

## Learning to Read Analog Gauges from Synthetic Data

Juan Leon-Alcazar

KAUST

Thuwal, Saudi Arabia

juancarlo.alcazar@kaust.edu.sa

Yazeed Alnumay

Aramco

Thuwal, Saudi Arabia

yazeed.alnumay@aramco.com

Cheng Zheng

KAUST

Thuwal, Saudi Arabia

cheng.zheng@kaust.edu.sa

Hassane Trigui

Aramco

Thuwal, Saudi Arabia

hassane.trigui@aramco.com

Sahejad Patel

Aramco

Thuwal, Saudi Arabia

sahejad.patel@aramco.com

Bernard Ghanem

KAUST

Thuwal, Saudi Arabia

bernard.ghanem@kaust.edu.sa

### Abstract

Manually reading and logging gauge data is time-inefficient, and the effort increases according to the number of gauges available. We present a pipeline that automates the reading of analog gauges. We propose a two-stage CNN pipeline that identifies the key structural components of an analog gauge and outputs an angular reading. To facilitate the training of our approach, a synthetic dataset is generated thus obtaining a set of realistic analog gauges with their corresponding annotation. To validate our proposal, an additional real-world dataset was collected with 4,813 manually curated images. When compared against state-of-the-art methodologies, our method shows a significant improvement of  $4.55^\circ$  in the average error, which is a 52% relative improvement. The resources for this project will be made available at: <https://github.com/fuankarion/automatic-gauge-reading>.

### 1. Introduction

Analog gauges are widespread across modern industrial facilities. Typically, these gauges monitor critical operational parameters such as temperature, pressure, and level indicators for active industrial processes. Therefore, keeping an accurate record of gauge reading data is essential to track an asset's trends and its conditions over time, which would aid in failure investigation if one occurs.

Manually reading and logging gauge data is time-inefficient, and increases according to the number of gauges available. To approach this problem, digital transmitters allow the remote monitoring of gauges. However, in many cases, the facilities are aged, or it is prohibitively expensive to upgrade numerous analog gauges to digital ones. Additionally, on-site readings of gauges serve as a direct verification means for validating the accuracy of digital transmit-



Figure 1. **Sample Results.** The top row shows samples of gauges captured on the field. The bottom row displays our model's predicted segmentation and regression of landmarks in these images. The segmentation result has three classes: background (not highlighted for easier visualization), gauge (cyan), and needle pointer (magenta). The gauge overlay also contains circles which visualize the estimated locations of the • start marker, • needle pointer, and • end marker.

ters, thus ensuring redundancy and process safety.

Since on-site readings can not be completely removed from the verification pipeline, some research efforts have explored the automatic reading of gauges from image data or video data captured directly in the facility [1, 17]. These studies are largely limited by the availability of datasets with real-world data and curated ground-truth annotations. Unlike most modern machine learning datasets [8, 15, 21], analog gauge reading data is still small and insufficient to train deep convolutional models. This data deficiency can be explained as it is expensive to collect and manually label a large number of real-world gauges. Moreover, privacy and security issues arise if these gauges are actively in use in an industrial facility.

In this paper, we approach the automatic analog gauge reading task and formulate a computer vision pipeline that

enables the direct reading of analog gauges by processing a single picture of the corresponding device. Thereby providing automated readings which facilitates the digitization process, and simplifies the operator intervention. Our approach focuses on generating synthetic gauge images as a surrogate training set, and formulating a training time strategy that enables generalization from synthetic training data into real-world gauges. Figure 1 contains some sample results from our method applied to real-world gauges.

We construct a synthetic gauge dataset by leveraging 3D modeling software. A single 3D model allows us to render images of various types of gauges from different view angles and simulates common noise sources. While synthetic gauge modeling mitigates the scarcity of training data, it poses a new challenge as it creates a domain gap between the rendered data (training) and the real-world examples (validation). With careful image augmentation, this domain gap can be effectively reduced, resulting in a model that can successfully operate in real-world scenarios. We validate the performance of our proposal on two real-world datasets containing over 7.500 labeled images.

Our paper brings 3 contributions. i) We develop a 3D model that simulates analog gauges and can render images that contain the spatial location of the gauge’s main components, and its reading. ii) We propose a two-stage (attention and classification) pipeline that identifies the key structural elements of the gauge and outputs a final reading. With our training data this two-step approach produces state-of-the-art results. iii) We create a real-world dataset with 4.813 gauge readings which are manually curated and include multiple capture conditions and noise sources.

## 2. Related Work

Automatic analog gauge reading is a machine vision task that involves object detection and landmark regression [12]. Despite having a well-defined structure and goal, its solution is non-trivial. Therefore, multiple research efforts have addressed this problem. Overall, these approaches can be split into two main categories: classical computer vision and Deep Learning methods.

Most of the traditional computer vision methods [1, 14, 17, 24, 28, 30] rely on the shape of the gauge (typically circular), and apply the Hough Transform [13] to estimate an initial gauge location. Upon this initial step, most of the approaches follow a secondary detection step which aims at establishing the location of key landmarks inside the gauge, commonly the pointer needle or the start and end markers [12]. The work of Bao *et al.* [2] uses adaptive thresholding to enable a second Hough transform step that approximates the location of the pointer needle. The approach of Yi *et al.* [31] proposed to use K-means clustering to identify the gauge elements, a similar method is outlined in [17] where PCA and clustering are used to estimate the needle location

on a binary image. The needle location step has also been approached by template matching and hand-crafted feature extraction (SIFT and RANSAC) [25]. Finally, some works have analyzed this task in video streams [5] and rely on the Optical Flow for temporal stability on the predictions.

Following the success of Deep Learning methods for natural image processing [11, 16], recent approaches have started training deep models that can directly regress the readings on a gauge. Some implementations rely on Deep Networks only for the detection procedure, commonly a CNN tuned for gauge detection [9, 29], and apply a Neural Network trained for Optical Character Recognition in order to parse the scale values [9, 20].

Currently, most of these Deep Learning approaches suffer from a lack of training data and can easily overfit the visual appearance of the gauges in the training set [22, 23]. This lack of generalization also holds for deep models that were trained on hybrid data, where the pointer needles of real gauges are artificially rotated [4, 29]. Recently, [12] devised a four-keypoint detection model to regress the minimum and maximum markers and the needle center and tip locations. Their model was trained on a large fully synthetic dataset and showed improved generalization to real images.

In contrast to the current approaches, we devise a simple, yet effective deep-learning pipeline with only two-trainable stages. We show that this strategy can be trained exclusively on synthetic data and achieve state-of-the-art results on real-world datasets.

## 3. Automatic Gauge Reading

In this section we outline the 3 stages of our approach: i) synthetic training data, ii) gauge component location, and iii) gauge landmark reading. The first component generates surrogate training data, by rendering synthetic gauges from a 3D model. The last two components are instantiated as convolutional neural networks [16, 18, 19]. These networks aim at estimating a fine-grained label of all the gauge components, and then approximate the location of 3 key landmarks in the gauge: start marker, end marker, and the pointer needle’s tip.

### 3.1. Synthetic Training Data

We decide against collecting training samples from gauges and manually labeling their readings. Such a process is expensive and slow and can be difficult to implement given the security constraints in some industrial facilities. Additionally, after gathering and labeling a large amount of data, we would still have to assure that it contains visually diverse gauges, and includes multiple capture conditions. Such diversity would mitigate the risk of overfitting to the global visual appearance of the collected gauges, instead of approximating the key landmarks on it.



Figure 2. **Synthetic Gauge Dataset.** We use the 3D creation suite Blender to generate diverse synthetic training samples. We design the rendered samples to include common artifacts and sources of noise, which makes the automatic gauge reading task more challenging. For example, we simulate light reflections on the gauge crystal (first row). We can include different patterns of dust/sand coverage (second row). We also control the needle’s shadow, which can generate alignment errors in the landmark detection task (third row).

For the training phase, we rely on synthetic data that approximates the visual appearance of an analog gauge. We name this data split the *synthetic dataset*. We approach this task with the open-source 3D creation suite Blender [7]. We design a 3D model of a circular gauge which allows us to manipulate key elements of its appearance like needle type, case style, sticker min and max values, size, tint on the front crystal, and simulated dust particles amongst many others. Figure 2 contains some samples of the synthetic dataset.

The flexibility of Blender allows us to freely mix and manipulate all the model elements in order to create multiple unique synthetic gauges. Moreover, we can adjust the camera parameters and lighting patterns while rendering an arbitrary view of the synthetic gauge. Finally, we can also control the needle orientation, which allows us to simulate multiple ground-truth values. This flexibility enables us to train convolutional networks in a fully supervised manner, since the ground-truth location of all the gauge components and the landmark angular locations are generated along with the rendered training data.

We build a Blender script to randomize all gauge attributes and render the training dataset. The resulting synthetic dataset contains 12,000 images. Along with the rendered image, the script generates the corresponding ground-truth segmentation mask with 3 classes (background, gauge, needle), the angular location of the needle’s tip, and the maximum and minimum markers of the gauge sticker.

### 3.2. Gauge Detection and Reading

Our approach for automatic gauge reading is a two-step classification pipeline. We first establish an approximate location of the most relevant gauge physical components, this

estimation is then used as a spatial guidance to a second stage that effectively estimates the location of the key landmarks in the gauge, thus enabling a final angular reading.

In practice, the first step is approached by performing semantic segmentation on the gauge images, we estimate a pixel-wise mask of the gauge, thus partitioning it into 3 classes: Gauge Case, Needle, and Background. The second step is approached as a classification problem, where we establish the angular location of three key-points in the gauge: start marker, end marker, and pointer needle location. The combination of these 3 key-points provides an estimation of the position of the needle relative to the start and end of the gauge scale.

**Segmentation Network** In the location step, we label every individual pixel in the input image. As a result, we obtain an approximate spatial location of the gauge in the picture (given by the gauge case class), and an initial location of the most relevant landmark in the process, namely the pixels assigned to the needle class. We approach the segmentation step using DeepLabv3 [6], with a ResNet-18 backbone [10]. At training time, we rely exclusively on the synthetic data and the associated ground-truth.

**Reading Network** In the reading step our main goal is to localize 3 landmarks: the start marker, the end marker, and the needle tip (see the blue dots in Figure 3 (g)). The ground-truth for these targets is provided in the synthetic dataset as angles between 0 to 359. We parameterize all three labels as the angle created between the respective gauge landmark and a horizontal line in the middle of the image, which we define as a 0° point (see the dashed pink lines Figure 3 (g)).

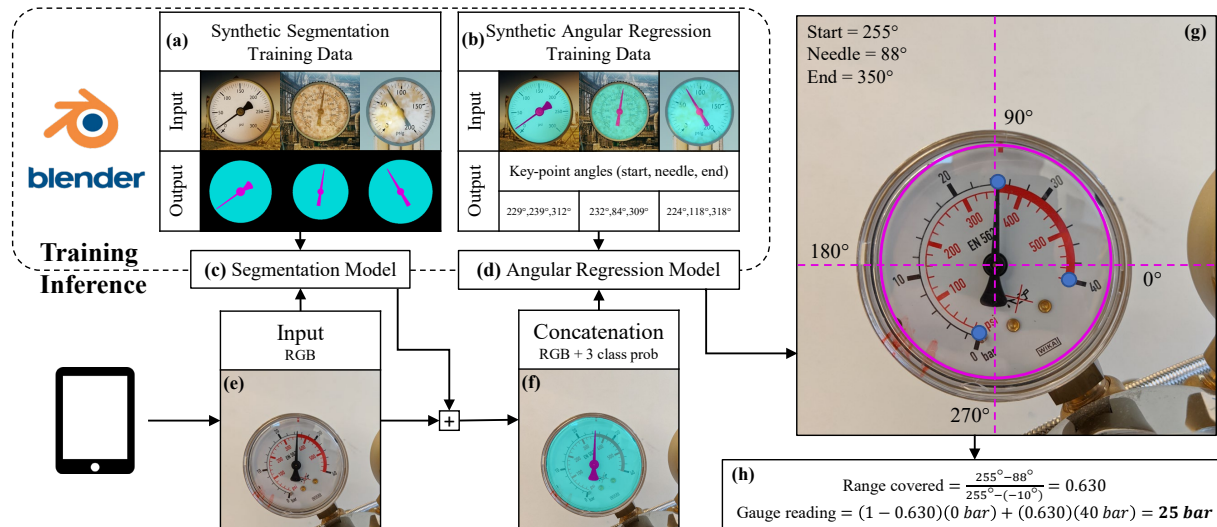


Figure 3. **Method Overview.** We generate synthetic training data with diverse appearances and levels of noise from the Blender component (a). Additionally, we render the ground-truth semantic segmentation and output the location of the gauge key-points For all the synthetic data (b). At training time, we first optimize a segmentation network (c) that provides a localization prior for the gauge components. Then we use a concatenated tensor of the synthetic image and its estimated segmentation mask (f) to train a reading network that approximates the angular location of the pointer needle and the sticker start and end points (d). The angular locations of these points along with the known gauge unit and range are used to determine the gauge’s reading (g, h).

We jointly model all these prediction targets with a single convolutional backbone augmented with three independent prediction heads (one head per prediction target). Although this task could be approached as a regression problem (*i.e* predict a number between 0 to 359), we favor the empirical stability of the gradients created by the cross-entropy loss and approach this problem as a classification task (*i.e* select a class between 0 to 359). Therefore, we quantize the ground-truth values into 360 bins (one bin representing one degree) and optimize the task as a classification problem. To this end, we chose the MobilnetV2 [26] or EfficientNet [27] as the backbones and complement each encoder with three linear heads, each head has an output size of  $1 \times 360$ .

At training time, we leverage the location prior provided by the segmentation step by appending the segmentation result to the input image in the channel dimension. This segmentation provides a form of visual guidance that allows us to direct the network’s attention toward the gauge instead of the background. We extend the filters on the first layer of each backbone to accept 4 channel inputs, the new filter data is initially set to the average of the original 3 channel filter. Additional training details are provided in the **supplementary material**.

## 4. Gauge Reading Datasets

In this section, we outline our second dataset contribution, namely two real-world datasets that contain manually curated ground-truth labels. The first one is the *land-mark dataset* which contains labels for the location of the start

and end markers, along with the pointer needle orientation on 4.813 real-world images. The landmark dataset serves as a held-out test-set to validate our proposed approach. The second dataset provides the semantic segmentation ground-truth for 59 real-world images, it serves as a validation set for the semantic segmentation step.

### 4.1. Landmark Regression Dataset

The landmark dataset consists of 5 video sequences captured using a GoPro camera and an analog gauge attached to the end of a Novint Falcon haptic controller, which was adapted to serve as a robotic arm. This mechanism enables us to control the location of the gauge in the Y-axis (up and down the camera plane) and Z-axis (away or close to the camera plane). This allows us to simulate the capture of gauge images from diverse points of view and varying distances. All the image data is captured and encoded as a video sequence at 720p ( $1280 \times 720$ ) resolution.

In addition to the capture viewpoint, the gauge is manipulated with two additional servomotors controlled by an Arduino Uno microcontroller. The first servomotor controls the orientation of the pointer needle allowing us to control the actual gauge reading. As the servomotor is not directly connected to the needle, but rather is attached to the internal mechanical structure of the gauge. The mapping of the servo motor rotation and the needle’s effective rotation is non-linear. To approximate this mapping, we sampled and manually recorded the amount of motor rotation input required for the needle to match a specific reading (tick

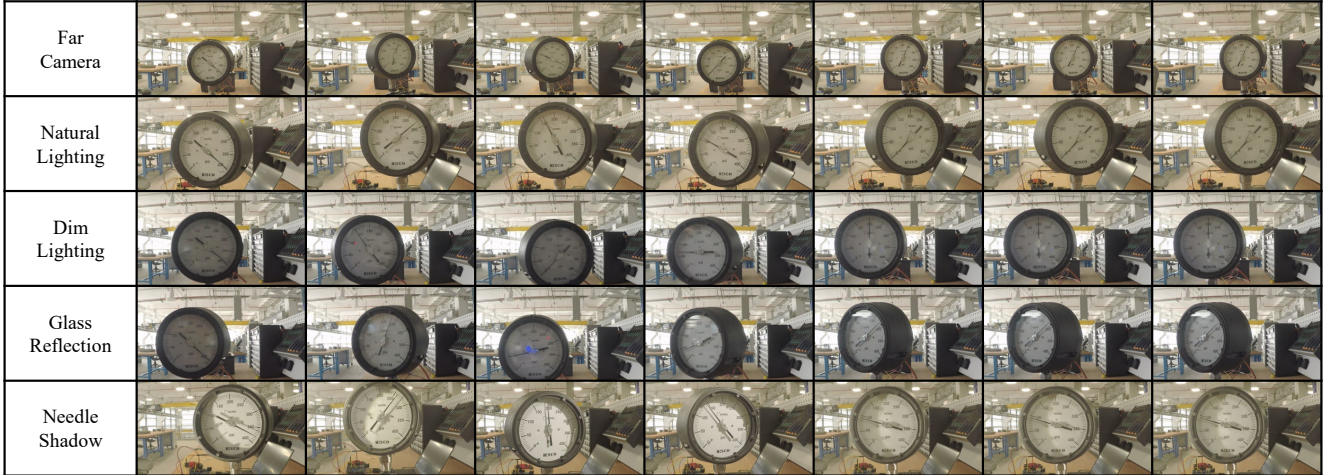


Figure 4. **Landmark Detection Dataset.** We show some samples of the Landmark detection dataset, each image has an associated ground truth that contains the needle angle, and the location of the start and end of the sticker. We create this dataset using a mechanical arm where we control the relative position of the gauge and the camera, along with the location of the pointer needle. By synchronizing the motion of the arm and the servomotors with the capture device, we can obtain an exact ground truth reading for every frame.

marker) on the gauge’s face. Then we linearly interpolate the motor inputs to approximate intermediate angles. The second servomotor can rotate the gauge’s face relative to the camera plane (the gauge face can be oriented towards the left or right of the Field of View, which represents the yaw rotation).

To account for some real-world capture conditions and noise, 5 different environments were simulated. These scenarios include varying light sources and camera locations. In total, each scenario contains between 960 to 966 images. A few samples from the landmark dataset in each of the 5 proposed scenarios are presented in Figure 4.

**Natural Lighting.** Serves as the baseline scenario without any complex conditions. We arrange some artificial lighting that approximates an outdoor setup by daylight and locate the camera close to the gauge (30 cm away). As a result, the case of the gauge occupies on average 30% of the image’s area. Clearly, the gauge is the main object in the image, but there is some background clutter on the capture.

**Far Camera.** Simulates the model’s performance when reading slightly more distant gauges. These images were captured 50 cm away from the gauge. On average, the gauge’s case occupies 10% of the image area, which means the global image features are dominated by the background.

**Dim Light.** In this scenario we turn off the artificial lighting, such that the capture condition simulates the use of the gauge reading tool indoors, or with diminished sunlight. The reduced lighting paired with the relatively high (25 fps) video capture introduces some noticeable motion blur artifacts in this scenario, making the location of the landmarks more challenging.

**Glass Reflection.** This scenario was captured with an open window allowing for direct sunlight towards the the gauge, additional objects were added to the scene to create reflections that could alter the appearance of the gauge.

**Needle Shadow.** For this scenario, we set the artificial light source much closer to the gauge and arrange the lighting direction at an angle (never fully frontal). This setup creates sharp shadows in the gauge sticker, mainly from the gauge case and the needle. These conditions often lead to erroneous estimations, as automatic reading models tend to miss-classify the needle’s shadow as the actual needle.

All the gauge readings in the landmark dataset are controlled by the servomotors and carefully synchronized with the video capture. Therefore, our landmark dataset offers a standard benchmark with accurate ground truth labels that enables the validation (and potentially training) of automatic gauge reading tools. Overall, the landmark dataset presents a diverse set of possible failure cases to test against, although it has a bias regarding the gauge type as the same gauge was used for all 5 capture scenarios.

## 4.2. Segmented Gauge Dataset

Our proposed approach relies on image segmentation of the background, gauge, and pointer needle. To validate the performance of our methods on this initial task, we manually segmented 59 real images. Of these 59 images, 43 are part of the Landmark Dataset, the remaining 16 were obtained from other gauges inside an industrial facility. The label distribution in this dataset is highly imbalanced, with the background, needle, and gauge classes representing 66.26%, 1.03%, and 32.71% of the pixels, respectively. We provide this dataset as the segmentation of



Figure 5. **Segmented Gauge Dataset.** We provide a dataset with 59 manually segmented real-world gauges. This dataset includes multiple indoor and outdoor gauges and supports the validation of the attention step in our pipeline. In this figure the Gauge pixels are highlighted in Bright cyan, while the pointer needle pixels are shown in magenta, we do not highlight the background pixels for easier visualization.

gauges is a key step in our approach. In addition, it will also support future automatic reading methods that build upon a location step. We visualize some samples of the segmentations contained in this dataset in Figure 5.

## 5. Results and Analysis

We now perform the empirical evaluation of our method, we begin with a direct comparison against the state-of-the-art in Howells *et al.* [12]. Then, we provide the evaluation of our approach in the proposed Landmark dataset. We conclude by assessing the effectiveness of the attention component using our segmentation dataset.

### 5.1. Comparison against the State-of-the-art

The work of Howells *et al.* [12] provides its own real-world angle reading dataset. We run our proposed approach in their test-set and perform a direct comparison. Their test-set consists of 6 visually distinct gauges, for each individual gauge there are 3 video clips containing a total of 450 frames (150 frames per video clip, 2,700 images total). Each clip shows a static capture (i.e. no camera or gauge motion, no illumination changes) of the gauge where only the pointer needle changes its orientation. The associated angle ground-truth is available for every clip.

The gauge reading dataset in [12] was deliberately collected with small gauges, occupying about 10-15% of the image, such that reading pipelines first crop out the target gauge and then regress the needle’s location. Since the semantic segmentation step in our pipeline works as an attention mechanism rather than a detection step, we use OpenCV’s implementation of the Hough Transform circle detection [3, 13] to extract an initial crop of the image. We keep the default settings for the algorithm and directly apply our method to the cropped detections. Table 1 summarizes the results of our approach in the test-set of [12].

Overall, we observe that our proposal outperforms the state-of-the-art by at least  $1.67^\circ$ , which represents a relative improvement of 17%. Although [12] outperforms our method in 3 sequences (c, e, f), we observe that the results are nearly identical in those scenarios as their absolute differences are  $0.03^\circ$ ,  $0.62^\circ$ , and  $0.33^\circ$  respectively. Meanwhile, our method largely outperforms their baseline in the more challenging sequences (b and d), reducing the error

Dataset	Ours		Howells <i>et al.</i>
	Small	Large	
meter_a	19.33	16.62	9.24
meter_a (rev.)	5.69	<b>2.50</b>	—
meter_b	<b>8.23</b>	8.29	26.06
meter_c	2.74	1.97	<b>1.94</b>
meter_d	7.68	<b>7.36</b>	11.70
meter_e	4.32	4.83	<b>3.70</b>
meter_f	4.64	4.69	<b>4.31</b>
<b>Average</b>	7.82	<b>7.30</b>	9.49
<b>Average (rev.)</b>	5.55	<b>4.94</b>	—

Table 1. **State-of-the-art Comparison.** We test our pipeline in the test-set of [12]. We report the mean absolute error of the predicted gauge angle readings in degrees. Our small model uses MobileNetV2, which is the same model used by [12]. The large model uses the EfficientNet-B2 backbone.

by  $17.83^\circ$  (68.4% relative improvement) in sequence b, and  $4.34^\circ$  (37.1% relative error reduction).

**Ground-truth and Predictions Design** We note that the test-set of [12] follows a different approach to our training set. While we use every angle from  $0^\circ$  to  $360^\circ$  in our labels and prediction heads, Howells *et al.* use labels and predictions between  $0^\circ$  to  $180^\circ$ . Their angle prediction formula from four key-points always takes the smallest of the two possible angles between the start marker and needle tip. A visualization of this discrepancy is presented in Figure 6.

Our model’s classification head with angles in the range  $[0^\circ, 360^\circ)$  overcomes the aforementioned limitation. Moreover, the same four key-point method proposed by [12] can be modified to incorporate the sign of the 2D cross-product of the central edges  $\vec{CS} \times \vec{CT}$  to determine if the angle  $\theta$  should be used or its explement ( $360^\circ - \theta$ ). We correct for this discrepancy on every ground-truth angle that is above  $180^\circ$ . Additional benchmark ground-truth discussion is provided in the **supplementary material**.

**Ground-truth Correction** In Table 1 we include two different evaluations for meter\_a, we identify that the sequence meter\_a, contains flawed ground-truth data that does not match the actual reading of the gauge. We manually inspect the entire sequence and identify that this annotation error

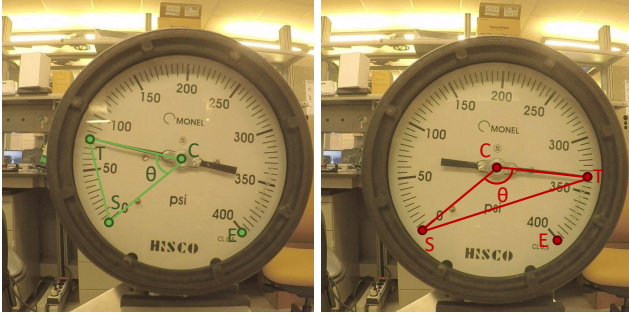


Figure 6. **Benchmark Ground-Truth Discrepancies.** The angle extraction method used in [12] extracts the explement angle between the start marker and needle tip when the angle exceeds  $180^\circ$ . The left figure shows angles for which both methods succeed. The right figure shows an apparent failure case of our proposed model, given the design of the ground-truth labels.

appears on 40 contiguous frames of the first video capture and on the entire 150 frames of the second. We detail the corrections made in the **supplementary material**.

We create a revisited (rev) version of meter\_a, with ground-truth values that match the needle location. For completeness, we report the performance of our method on both the full sequence and the revisited subset. We observe a significant gap when the revisited data (rev) is used. Furthermore, we also break down the average performance according to the selected sequence\_a data (original or revisited). With the revisited data, our approach increases the performance gap against the state-of-the-art by up to  $4.55^\circ$  in the average error (52% relative improvement). Figure 7 shows our model’s performance on two videos of the meter\_a, video 3 uses the original ground truth, while video2 uses our revisited version. In the **supplementary material**, we include a similar analysis for every video in the dataset.

## 5.2. Landmark Dataset Evaluation

We now assess the effectiveness of our pipeline in the landmark dataset. Unlike the test-set of [12], we directly apply our method to the test data, *i.e.* we do not perform any additional detection step or image cropping. For this evaluation, we also test how our synthetic dataset enables the fine-tuning of models deeper than MobileNetV2 [26]. To this end, we choose to fine-tune the EfficientNet [27], Table 2 summarizes the results.

The best results are obtained with the EfficientNet B2 Network, which corresponds to our Large model in Table 1. Although the MobilnetV2 (small model) underperforms the large model in every scenario, the small model shows competitive results in the presence of shadows or low light. Our benchmark shows that the baseline scenario (Natural Lighting) reports the best results, whereas our pipeline reduces its performance in the presence of low light or strong shadows.

We observe that the EfficientNet family can be fine-tuned

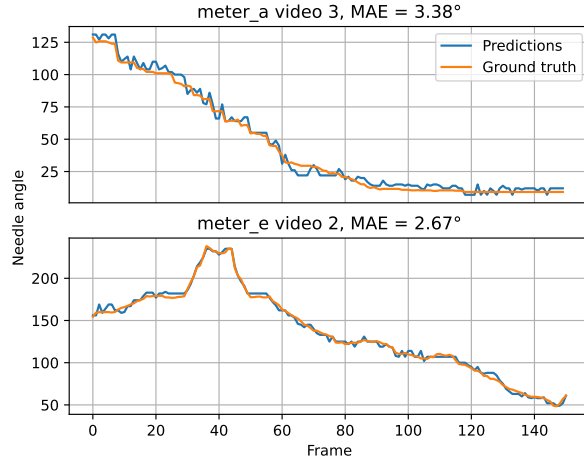


Figure 7. **Benchmark Time Series Prediction Samples.** Our small (MobileNetV2) model’s performance on two sample videos from the benchmark dataset. Our model’s needle angle prediction closely matches the ground-truth. All the time series results are presented in the supplementary material.

with our synthetic data up to the B2 size without showing overfit. The larger B3 model does not report any empirical improvement in any scenario. Our best model reports an average angular distance across all scenarios of  $2.69^\circ$ .

## 5.3. Qualitative Analysis

We tested our large (EfficientNet-B2) and small (MobileNetV2) models on real-world gauges located in an industrial facility. These gauges are mostly outdoors and exhibit various signs of environmental degradation. Figure 8 shows a set of 18 field gauges that were captured under various visual conditions such as extreme brightness, low lighting, reflections, glare, etc. Many gauge glasses were dusty, tinted, cracked, or opaque due to sunlight exposure. Some of the gauges were not easily accessible and were captured from a distance or from non-frontal view angles.

Scenario	Mobile-NetV2	EfficientNet			
		B0	B1	B2	B3
Natural Lighting	2.83	2.62	2.59	<b>2.29</b>	2.31
Far Camera	5.15	2.81	2.93	<b>2.48</b>	2.52
Dim Light	3.62	5.50	4.56	<b>3.20</b>	4.10
Glass Reflection	2.88	2.60	2.64	<b>2.37</b>	2.44
Needle Shadow	3.15	3.36	<b>2.85</b>	3.11	3.14
<b>Average</b>	3.53	3.38	3.11	<b>2.69</b>	2.90

Table 2. **Performance on the Landmark Dataset.** We evaluate our proposed pipeline on the Landmark Dataset using MobilnetV2 and various sizes of the EfficientNet. The cell values represent the mean absolute angular errors in degrees. The best results are obtained with the EfficientNet B2, nevertheless, the other models remain close in terms of absolute distance.



Figure 8. **Qualitative Field Results.** Field samples were used for field sample analysis in Section 5.3 showing our models’ performance on difficult gauges with odd view angles, harsh lighting conditions, and various obstructions. Using the EfficientNet-B2 model, the field gauge images were correctly predicted within 2% of the gauge’s range or two marker ticks of the ground-truth.

Class	Percentage of Dataset	Mean IoU
Background	66.26 %	94.60 %
Gauge Case	32.71 %	88.13 %
Needle Pointer	1.03 %	76.28 %

Table 3. **Segmentation Performance.** Intersection over the union of our method on the segmented dataset from Section 4.2. The segmentation classes are highly imbalanced, with the needle class accounting only for just 1% of the labeled pixels. Consequently, this is the class that is harder to segment.

For this assessment, we include the gauge min and max values, and manually estimate the gauge actual reading (not only the angular locations). Our EfficientNet-B2 model was able to successfully predict all the field samples in Figure 8 with an error of less than 2% of the gauge’s range or two marker ticks of the ground-truth, while the MobileNetV2 model predicted 14 out of the 18 sample images correctly under the same tolerance criterion. The range percentage and marker ticks criterion is commonly used in the field instead of the angular error, it is faster and easier to compute and compare. The EfficientNet-B2 and MobileNetV2 models have achieved 2.3° and 4.46° MAE respectively on this sample of field images.

#### 5.4. Segmentation Results

Finally, we validate the segmentation step using our real-world segmented dataset from Section 4.2. The evaluation of the segmentation performance is presented in Table 3. Overall, the segmentation network achieves a very good result at locating the gauge case, with an IoU over 94% IoU for the background class, and the case IoU just below 89%. Although the needle pointer also shows satisfactory performance, it is the lowest among the 3 classes at 76.3%.

We complement this result with the analysis of the error distribution for each individual class in Figure 9. The gauge case (cyan line) and the background show that most of the images report an IoU over 85% for these two classes. How-

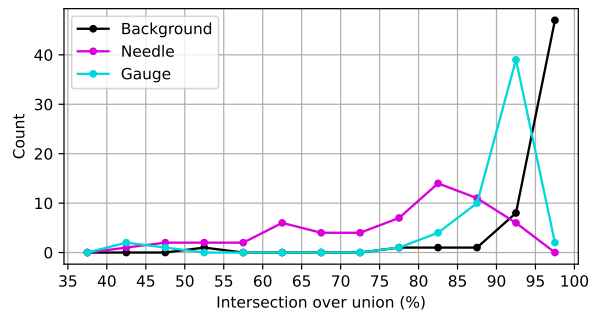


Figure 9. **Segmentation Model Results.** Distribution of segmentation model IoU results on real images from the segmented dataset in Section 4.2. IoU values are binned into 5% ranges, i.e. 90-95%. Naturally, the sparse nature of the needle class resulted in the poorest results.

ever, the needle has a long tail that extends to just above 55%. Although the needle performance is lower in comparison to the other 2 classes, most of the images obtain an approximate segmentation of the needle pointer.

## 6. Conclusion

We have introduced a simple yet effective pipeline for the automatic gauge reading task. By means of a two-stage CNN network that is trained on synthetic data. Our proposal is validated on two real-world datasets showing that synthetic data and augmentations can effectively bridge the domain gap. Our method provides a robust technique to read and directly digitize analog gauges, showing strong performance even in challenging and noisy real-world scenarios.

**Acknowledgment** This publication is based on work funded and supported by the SDAIA center of excellence, and by the Research & Development Center of Saudi Aramco, the authors would like to acknowledge their support in the development of this paper.



## References

- [1] E Corra Alegria and A Cruz Serra. Automatic calibration of analog and digital measuring instruments using computer vision. *IEEE transactions on instrumentation and measurement*, 49(1):94–99, 2000. 1, 2
- [2] Haojing Bao, Qingchang Tan, Siyuan Liu, and Jianwei Miao. Computer vision measurement of pointer meter readings based on inverse perspective mapping. *Applied Sciences*, 9(18):3729, 2019. 2
- [3] G. Bradski. The OpenCV Library. *Dr. Dobb's Journal of Software Tools*, 2000. 6
- [4] Weidong Cai, Bo Ma, Liu Zhang, and Yongming Han. A pointer meter recognition method based on virtual sample generation technology. *Measurement*, 163:107962, 2020. 2
- [5] Shruti Chavan, Xinrui Yu, and Jafar Saniie. High precision analog gauge reader using optical flow and computer vision. In *2022 IEEE International Conference on Electro Information Technology (EIT)*, pages 171–175, 2022. 2
- [6] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018. 3
- [7] Blender Online Community. *Blender - a 3D modelling and rendering package*. Blender Foundation, Stichting Blender Foundation, Amsterdam, 2018. 3
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 1
- [9] Stefan Dumberger, Raimund Edlinger, and Roman Froschauer. Autonomous real-time gauge reading in an industrial environment. In *2020 25th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA)*, volume 1, pages 1281–1284, 2020. 2
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015. 3
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2
- [12] Ben Howells, James Charles, and Roberto Cipolla. Real-time analogue gauge transcription on mobile phone. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 2369–2377, June 2021. 2, 6, 7
- [13] J. Illingworth and J. Kittler. The adaptive hough transform. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-9(5):690–698, 1987. 2, 6
- [14] P Jinuntuya, W Sudatham, and N Jinuntuya. Coordinates transformation method for pointer gauge reading by machine vision. In *Journal of Physics: Conference Series*, volume 2431, pages 012–019. IOP Publishing, 2023. 2
- [15] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 1
- [16] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017. 2
- [17] Jakob S Lauridsen, Julius AG Graasmé, Malte Pedersen, David Getreuer Jensen, Søren Holm Andersen, and Thomas B Moeslund. Reading circular analogue gauges using digital image processing. In *14th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (Visigrapp 2019)*, pages 373–382. SCITEPRESS Digital Library, 2019. 1, 2
- [18] Yann LeCun, Bernhard Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne Hubbard, and Lawrence D Jackel. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989. 2
- [19] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 2
- [20] Zhu Li, Yisha Zhou, Qinghua Sheng, Kunjian Chen, and Jian Huang. A high-robust automatic reading algorithm of pointer meters based on text detection. *Sensors*, 20(20):5946, 2020. 2
- [21] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. 1
- [22] Yue Lin, Qinghua Zhong, and Hailing Sun. A pointer type instrument intelligent reading system design based on convolutional neural networks. *Frontiers in Physics*, 8:618917, 2020. 2
- [23] Yang Liu, Jun Liu, and Yichen Ke. A detection and recognition system of pointer meters in substations based on computer vision. *Measurement*, 152:107333, 2020. 2
- [24] Yifan Ma and Qi Jiang. A robust and high-precision automatic reading algorithm of pointer meters based on machine vision. *Measurement Science and Technology*, 30(1):015401, 2018. 2
- [25] Xiaoming Mai, Wensheng Li, Yan Huang, and Yingyi Yang. An automatic meter reading method based on one-dimensional measuring curve mapping. In *2018 IEEE International Conference of Intelligent Robotic and Control Engineering (IRCE)*, pages 69–73, 2018. 2
- [26] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018. 4, 7
- [27] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019. 4, 7

- [28] Erlin Tian, Huanlong Zhang, and Marlia Mohd Hanafiah. A pointer location algorithm for computer visionbased automatic reading recognition of pointer gauges. *Open Physics*, 17(1):86–92, 2019. [2](#)
- [29] Visarut Trairattanapa, Sasin Phimsiri, Chaitat Utintu, Riu Cherdchusakulcha, Teepakorn Tosawadi, Ek Thamwivatthana, Suchat Tungjitnob, Peemapol Tangamonsiri, Aphisit Takutruera, Apirat Keomeesuan, Tanapoom Jitnaknan, and Vasin Suttichaya. Real-time multiple analog gauges reader for an autonomous robot application. In *2022 17th International Joint Symposium on Artificial Intelligence and Natural Language Processing (iSAI-NLP)*, pages 1–6, 2022. [2](#)
- [30] Junzhe Wang, Jian Huang, and Rong Cheng. Automatic reading system for analog instruments based on computer vision and inspection robot for power plant. In *2018 10th International Conference on Modelling, Identification and Control (ICMIC)*, pages 1–6. IEEE, 2018. [2](#)
- [31] Ming Yi, Zhenhua Yang, Fengyu Guo, and Jialin Liu. A clustering-based algorithm for automatic detection of automobile dashboard. In *IECON 2017-43rd Annual Conference of the IEEE Industrial Electronics Society*, pages 3259–3264. IEEE, 2017. [2](#)