Making Analog Water Meter Smart using ML and IoT-based Low-Cost Retrofitting

by

A. Kumar Lall, Ansh Khandelwal, Rhishi bose, Nilesh Bawankar, Nitin nilesh, ayush.dwivedi, Sachin Chaudhari

Report No: IIIT/TR/2021/-1



Centre for Communications International Institute of Information Technology Hyderabad - 500 032, INDIA August 2021

Making Analog Water Meter Smart using ML and IoT-based Low-Cost Retrofitting

A. Kumar Lall, A. Khandelwal, R. Bose, N. Bawankar, N. Nilesh, A. Dwivedi, S. Chaudhari

International Institute of Information Technology Hyderabad (IIIT-H), India

Email: {ayush.lall, rhishi.bose, nilesh.bawankar, ayush.dwivedi, nitin.nilesh}@research.iiit.ac.in,

ansh.khandelwal@students.iiit.ac.in, sachin.chaudhari@iiit.ac.in

Abstract—This paper introduces an internet-of-things (IoT) based economic retrofitting setup for digitising the analog water meters to make them smart. The setup contains a Raspberry-Pi microcontroller and a Pi-camera mounted on top of the analog water meter to take its images. The captured images are then preprocessed to estimate readings using a machine learning (ML) model. The employed ML algorithm is trained on a rich dataset that includes digits from the images of water meters captured by the hardware setup for ten days. The readings are posted on a cloud server in real-time using Raspberry-Pi. High temporal resolution plots of flow rate and volume are generated to derive inferences. The collected data can be used for deriving water consumption patterns and fault detection for efficient water management.

Keywords- IoT, ML on the node, Retrofitting, Smart water meter

I. INTRODUCTION

The conventional method of manually reading analog meters to calculate a consumption trend is cumbersome and expensive. This approach is also incapable of effectively managing sustainable water supplies, as it needs accurate monitoring techniques that enable the consumers to know the level of water usage in real-time. The traditional analog water meters have a long life, and removing them for digitization is a waste of resources. Although digital water meters are introduced in recent times and are used in workplaces like government institutions, hospitals, they are very expensive. Also, the digital water meter themselves do not give any inference or do not do any analytical analysis on the water consumption patterns. A smart device for water monitoring can make users reduce their use of water to conserve it.

There have been few works in the literature for retrofitting the existing analog utility meters for making them smart [1], [2]. In [1], a wireless sensor network-based water management system is proposed, where a analog water meter with Reed switch is used, which gives pulse output. The pulse output is processed and transmitted to the server using IEEE 802.15.4 protocol. The data is visualized and monitored using a webbased system developed for this purpose. In [2], authors have proposed a system to convert the image of utility meter into numeric data using the convolutional neural network (CNN), which is a deep learning based method. The paper compared the algorithms based on two architectures: You Only Look Once (YOLO) [3] architecture and LeNet [4] architecture. The comparison was carried out on images of gas and water meters. The work in [5] takes the reading problem from gas meter. Each digit is recognized using a Support Vector Machine due to which complexity of the calculation is high. In [6], a meter image capturing system (MICAPS) is proposed to read data from meter images. The authors used a K-Nearest Neighbor (KNN) based ML algorithm to predict digits from the meter.

In this paper, an ML and IoT-based low-cost retrofitting of existing analog water meters is proposed for making them smart. Specific contribution of this paper are

- An IoT-based sensor node is designed and deployed in field to take water meter images, convert them into digits and send it to cloud.
- A retrofitting approach is used which does not temper the existing meter in any physical form. It is assumed that the meter does not have any pulse output and relies completely on the conversion of images to digits.
- The designed node can provide real-time data with high temporal resolution (readings every few secs) and is equipped with lighting to enable readings even in nights. This way the meter readings can also give accurate estimates of derived parameters such as flow rate.
- Simple ML algorithm is used to recognize the digits from the meter image, which requires low processing and is implemented at the node itself.
- The performance of the ML algorithm is further improved by using specific constraints related to the water meters.
- The proposed approach is evaluated based on the data of over 10,000 images collected from the field deployment of 10 days. The collected data will be made public for further research in future.

Note the novelty of the work done in this paper as compared to the existing literature [1], [2]. The work in [1] assumes pulse output and does not rely on the images of the meter as we do. The work in [2] uses computationally heavy CNN models while we use computationally light ML algorithms. Thus, ML algorithms are energy efficient to be deployed on the edge or even node itself. The comparison in [2] is done on 100 images, while we have evaluated over a large dataset. The work in [2] and [6] focus only on the conversion of images to digits and do not take into account the use-case specific constraints. Also, in both of them, there is no clarity on the sensing interval or facility to provide light in low light times.

The remainder of this paper is organized as follows: details

of the hardware implementation are provided in Section II. Section III discusses dataset preparation while Section IV details the methodology of the proposed algorithm. Experimental results and evaluations are reported in Section V. Section VI concludes the paper.

II. HARDWARE DESCRIPTION

A. Circuit description

Figs. 1(a) and 1(b) shows the block architecture and circuit design of the proposed model. The model comprises a Raspberry-Pi 3B+ microcontroller [7], a Raspberry-Pi V2.1 camera module, an LED ring for illumination and an activehigh relay module to control the LED. The hardware is powered by a Li-ion power bank which is connected to AC mains for charging. The power bank enables the device to function without interruption, even if the primary AC power supply is unavailable temporarily. The model is also equipped with a lighting feature, such that the readings can be obtained even at night when there is no ambient light available. The lights are controlled using an active high relay module to operate only when the camera captures an image. This feature helps in extending the battery life of the model. The numeric values are extracted from the images of the water meter dial captured by the camera. The micro-controller executes the ML-based image processing algorithm to detect the reading on the meter. This reading is transmitted in real-time to ThingSpeak [8] using an LTE based portable WiFi hotspot [9]. ThingSpeak is a cloud-based IoT platform for aggregating and processing data. The POST method of the HTTP protocol is used to write data on the ThingSpeak server. This setup can sense the reading at a high frequency (as quick as every 5 seconds), even when the flow is at its peak. Such high-frequency information can be used to derive valuable insights about the community's consumption patterns and timely detection of leaks or faults.

B. Retrofit structural view

Fig. 2 shows the structure of the developed enclosure. It is a 3D-printed multi-layer structure made up of Poly-lactic Acid (PLA) material which offers protection against various weather conditions in outdoor deployment. It consists of a four-layered stack which separates the various hardware component for comfortable placement. The first layer (bottom most) consists of an LED ring for providing adequate illumination to capture good quality images. In the second layer, the camera module is placed facing the dial of the meter. The camera is configured at a focus of 4 cm with the maximum resolution of 3280×2464 and pixel size of $1.12 \times 1.12 \ \mu$ m. The Raspberry-Pi microcontroller is placed at the third layer. It is interfaced with an active high relay module using the Raspberry-Pi GPIO pins to control LED switching. At the fourth layer (topmost), the power bank is placed. The whole setup is mounted on top of the analog water meter without altering or tempering the analog meter in any sense. It is important to note that the discussed water meters are ISI marked [10], a standardscompliance mark for India's industrial products provided by the Bureau of Indian Standards (BIS). Hence any technological



(a) Block architecture of the model







Fig. 2. 3D structure and deployed model

intervention made by physically altering the meters would lead to loss of standardization.

III. DATASET

With the help of the above-described hardware set up on the water meter, two separate datasets were created—one for training the model and the other for analyzing the volume flow and its flow rate. As explained in the hardware section, the camera position is fixed, and the orientation of the images remains the same for all captured images. We specified the



Fig. 3. Algorithmic pipeline for the proposed method. Notice that last digit of the meter is discarded due to digit image ambiguity. Both the training and detection phases has shown. In training phase, we train the ML model on digit features and their corresponding labels. In the detection phase, each segmented digit image feature is predicted using the trained ML model. (*Best viewed on screen*)

coordinates where the meter reading was present and extracted the region of interest (RoI). The resultant RoI image is shown in Fig. 3(b). As shown in Fig. 3(a), the meter reading is tilted towards the right side. In order to straighten the RoI with respect to specified width and length, perspective transformation was used. The transformation matrix is calculated from the manually selected source and destination coordinates. Using these transform coefficients, warp perspective [11] was applied to get the transformed image. This step helps us to form nice rectangle-shaped contours.

As the next step, all the digits in the complete reading are segmented into individual digit images and are stored separately as per their numerical value. This makes the dataset suitable for supervised learning. These digit images were further used to train the ML model. The dataset consisted of digits from 0 to 9. These digits were collected from the meter images captured using the Pi camera for 10 days, out of which 7 days data was used to analyze the water flow and the rest for training.

One of the significant advantages of the dataset is that the digits on the water meter have the same font style. Hence there are significantly fewer variations in the orientation while capturing the digit images. As shown in Table I, we need less numbers of individual digit images (ranging from 0 to 9) to cover all possible digit variations. So training the ML model once on this dataset could be used for other meters also that were present on the campus. Below is the estimate of the number of digit images extracted for the ML model training.

 TABLE I

 FREQUENCY OF THE INDIVIDUAL DIGIT IN THE TRAINING DATASET

Digit	0	1	2	3	4	5	6	7	8	9
Frequency	148	110	77	97	45	98	86	85	63	50

We tried to cover all possible variants of the digit images that occurred inside the water meter (RoI), i.e., since it is an analog meter, the digits are present on rolling wheels. As the water is flowing, these wheels will rotate for the next digit to come. It can be observed from Table I that the distribution of the frequency of digits is not uniform. The primary reason for that is we made our dataset with the help of a real-time water meter, and thus the rate at which the digits of the meter changed was also not uniform. We also observed ambiguity in digit images (half of the previous digit and half of the incoming digit) cause of rotating-disc type meters. Simultaneously, the images were being captured, and such ambiguous digits were discarded for the dataset collection part.

After creating the digit dataset, water meter images were collected to study the volume flow and its corresponding rate with which the water is flowing. For this purpose, Raspberry Pi captured meter images for around two days at every minute, which resulted in a total of 10,508 images. These images were then analyzed with the algorithm's help to get the volume flow and its corresponding flow rate.

IV. METHODOLOGY

A. Training

The digit dataset, explained in the dataset section, was used to train a ML model to classify the given digit image from 0 to 9. As this is a supervised learning problem in ML, a classification-based method was used to train the model. There are several classification algorithms in this particular area. However, we chose tree-based method to solve the problem at hand.

In this paper, Random Forest (RF) [12] classifier was used to train on the digit dataset. RF is a supervised learning algorithm that combines multiple decision trees [13] and is trained together with bagging [13]. The bagging method uses the idea that the combination of several weak learning models increases the overall result. To explain the RF model more straightforwardly, it can be said that it combines multiple decision trees and merges the results obtained by them to get a more accurate and stable prediction. The idea behind choosing this particular classifier is to reduce the number of hyperparameters while training the model. It is also one of the most used classifiers for these problems because of its simplicity and diversity.

The training procedure follows these steps:

- 1) For each digit in the dataset, the digit's Histogram of Oriented Gradients (HOG)-based features [14] were computed, which is a two-dimensional matrix as shown in Fig. 3(f). This matrix was flattened and converted to a one-dimensional feature vector with size $n \times 1$. As the image height and width remains the same for all the digit images, we always get a one-dimensional $n \times 1$ feature vector for each digit image. At the end of this step, we have $m \times n$ sized data matrix M where m is the number of samples in the digit dataset.
- 2) In the second step, the corresponding labels were collected for each digit image in the dataset. The label ranges from 0 to 9. After this step, we have $m \times 1$ label vector.
- 3) The data matrix M was split into two parts: training and validation with the ratio 80:20, respectively. This is a standard paradigm used in ML methods to test the generalization of the trained model. The RF model is trained on training data only and tested on validation data to find the error rate. The training data was used to train the RF model as a classification supervised learning problem. The RF model's input is the training data matrix M' (training part of M) and the corresponding label vector. The training phase is shown in Fig. 3(g). Finally, this trained model was saved for further detection tasks.

B. Detection

To detect the digits of the water meter we present a 4-step structure: i) RoI Extraction, ii) Image Preprocessing, iii) Digit Image Segmentation, iv) Digit Recognition and Correction.

1) Region of Interest (RoI) Extraction: The specific region in the image, which consists of the digits, i.e., the object area, is manually extracted by inputting the RoI's coordinates in the algorithm. We have leveraged the fact that the camera is permanently fixed in one position, and hence setting the coordinates once is sufficient to get the fixed RoI. Notice that the last two digits of the meter reading are decimal places. We have not included the last digit into RoI as this part is the most ambiguous and sometimes even hard for humans to detect the reading. Again, as mentioned in the dataset section, we use perspective transformation to straighten the RoI.

2) *Image Pre-processing:* Recognizing the location of digits from the RoI is a challenging task. The RoI can be noisy and blurred because of the dusty environment and presence of dew on the water meter. Hence we need to pre-process the image beforehand to find the location of the digits. We use the following methods to pre-process the image: i) grayscale [11], ii) median blur [11] iii) adaptive thresholding [11].

Initially, the RGB colored image is converted to a grayscale image to reduce the complexity and computation overhead: from a 3-dimensional pixel value (R, G, B) to a 1-dimensional value. Later on, the grayscaled image is median blurred to smoothen out the edges. Hence, all the high-frequency components (noise) of the image will be removed. Finally, adaptive thresholding followed by morphological operation (dilation) [11] was performed on the blurred images to separate desirable foreground image objects (digits) from the background based on the difference in each region's pixel intensities. After this step, from Fig. 3(c), it can be noticed that a closed curve is successfully formed around the digits, which will further help in determining the location of the digits.

3) Digit Image Segmentation: The preprocessed image was used to find the location of the digits in the RoI. We have used the closed curves formed outside the digits on the water meter to determine the location. In any image, Contours are curves or the continuous lines that join all the continuous points, having the same color or intensity, to bound an object's complete boundary in the image. As in our case, the curves are formed around digits, finding them will provide the digit's locations. The contour retrieval mode was set to retrieve only the outer contours, so only the outermost is given, in case we have one contour enclosing another (like concentric circles). The contour approximation method is set to remove all redundant points and compress the contour to save memory. However, there are other contours as well in the preprocessed image that are not formed around the digits and they are discarded based on the contour area. The contours formed on the RoI image is shown in Fig. 3(d).

After image-processing and detecting the selected region, each digit present in that contour is segregated and extracted. This was made possible cause the contour stores the coordinates information as well. This process is called image segmentation, in which the image is partitioned into different regions based on a common feature, in our case, the digits in the selected region. As each contour is formed around digits only, the contour coordinates are sorted from left to right based on their position. Fig. 3(e) shows the segmented digit images. In the next step, each extracted digit image is detected using the trained ML model.

4) Digit Recognition and Correction: The last step of the proposed system is the recognition process for the meterreading digit images obtained in the previous step. The digit images that were segmented with the contour method are now passed to the trained RF model that we created while training. The digit image features were computed using the HOG feature extractor whose output is $n \times 1$. As shown in Fig. 3(h), the RF model's input is one digit image feature at a time, and the output is the corresponding digit.

As the digits may be detected wrongly, two digit correction mechanisms as postprocessing are applied to the number collected after combining the predicted digits. These mechanisms are defined as follows:

 As this is continuous real-time chronological data, it is assumed that the current value must be greater than or equal to the previous value. This helps us to mitigate the common detection error that occurred while detecting the digits. Another assumption is made that the flow of the water can not increase suddenly by a huge number. Leveraging this fact, the digits detected are adjusted based on the previous flow.

V. RESULTS AND OBSERVATIONS

The developed model was deployed for ten days (Mar 11 -Mar 21, 2021) on a water meter at the pump-house of IIIT-H campus. Water from this pumping station is delivered to a residential area of forty families where it is stored in the overhead tanks. During the initial three days, the captured images were used for training the ML model. The trained model was later used for the next seven days to collect the flow data in real-time. The HOG-based image feature extraction was done using *skimage* [15]. The orientation value was set to 9 with 8 pixels per cell and two cells per block. The RF classifier was implemented using Scikit Learn [16], a popular python-based ML library. The criterion hyperparameter of the RF classifier, which measures the quality of split while training was set to entropy. For all kinds of image pre-processing, e.g., conversion to grayscale, median blurring (with window size 15), and adaptive thresholding, the popular computer vision library OpenCV [17] was used.

Figs. 4(a) and 4(b) show the sensed meter readings and flow rate w.r.t time for the retrofitted water meter, respectively. The time axis is plotted such that every tick marks the beginning of the day with 0000 hrs. Therefore the time elapsed between two consecutive ticks is 24 hrs. Every pumping instance is marked by a rise in the meter reading as well as the flow rate. The meter reading remained constant, and the flow rate remained zero for the duration when the pump was not running and, consequently, no water was flowing. It is observed that the water was pumped four times in the period of observation. The duration for which the water is usually pumped is significantly less when compared to the total duration of observation. Hence, although the pumping duration is few tens of minutes for every instance, it appears as a sudden step in the meter reading or a sharp peak in the flow rate. This is a valid observation because the campus has an automated pumping system. The water is pumped to the residential area's overhead tanks only when the tank's levels dip below a certain threshold. The various instances of water flow along with the duration and volume are shown in Table II. It is worth observing from Fig. 4(b) (Flow rate plot) that the flow rate does not remain constant for the entire duration of water flow; instead, it usually keeps varying between 0.1 KL/minute and 0.2 KL/minute. It means that there is no deterministic linear relationship between the duration and the volume of water flow. Instead, it is governed by the entirely random demand for water. The observations of Table II validate the same.

The proposed image processing algorithm is tested on the manually annotated test dataset as explained in Section III. The trained RF model achieved a validation accuracy of 97.69%. Even though the RF model is trained at very high accuracy, some digit recognition errors were noticed. Most recognition errors occur for the digits appearing at the tenths (first digit



(a) Meter reading (volume of water flown) w.r.t time, assuming 0000 hrs as the start of the day on time axis



(b) Flow rate (KL/minute) w.r.t time, assuming 0000 hrs as the start of the day on time axis

Fig. 4. Plots of meter reading (volume in KL) and flow rate (KL/minute) for the observation period from 14 Mar 2021, 1525 hrs to 21 Mar 2021, 2359 hrs.

TABLE II WATER FLOW TIME, DURATION AND VOLUME DURING THE PERIOD OF OBSERVATION

S.No.	Start Time (dd-mm, hrs)	Stop Time (dd-mm, hrs)	Duration (mins)	Volume (KL)		
1	15-03, 1239	15-03, 1321	42	7.5		
2	17-03, 1856	17-03, 1908	12	2		
3	18-03, 1544	18-03, 1606	22	3.7		
4	19-03, 1654	19-03, 1708	14	2.1		

TABLE III DIGIT ERROR RATE

Digit	0	1	2	3	4	5	6	7	8	9
Error Rate (%)	0.0	0.0	0.0	1.3	0.6	0.4	5.0	5.2	3.2	0.0

after the decimal) place or units place as they change most frequently. It is reasonable in the rotating-disc type meters. Images captured during the digit change process will lead to an ambiguous situation, where half of the previous digit and half of the incoming digit will be visible in the image. The errors due to rotating-disc issues are efficiently resolved by the digit correction techniques discussed in Section IV-B4.

Next, the proposed solution's performance is evaluated by calculating the following three parameters.

- Digit Error Rate (DER): It is defined as the number of times a particular digit is unrecognised by the total number of times a particular digit appears for recognition. This parameter helps evaluate the ability of the ML model in recognising the digit's image correctly. The obtained DER for all the digits is shown in Table III. It can be inferred that the DER is comparatively higher for digits whose shape resembles any other digit's shape.
- 2) Value Error Rate (VER): It is defined as the number of times the complete meter reading or value is incorrectly recognised by the total number of readings recognised. A meter reading is considered incorrectly recognised if it has one or more unrecognised individual digits. VER helps in evaluating the performance of the applied postprocessing techniques. Since the obtained meter values can only increase or remain unchanged, digit errors can be corrected to improve VER. The proposed postprocessing ideas achieve a VER of 4.49%.
- 3) Root Mean Squared Error (RMSE): It is defined as the square-root of the average squared difference between the actual meter reading and the reading obtained from the ML algorithm. This parameter helps in evaluating the severity of the unrecognised meter reading. The RMSE is calculated before and after the postprocessing to check the effectiveness of the technique. The proposed ML algorithm achieved an RMSE of 0.0748 KL before postprocessing and 0.0361 KL after postprocessing. It can be inferred that the postprocessing technique has improved the RMSE. Moreover, it is observed that the RMSE is significantly low when seen in the light of the discussed application.

The above results and observations show that DER does not directly translate into high VER. Moreover, the achieved VER does not lead to high errors in the analysis as most of the errors occur in at least significant digits.

VI. CONCLUSION

This paper presented an IoT-based method for retrofitting analog water meters to convert them into smart digital meters. An image of the analog meter is captured at high frequency to extract the meter reading in real-time. The RF learning method is used to perform image classification for extracting the meter readings at the edge. The sensed readings are stored on a cloud-based platform for analysis purposes. The proposed model was deployed at the IIIT-H campus for realtime analysis and data collection. The image processing algorithm achieved an accuracy of 97.69% for digit recognition. Application-specific post-processing mechanisms achieved a low VER of 4.49% and low RMSE of 0.0361 KL. The measurements done over ten days clearly show the effectiveness of the deployed system in generating flow-volume and flowrate data. In future, it is intended to expand the deployment over all the water meters on the campus. Deep learning-based image processing techniques can also be developed to improve digit recognition further. The collected data can generate water consumption patterns of the community and help in the timely detection of leaks/faults.

ACKNOWLEDGEMENT

This research was supported partly by National Geospatial Programme (NGP), India, under grant no. 2073 (2020), PRIF Social Incubator Program (2019) and the Ministry of Electronics and Information Technology (MEITY), Govt. of India under grant no. 3070665 (2020), with no conflict of interests.

REFERENCES

- M. J. Mudumbe and A. M. Abu-Mahfouz, "Smart water meter system for user-centric consumption measurement," in 2015 IEEE 13th International Conference on Industrial Informatics (INDIN), 2015, pp. 993–998.
- [2] H. J. C. Barreto, Kurtulan, S. İnci, M. Nakıp, and C. Güzeliş, "Converting utility meters from analogue to smart based on deep learning models," 2020 Innovations in Intelligent Systems and Applications Conference, 2020.
- [3] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016.
- [4] Y. LeCun et al., "Lenet-5, convolutional neural networks," URL: http://yann. lecun. com/exdb/lenet, vol. 20, no. 5, 2015.
- [5] M. Vanetti, I. Gallo, and A. Nodari, "Gas meter reading from real world images using a multi-net system," *Pattern Recognition Letters*, vol. 34, no. 5, pp. 519–526, 2013. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167865512003832
- [6] S. Bajaj, C. Teja S., and P. K. Yemula, "Computer vision based energy monitoring system using meter image capturing system (MICAPS)," in 2020 First International Conference on Power, Control and Computing Technologies (ICPC2T), 2020.
- [7] Raspberry-Pi micro-controller module, accessed 21 Feb. 2021. [Online]. Available: https://raspberrypi.org/
- [8] ThingSpeak: Cloud data collection and analysis using MATLAB, accessed 27 Feb. 2021. [Online]. Available: https://thingspeak.org
- [9] JioFi JMR540: LTE based portable WiFi hotspot, accessed 27 Feb. 2021. [Online]. Available: https://www.jio.com/shop/en-in/ router-jmr540-black-/p/491193576
- [10] Bureau of Indian Standards: Product certification, accessed 12 Mar. 2021. [Online]. Available: https://bis.gov.in/index.php/ product-certification/
- [11] R. C. Gonzalez and R. E. Woods, Digital Image Processing (3rd Edition). USA: Prentice-Hall, Inc., 2006.
- [12] L. Breiman, "Random forests," Machine learning 45.1 (2001): 5-32, 2001.
- [13] R. O. Duda, P. E. Hart et al., Pattern classification. John Wiley & Sons, 2006.
- [14] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2005.
- [15] S. van der Walt, J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, T. Yu, and the scikit-image contributors, "scikit-image: image processing in Python," *PeerJ*, accessed 10 Feb. 2021. [Online]. Available: https://scikit-image.org
- [16] Scikit-Learn: Machine learning in Python, accessed 10 Feb. 2021. [Online]. Available: https://scikit-learn.org
- [17] OpenCV: Python library for real-time computer vision, accessed 10 Feb. 2021. [Online]. Available: https://opencv.org