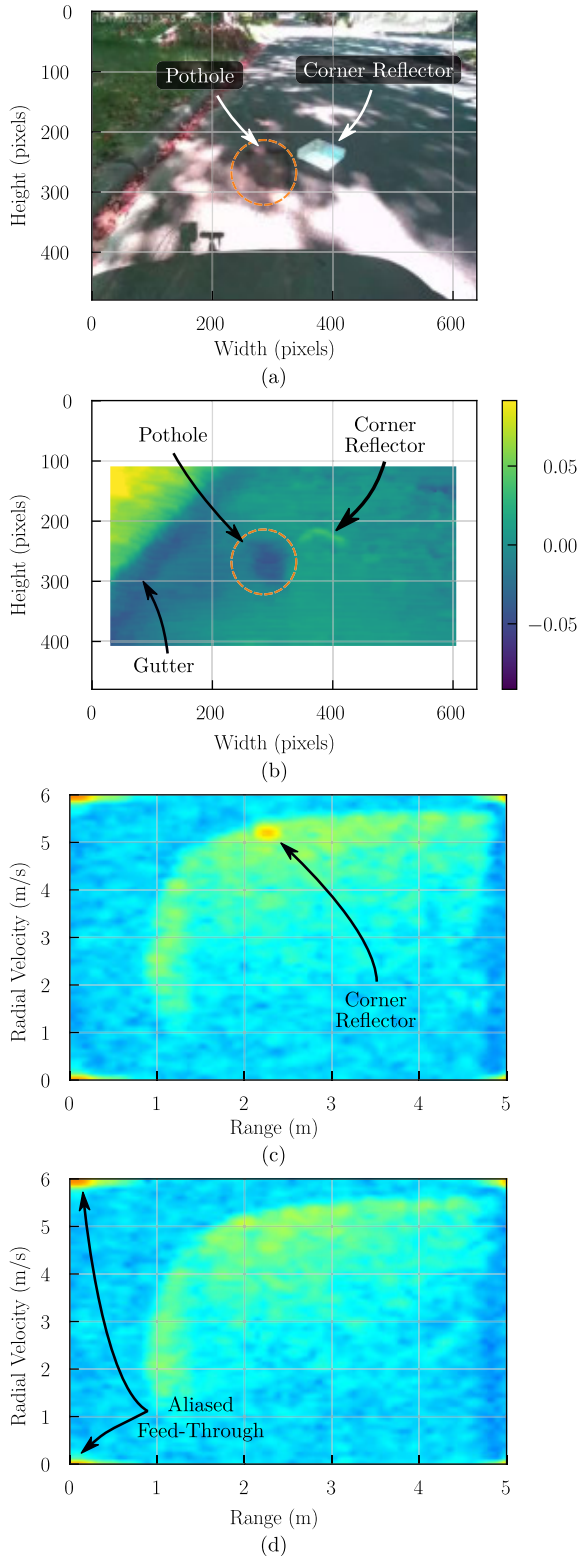


FIGURE 8. (a) RGB image, (b) disparity map, and (c) RDM captured when a corner reflector was placed next to the pothole of interest. For comparison, the corner reflector has been removed to produce the RDM in (d).

the way that the human visual cortex processes information. Furthermore, the current generation of computational resources, such as those available on the Jetson Nano, enable

deployment of deep layer CNNs at the edge. Lastly, for classification tasks where objects are perceivable by humans, collection of massive amounts of tagged training datasets has become relatively easy, which has fuelled the training of deep CNNs. These three factors have made CNNs the best pattern classification machines for image-based object classification where the images are in the bandwidth that is human perceivable [17].

B. RADAR

As discussed previously, it was decided that the RDM was the most appropriate option for pothole detection given the platform limitations. This decision was based on the ability of the RDM to separate the response of targets from clutter based on radial velocity. The radial velocity associated with stationary targets with respect to the moving radar is a function of their angle in both elevation and azimuth, measured relative to the antenna boresight. Having orientated the radar in the vehicles’ forward looking direction, objects located in the centre of the road are perceived as having a greater radial velocity than those located at the edges. Under the condition that targets of interest are located within the bounds of the road, the RDM can therefore be used to isolate targets based on velocity and improve signal-to-clutter ratio (SCR).

Unfortunately, but not surprisingly, the results obtained with the forward-looking radar are in line with the simulations obtained in [11], and the RCS of the pothole proved too small to be reliably detected in the RDM. This was determined by comparing the RDM produced for a scenario in which a small corner reflector was positioned alongside the pothole, to the RDM produced for a standalone, unmodified pothole. Fig. 8 presents this comparison of RDMs along with the associated RGB image and disparity map which were captured by the D455 simultaneously.

The location of the pothole and corner reflector have been marked in the RGB image and disparity map of Fig. 8(a) and 8(b) respectively. Note that only the vertical sides of the corner reflector are visible in the disparity map. At the vehicle’s velocity of 5.5 m/s the Doppler spectrum was aliased around the velocity axis. As such, the velocity axes of the RDMs in Fig. 8(c) and 8(d) have been unwrapped. This is evidenced by the aliased feed-through component. Comparison between the RDMs reveal that the RCS of the pothole is not significant enough to be successfully detected by the radar.

VI. CONCLUSION

This paper has provided an overview of the PDCL system and presented results from initial measurements. A detailed comparison of the error in depth map recovery for several video encoders revealed that Nvidia’s H.265 encoder should be used for all future measurement campaigns. Furthermore, the disparity maps produced through depth map flattening show great promise for CNN-based detection. Unfortunately, however, the RDMs produced by the radar proved ineffective in the detection of potholes.

Future work includes the implementation of plane detection methods for depth maps, such as **random sample consensus (RANSAC)**. Alternative mappings for depth map colourisation should be investigated to potentially improve dynamic range. Finally, the temporal performance of the videos encoders should be analysed for a more comprehensive comparison.

ACKNOWLEDGMENT

The authors would like to thank Sensorit.io and the **Technology Innovation Agency (TIA)** of South Africa for the funding received to facilitate this research.

REFERENCES

- [1] J. S. Miller and W. Y. Bellinger, "Distress identification manual for the long-term pavement performance program," Federal Highway Admin., Richmond, VA, USA, Tech. Rep. FHWA-HRT-13-092, 2014.
- [2] P. Paige-Green, A. Maharaj, and J. Komba, "Potholes: A technical guide to their causes, identification and repair," Council Sci. Ind. Res., Pretoria, South Africa, Tech. Rep., 2010. [Online]. Available: <https://www.yumpu.com/en/document/view/34833737/potholes-technical-guide-to-their-causes-identification-csir>
- [3] T. Kim and S. K. Ry, "Review and analysis of pothole detection methods," *J. Emerg. Trends Comput. Inf. Sci.*, vol. 5, no. 8, pp. 603–608, Aug. 2014.
- [4] R. S. Rodrigues, M. Pasin, A. Kozakevicius, and V. Monego, "Pothole detection in asphalt: An automated approach to threshold computation based on the Haar wavelet transform," in *Proc. IEEE 43rd Annu. Comput. Softw. Appl. Conf. (COMPSAC)*, Jul. 2019, pp. 306–315.
- [5] M. Abdellatif, H. Peel, A. G. Cohn, and R. Fuentes, "Pavement crack detection from hyperspectral images using a novel asphalt crack index," *Remote Sens.*, vol. 12, no. 18, p. 3084, Sep. 2020.
- [6] A. Akula, Y. Bhatia, R. Rai, V. Gupta, N. Aggarwal, and A. Akula, "Convolutional neural networks based potholes detection using thermal imaging," *J. King Saud Univ., Comput. Inf. Sci.*, vol. 34, no. 3, pp. 578–588, Mar. 2022.
- [7] A. Akagic, E. Buza, and S. Omanovic, "Pothole detection: An efficient vision based method using RGB color space image segmentation," in *Proc. 40th Int. Conv. Inf. Commun. Technol., Electron. Microelectron. (MIPRO)*, May 2017, pp. 1104–1109.
- [8] B.-H. Kang and S.-I. Choi, "Pothole detection system using 2D LiDAR and camera," in *Proc. 9th Int. Conf. Ubiquitous Future Netw. (ICUFN)*, Jul. 2017, pp. 744–746.
- [9] D. Joubert, A. Tyatyantsi, J. Mphahlehle, and V. Manchidi, "Pothole tagging system," in *4th Robot. Mechatronics Conf. South Afr. (RobMech)*, Nov. 2011. [Online]. Available: https://researchspace.csir.co.za/dspace/bitstream/handle/10204/5384/Joubert_2011.pdf?sequence=1&isAllowed=y
- [10] S. Li, C. Yuan, D. Liu, and H. Cai, "Integrated processing of image and GPR data for automated pothole detection," *J. Comput. Civil Eng.*, vol. 30, no. 6, Nov. 2016, Art. no. 04016015.
- [11] A. Srivastava, A. Goyal, and S. S. Ram, "Radar cross-section of potholes at automotive radar frequencies," in *Proc. IEEE Int. Radar Conf. (RADAR)*, Apr. 2020, pp. 483–488.
- [12] NVIDIA. (2021). *NVIDIA Jetson Nano Developer Kit*. [Online]. Available: <https://developer.nvidia.com/embedded/jetson-nano-developer-kit>
- [13] Texas Instruments. (2021). *AWR1843BOOST FMCW Radar Evaluation Board*. [Online]. Available: <https://www.ti.com/tool/AWR1843BOOST>
- [14] Intel. (2021). *Intel RealSense D455 Camera Module*. [Online]. Available: <https://www.intelrealsense.com/depth-camera-d455>
- [15] T. Sonoda and A. Grunnet-Jepsen. (2020). Depth image compression by colorization for Intel RealSense depth cameras. Intel. [Online]. Available: <https://dev.intelrealsense.com/docs/depth-image-compression-by-colorization-for-intel-realsense-depth-cameras>

[16] Texas Instruments. (2020). *mmWave Demo Visualizer*. [Online]. Available: <https://www.ti.com/lit/ug/swru529c/swru529c.pdf>

[17] P. Wang, E. Fan, and P. Wang, "Comparative analysis of image classification algorithms based on traditional machine learning and deep learning," *Pattern Recognit. Lett.*, vol. 141, pp. 61–67, Jan. 2021.



DARRYN ANTON JORDAN received the B.Sc. and Ph.D. degrees from the University of Cape Town, in 2015 and 2021, respectively. His research interests include frequency-modulated continuous-wave synthetic aperture radar systems, radar signal processing, and high-performance computing.



STEPHEN PAINE received the M.Sc. degree in radar and electronic defense specializing in RF and antenna design, in 2016, and the Ph.D. degree in passive radar jamming techniques from the Radar and Remote Sensing Group (RRSG), UCT. He is currently a Lecturer with the Department of Electrical Engineering, University of Cape Town, where he is also the Head of the Radar and Remote Sensing Group (RRSG). Previously, he was worked in industry both locally and internationally with research interests include the radar, passive radar, EW, signal processing, and automotive systems. He is a Junior Member of the IEEE.



AMIT KUMAR MISHRA (Senior Member, IEEE) received the Ph.D. degree from the University of Edinburgh. He is currently a Professor with the Department of Electrical Engineering, University of Cape Town (ranked at 157 by THE, in 2021). Before Cape Town, he worked at Australia and India. His research interests include radar system design and applied machine learning. He has successfully supervised nine Ph.D. students so far and holds five patents. He is currently the Head of the Centre for 5G in SDG which works on projects related to the application of 5G in challenges related to sustainable development goals.



JAN PIDANIC (Senior Member, IEEE) received the M.Sc. and Ph.D. degrees from the University of Pardubice, in 2005 and 2012, respectively. His research interests include signal processing in passive radar systems, bistatic radars, clutter modeling, and optimization of signal processing algorithms with parallel processing techniques.