Robust Angle Invariant GAS meter reading

Ignazio Gallo, Alessandro Zamberletti and Lucia Noce
University of Insubria,
Department of Theoretical and Applied Science,
Via Mazzini, 5, 21100 Varese, Italy
Email: {ignazio.gallo, a.zamberletti, lucia.noce}@uninsubria.it

***Abstract*—In this work we propose a novel method for automatic gas meter reading from real world images. In a wide range of countries all over the world, the existing automatic technology is not adopted, usually the reading is manually done on site, and a picture is taken through a mobile device as a proof of reading. In order to confirm the reading, a tedious work of checking the proof images is commonly done offline by an operator. With this contribution we aim to supply an effective system, able to provide a real support to the validation process reducing the human effort and the time consumed. We exploit both region-based and Maximally Stable Extremal Regions techniques, during the phase involving the localization of the meter area and to detect the meter counter digits in the detection step respectively. The evaluation has been carried out on every step of our approach, as well as on the overall assessment; although the problem is complex, the proposed method leads to good results even when applied to degraded images, it represents an effective solution to the gas meter reading problem and it can be utilized in real applications.**

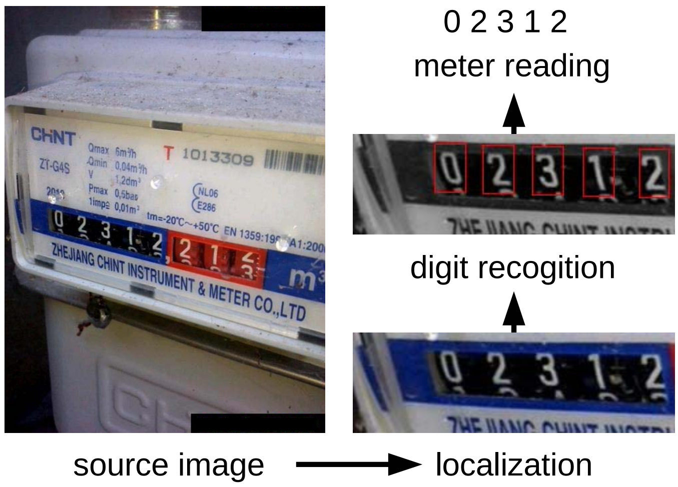

Fig. 1. An overview of the steps performed by the proposed model in order to read the gas meter counter that appear in the source image. The localization phase rotates and cuts the region of interest; the recognition phase identifies all potential digits of interest, while the reading step produces as output the final reading.

I. Introduction

Localizing and detecting text in natural images is a challenging problem that finds interests and applications in multiple areas such as, assisted driving [1], text based image retrieval [2], and multi-class object recognition [3]. Also a wide range of tasks belonging to commercial domains may be solved exploiting text localization and recognition from scene images; for example a lot of repetitive and monotonous data entry works, that are manually performed, can be replaced by automated systems that exploit such techniques. With this work we analyze the specific and concrete task of gas flow meter reading from real world images, proposing a novel method able to automatically localize the meter area within an image and perform the reading.

The motivation behind this work is that nowadays in a wide range of countries all over the world, the existing automatic meter reading technology is not widespread [4], for this reason the gas meter reading is usually done on site by an operator and a picture is taken from a mobile device as reading proof; some examples of these photos are shown in Fig. 2. Since this operation is prone to errors, to confirm the reading, the same proof image is manually checked offline by another operator. The validation phase is very expensive because of its time consuming, the method that we propose could represent a real solution to this particular problem, allowing a considerable saving of money and time.

Text localization has been the object of many recent studies and a wide range of methods are reported in the literature [5]–[11]. These techniques can be classified either as region-based or connected component CC-based.

Region-based [5]–[8] approaches view text as a special trait of the image that is distinguishable from the background; typically local features and sliding window classifiers are exploited to identify the existence of text regions within the processed image.

Different from the region-based, CC-based [9]–[11] methods handle the text components as elements with uniform characteristics, and extract connect components usually using Maximally Stable Extremal Regions (MSER) [12].

The proposed method exploits both the approaches. More specifically the meter reader algorithm can be divided into three different phases: the localization of the meter area, the detection of the meter digit and the reading of the meter counter. We use a region-based method during the localization step, focusing on the characteristic features that most of the meter counters images have; and exploit MSER to detect the meter digit during the detection phase, all the details are following reported.

Our method has been evaluated both in the overall assessment and in every step, on a huge dataset containing a wide range of real images. The images differ from rotation, gas meter size, color, blurring, illumination, meter type, etc.; the differences between the images make the task not trivial to solve.

Although the complexity of the problem, the proposed method leads to good results even when applied to degraded images, it represents an effective solution to the gas meter reading problem and it can be utilized in real applications.

II. Related Works

Even though the literature is full of works on text detection in real world images, the task of automatically recognize consumption digits of meters such as gas meters and electricity meters is not widespread. However a satisfactory number of works have been initiated in the last few years.

Cai *et al.* [13] studied the problem of automatic electric meter readings. They defined the region of interest (ROI) as the area within the image occupied by the meter, and detected it exploiting colors features. A post processing step, that takes into account the format of the display of the consumption digits, allows to obtain a finer detection. At the last phase, meter readings are recognized using a neural network. The presented results show a recognition rate greater than 90%, however this results are based on very specific meters. The dataset utilized is not provided, so it is impossible to make a comparison.

A recent work proposed by Vanetti *et al.* [14] tackle the problem of automatic gas meter reading from real world images, proposing a method based on an ensemble of neural networks. In the final classification of the digits phase, an SVM with radial based function kernel is used. The final accuracy of their proposed system, reached an high overall accuracy value, equal to 87%. Their system was evaluated on a very small dataset composed only by 100 images. In Section IV we provide a comparison on the *meter-integration* dataset provided by authors, achieving competitive results.

Another recent work is from Grafmller and Bayerer [15]; authors discussed about the improvement of the performance of character recognition algorithms for industrial applications. They show how to improve classification accuracy and speed using prior knowledge, e.g., number of lines or characters, in character segmentation; and compared different combinations of features and classifiers to show differences in accuracy and processing time.

In 2014, Chouiten and Schaeffer [16], proposed a mobile application able to read and send the current gas meter consumption, acquired through a smart-phone camera. Authors proposed a semi-automatic algorithm; using the mobile application, the user localizes the meter region taking a picture of it, then the system provide the detection of the digit and the reading. In the last step the user is allowed to manually modify the reading in order to correct it through the graphical interface. The results achieved by their method are user dependent, for this reason we did not compare our work with this one.

Henning *et al.* [17] in their recent work proposed an approach for automatic meter reading that can be applied

Fig. 2. Some representative images extracted from the dataset of natural images used in this article. The dataset contains images rotated in all directions.

to water, gas, and electric meters. The detection phase is implemented exploiting the vertical edges in the image, suggesting that the approach is not able to deal with rotated images, and multilayer perceptron is employed during the digit classification step. Although they reached an high evaluation score, they tested a prototype implementation of it on a small dataset composed of only 180 images.

Analyzing recent works that approach the task of automatic meter reading, we notice that most of the systems are tested on a reduced dataset, as a consequence images could be very similar and could present lots of common features, leading the system to reach excellent results on that specific subset of examples. Thinking about a real applications, these small datasets could be not enough representative, and may do not reflect a real evaluation of the system itself. Unlike them, we tested our method on a large dataset composed of more then 160000 images, achieving very reliable results. Moreover, due to the heterogeneity of the set of images adopted in our work, the system must be able to manage several aspects such as: different rotations or sizes, blurring, illumination and various meter types.

III. Proposed Model

The proposed method consists of three major phases. The first phase is dedicated to localize the area of interest: it begins with the source image I in input, and produces in output a crop I_c of the area of interest, rotated in such a way that the consumption digits of the gas meter are arranged horizontally from left to right. During the second phase, the algorithm extracts all the regions r_i from I_c and classifies them as potential digits of the gas meter counter. The third step determines which are the valid regions r_i, of the gas meter counter, that compose the final reading. These main steps are summarized in Fig. 1.

A detailed description of all the phases of the proposed method are given in the following paragraphs.

A. Localization

The main characteristic feature that most of the gas meter counters have in common is the red color of the decimal digits. In this work we exploit this feature to locate the area of interest within the image (see an example in Fig. 3-(e)). A very simple solution to this problem is to segment the color image using a range of valid reds colors. Unfortunately, as demonstrated in our experiments, due to the high number of artifacts in the real images this way introduces many false positives that make difficult to locate the area of interest. A valid alternative is to train a very simple neural model in order to segment the image. In this work, to locate the area of interest in which seeking the digits, a MultiLayer Perceptron (MLP) [18] was trained; the model classifies each pixel of I as belonging to one of the decimal red digits or not.

The meter counter localization algorithm starts with the creation of a soft map of "red areas", doing regression with a 3×3 sliding window over all the image I. The MLP model reads the content of a window directly from an RGB image and writes the single output value in an equivalent window of the output image. The output map is converted into a binary image I_B using a threshold t. Starting from I_B, the next step is to unify, potential red areas that are close and that normally represent the decimal digits (see an example in Fig. 3-(b)), in one area. In this way it's possible to identify the decimal separator and the rotation angle. At each iteration the morphological dilate operation is performed to try to obtain a single region. The morphological dilate operation uses a 5×5 structuring element for probing and expanding shapes contained in the input image. The iterative algorithm stops as soon as the number of objects in the map becomes equal to 1 or it terminates when reached the maximum number of iterations. At each iteration the binary mask is blurred with a 7×7 gaussian filter, followed by a dilate operation. In many cases the localization algorithm finds more than one region, and then we must select the most likely. The final step of this localization algorithm is to identify the most likely region in order to obtain the rotation angle to rotate the source image, and the area of interest to crop the source image. The selection strategy used in this work consists in applying the *Digit Recognition* phase, described in the following section, and to select the region that contains the largest number of regions classified as digit of interest.

B. Digit Recognition

The first step of this phase transforms the cropped image I_c into a gray levels image, then MSER [12] are computed. MSER algorithm is a salient region detector that we used to find regions r_i that surround a gas meter digit. Each identified region r_i must be classified by an appropriate

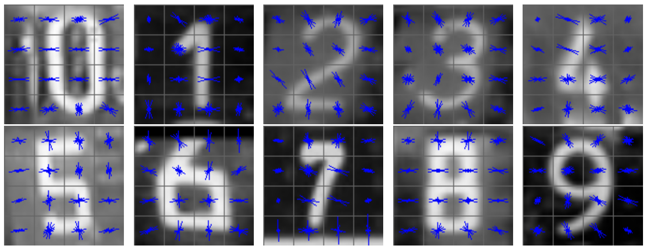

Fig. 4. HOG feature example, one for each class of interest. In this example, each patch 32×32 is divided into cells of size 8×8 and for each cell the gradient strengths are then visualized by line lengths in the direction of the corresponding gradient bin.

Support Vector Machine (SVM), trained to recognize and classify each region as belonging to one of the following classes $c_i \in \{0, 1, 2, 3, 4, 5, 6, 7, 8, 9, n\}$ where n means *not a digit*. Before passing a MSER region to the classifier, the same region is dilated by $1/6$ and then it's squarified in order to turn a rectangle into a square. Each square extracted is scaled to the $N \times N$ size, on it HOG [19] feature is calculated; it represents the input of the SVM classifier. In Fig. 4 the HOG feature is graphically represented for some representative MSER patches.

This algorithm phase produces a set of regions r_i, each region is classified as containing a digit or not. The output of this step consists of a set of digits d_i extracted from the regions r_i that contain them.

C. Meter Reading

Normally on every single digit of interest, we obtain lots of MSER regions (see Fig. 5-(a)). To be able to correctly read a digit we must group the MSER regions and find the representative label of each cluster. The basic steps of the meter reading algorithm are shown, from top to bottom, in Fig. 5, and are described in the following points:

(a) The starting point is represented by all the bounding boxes of the digit d_i associated with the MSER region r_i found as described in Section III-B.
(b) From these bounding boxes every potential d_i found is first deflated of τ pixels to better separate the overlapping clusters of MSER. After that, every d_i is drawn on a mask M.
(c) From M we seek all the contours, using the algorithm proposed in [20], in order to identify clusters on which estimating the dominant digit $\overline{d_j}$ that will become the label associating with the cluster centroid. In this step the digits $\overline{d_j}$ found are sorted according to their position in the cropped image.
(d) The searching of clusters often introduces false positive values that must be eliminated to avoid compromising the final reading. The approach used to eliminate false positives is a method that uses a M-estimator algorithm [21] to try to pass a line L on a set of 2D points. The points in this case are the first n centers of the bounding box of each $\overline{d_j}$, sorted in descending order according to their

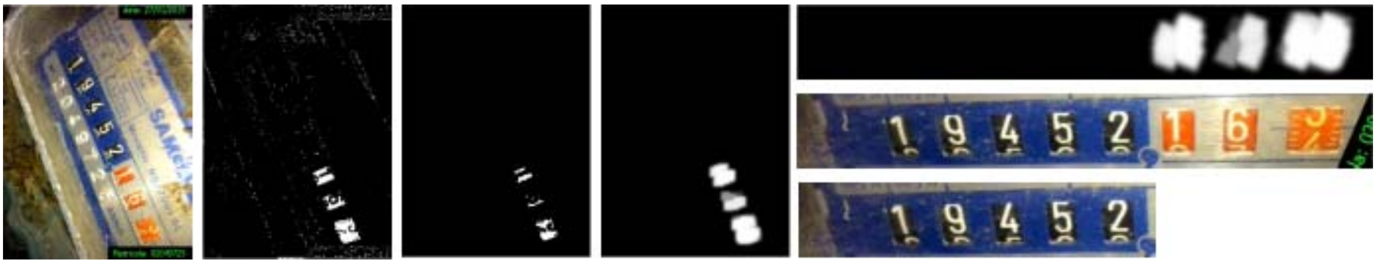

(a) (b) (c) (d) (e)

Fig. 3. Gas meter counter localization example. From left to right: source image (a), soft mask of red areas generated by the MLP (b), first (c) and last (d) image of the iterative inflate of the white regions, rotation and crop of the area of interest (e).

size. The algorithm fits a line to a 2D point set by minimizing $\sum_i \rho(q_i)$ where q_i is a distance between the i^{th} point and the line L. The distance function $\rho(q_i)$ used here is the following:

$$\rho(q) = 2 \cdot (\sqrt{1 + \frac{q^2}{2}} - 1) \qquad (1)$$

This part of the algorithm selects the subset D_f of bounding boxes $\overline{d_j}$ that have non-zero intersection with the line L.

(e) Considering that the digits of the meter counter are located at a fixed distance, to avoid including false positives or losing false negatives, we estimate the size S_i of a generic digit $\overline{d_j}$, and the distance σ_{ij} between digits belonging to the set D_f. If the estimated bounding box has non-empty intersection with the bounding box belonging to D_f then we use the same label $\overline{d_j}$, otherwise the bounding box is classified again in $\hat{d_j}$ using the same SVM described in Section III-B. The number of digits to be searched can be fixed in advance or can be estimated in the step (d).

(f) The final reading of the gas meter counter is obtained automatically from the set of bounding box $\hat{d_j}$ and $\overline{d_j}$ estimated in the previous step.

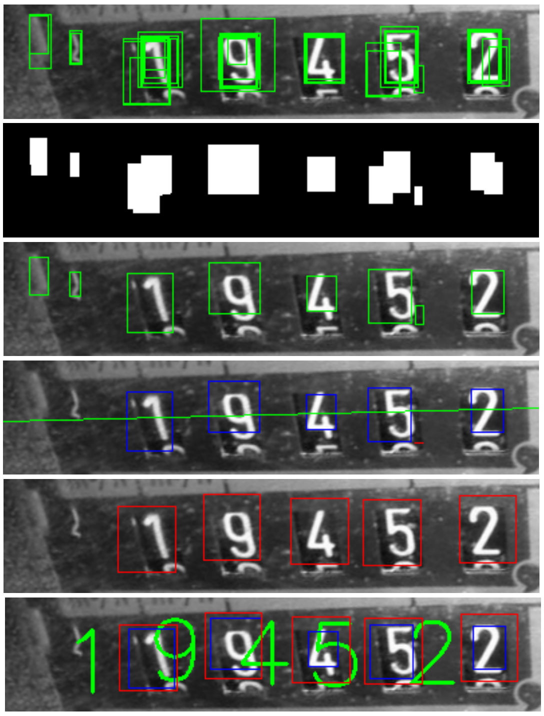

(a)

(b)

(c)

(d)

(e)

(f)

Fig. 5. Meter reading example. From top to bottom: MSER bounding box (a), the activation map (b), the bounding box of the clusters found (c) fit line to eliminate false positives (d), digit hypotheses to eliminate false negatives (e), final reading (f).

IV. Experiments

Experiments described in this section are intended to assess each phase of the proposed algorithm. In particular in this work we conducted three main experiments each one designed to evaluate all the steps described in Section III. For these experiments we used a large dataset D_R having $168,958$ natural images of size 640×480, captured by a mobile device with which the performance of the system was evaluated. For some of the following experiments we selected some subsets of random images from D_R. In Fig. 2 some representative images, of the dataset used, are shown. All the experiments were performed on an Intel Xeon E5-1620 at 3.5 GHz using the C++ library OpenCv 2.4.4.

A. Localization

To evaluate the first step of the proposed method we built and used a dataset containing 7161 points belonging to the classes "*red area*" and "*background*", equally distributed and extracted from 100 different images of the D_R dataset. The training set contains 4726 points extracted from these images while the test set contains 2435 points. The MLP is configured with an input layer having 27 neurons, used to read the contents of a 3×3 window from the RGB input image, and a single neuron for the output layer which represents the degree of membership to the "*red area*" class. The network has only one hidden layer with 27 neurons. To train the neural model we used the Resilient BackPropagation [18] algorithm with the following parameters: $\Delta_0 = 0.1$, $\eta^+ = 1.2$, $\eta^- = 0.5$, $\Delta_{min} = 0$, $\Delta_{max} = 50$. The training was stopped after 600 epochs or as soon as the error become less than 0.001;

The testing accuracy of the trained model is 95.1% on single pixels classification. The test accuracy is high despite the classification of the red areas is an ambiguous problem since the images contain many false positives examples.

To understand the goodness of the obtained result we compared our localization algorithm with the results obtained by a simple segmentation based on a range of red colors. We check if RGB values of an image point $I(x, y)$ lie between the elements of the lower red color CL and the upper red color CU checking the following three conditions:

$$\begin{aligned} CL_R \leq I_R \leq CU_R \\ CL_G \leq I_G \leq CU_G \\ CL_B \leq I_B \leq CU_B \end{aligned} \tag{2}$$

The lower and upper red color boundary was computed using all the positive points of the training set. In particular, $CL = [131, 32, 11]$ and $CU = [195, 82, 77]$. Applying this simple segmentation to all the points of the test set we got to an accuracy equal to 52.5%

This result is good for us because classifying all the pixels of an image with this trained MLP we can achieve a very high localization accuracy, comparing the accuracy achieved with the one obtained with a simple segmentation. Indeed, using a dataset consisting of 6000 test images randomly extracted from D_R, we measured 95% of accuracy in localization. This means that with a map created with the MLP described above, in 95% of cases, we are able to identify, rotate and crop properly the area of interest where we will try to read the meter counter, using the following "*Digit Recognition*" and "*Meter Reading*" phases.

B. Digit Recognition

This experiment aims to evaluate the best configuration of the HOG feature introduced in Section III-B, in order to maximize the test accuracy of the SVM model. For this purpose we changed two of the most important parameters of the HOG feature as described in Table I. These parameters are the *window size* and the *cell size* described in [19].

The dataset used for this phase contains clippings digits extracted from the images and in Fig. 6 some examples for each class are shown. The training patterns, used to train the SVM model, were extracted in order to leave the dataset balanced. The entire dataset was divided into 49836 training patterns and 24917 test patterns.

Guided by the results we set the *window size* to 32×32 and the *cell size* to 8×8 (see some examples in Fig. 4), leaving all other parameters with their default value. This means that in our proposed model we reported the size $N \times N$ of each r_i to 32×32 and then on it we calculate the feature HOG with the following main parameters: window size 32×32, cell size 8×8, block size 16×16, block stride 8×8 and number of bins 9.

The SVM model was configured with a kernel RBF. We automatically choose the optimal parameters C, γ, considering the minimal cross-validation estimate of test error. The model takes as input a vector with 324 members represented the feature HOG.

TABLE I

COMPARISON ANALYSIS OF THE MAIN PARAMETERS OF THE FEATURE HOG. EACH ROW SHOWS THE TEST RESULT OBTAINED WITH THE SVM TRAINED WITH A PARTICULAR CONFIGURATION OF THE HOG FEATURE.

HOG params		SVM accuracy	
cellSize	winSize	train	test
8x8	16x16	86.84	86.10%
8x8	32x32	99.97%	**99.05%**
8x8	48x48	99.97%	98.67%
8x8	64x64	99.97%	98.76%
4x4	16x16	99.97%	98.57%
4x4	32x32	99.97%	98.69%
4x4	48x48	99.97%	98.69%
4x4	64x64	99.97%	98.76%

Fig. 6. Training examples of the digit recognition step. From each scaled patch the HOG descriptor is extracted, which become an input pattern to the SVM model. The *not a digit* class is not included and contains regions that do not contain meter's digit.

C. Meter Reading

The main objective of this last experiment is to understand what are the optimal values of the main parameters that affect some of the points (a)-(f) summarized in Fig. 5. To analyze these parameters we used a small subset $\bar{D}_R \subset D_R$, consisting of 1000 images randomly selected.

The first parameter δ that we want to analyze is responsible for the number of regions extracted by the MSER algorithm. The parameter δ indicates through how many different gray levels does a region need to be stable to be considered maximally stable. For a larger δ , we will get less regions. For this reason we have assigned to the parameter δ values ranging from 0 up to 9 and for each value we evaluated the reading accuracy and the average execution time. Fig. 7 summarizes the results obtained. From Fig. 7 we know that when we increase the δ value, the execution time decreases just because the system must classify a smaller number of regions. The δ value that leads to maximum accuracy in this case corresponds to $\delta = 6$.

The second parameter that we experimentally estimate is τ, the deflating size in pixels of the bounding boxes r_i, from which we identify the centroids in the mask M, as shown in Fig. 5-(b). We varied the number τ of pixels eroded from each bounding box, and for each value of τ we evaluated the

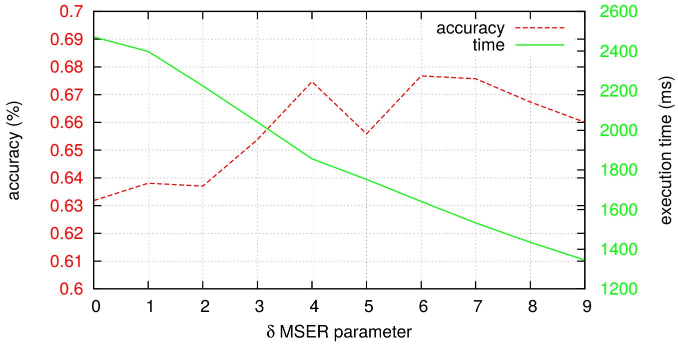

Fig. 7. Analysis of the δ parameter of the MSER. The graph shows the resulting accuracy of meter readings and the average execution time for a single input image when varying the δ parameter.

Fig. 9. Some problematic images of the dataset used. In some images the red decimal digits are missing and then the localization fails, while in the others due to poor lighting conditions, occlusions and perspective distortions, the MSER algorithm fails to find the regions containing the digits of interest.

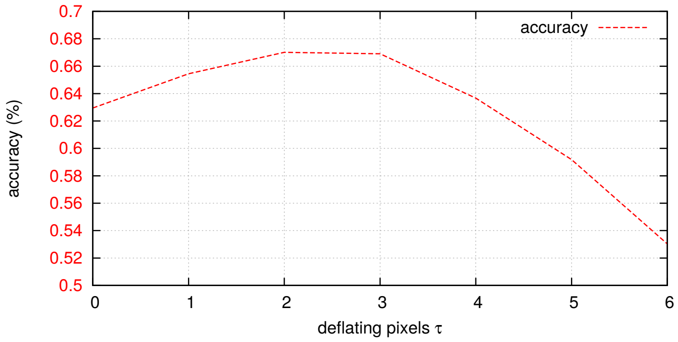

Fig. 8. Analysis of the τ parameter, representing the deflating size in pixels of a MSER bounding box r_i, from which the centroids in the mask M are identified.

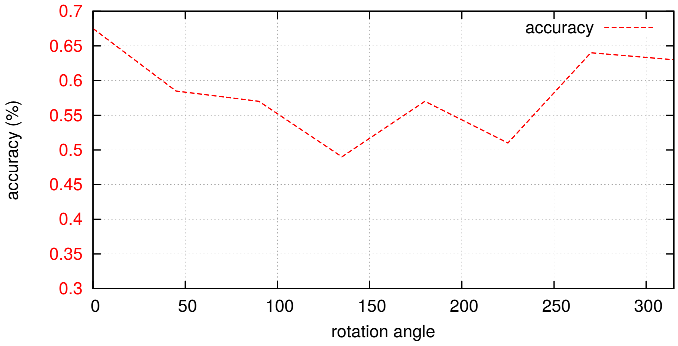

Fig. 10. Analysis of the meter reading accuracy while varying the rotation angle (in degree) of the gas meter counter.

overall accuracy using the $\bar{D}_R$ dataset. As shown in the graph of Fig. 8, the best result obtained is with $\tau = 2$.

D. Overall Evaluation and Comparison

The final meter reading accuracy reached with the full D_R dataset, excluding the images used for training and using the best parameter configuration, is equal to 67.5%.

To better understand where our system fails, we analyzed the discarded images. In Fig. 9 some representative examples of discarded images are shown. Some of these images have problems of blur, lighting, perspective distortion, etc., due to these issues the MSER algorithm does not correctly identify the regions of interest. Other images show how, due to the lack of red decimal numbers, the localization algorithm fails and consequently it becomes difficult to find the digits of interest.

To analyze the robustness of the proposed method when varying the angle of rotation of the meter counter, we made a further experiment. We used the dataset D_R selecting all the images in which the meter counter was horizontal (angle equal to zero degrees). The set of images obtained was manually rotated incrementally at intervals of 45 degrees, measuring the reading accuracy at every rotation angle. Fig. 10 shows the obtained accuracies curve. This result obtained, even if a small loss in accuracy was measured for angles different to zero, shows how the proposed method is robust while varying the reading angle.

To obtain a goodness measure of our system, regardless of the dataset used for training, we compared the proposed method with another method [14] available in literature using their dataset *meter-integration* available online. On this dataset we got the 85% of accuracy against the 87% stated by the authors. Our algorithm was trained on completely different images and then the result shows the validity of what we propose in this work.

V. Conclusion

The localization and recognition of digits in a gas meter counter is a difficult task due to the quality of images taken by mobile devices. The proposed model, based on MSER and neural network regression and classification, is able to detect the area of interest and to recognize every single digit in an average of 1.6 seconds per image. Starting from the consideration that about 99% of the images has red colored decimal digits, the localization approach that is anchored to this information becomes robust and reliable. Finally, the HOG and SVM used to classify each single digit has been shown to be an excellent and accurate solution. For these reasons the proposed method can be applied in a real application context for the text reading in natural images, our solution can be applied in support to humans, to validate the reading that currently operators do manually.

We conducted an experimental phase in a real scenario where the solution is applied to a large database of images reaching a very interesting accuracy. As future work, we are working to extend the proposed method to locate the meter counter even in the absence of red digits.

References

[1] L. Neumann and J. Matas, "Real-time scene text localization and recognition," in *Proceedings of CVPR*, 2012.

[2] N. Goel and P. Sehgal, "A refined hybrid image retrieval system using text and color," *International Journal of Computer Science Issues*, vol. 9, no. 4, pp. 48–56, 2012.

[3] S. Karaoglu, J. C. van Gemert, and T. Gevers, "Object reading: Text recognition for object recognition," in *Proceedings of ECCV*, 2012.

[4] G. Deconinck, "Metering, intelligent enough for smart grids?" *Securing Electricity Supply in the Cyber Age*, vol. 15, pp. 143–157, 2010.

[5] R. Minetto, N. Thome, M. Cordb, N. J. Leitec, and J. Stolc, "T-HOG: An effective gradient-based descriptor for single line text regions," *Pattern Recognition*, vol. 46, pp. 1078–1090, 2013.

[6] C. V. J. Anand Mishra, Karteek Alahari, "Top-down and bottom-up cues for scene text recognition," in *Proceedings of CVPR*, 2012.

[7] K. Wang, B. Babenko, and S. Belongie, "End-to-end scene text recognition," in *Proceedings of ICCV*, 2011.

[8] A. Mishra, K. Alahari, and C. Jawahar, "Scene text recognition using higher order language priors," in *Proceedings of BVMC*, 2012.

[9] H. I. Koo and D. H. Kim, "Scene text detection via connected component clustering and non-text filtering," *IEEE Transactions on Image Processing*, vol. 22, no. 6, pp. 2296–2305, 2013.

[10] C. Shi, C. Wang, B. Xiao, Y. Zhang, and S. Gao, "Scene text detection using graph model built upon maximally stable extremal regions." *Pattern Recognition*, vol. 34, no. 2, pp. 107–116, 2013.

[11] L. Neumann and J. Matas, "A method for text localization and recognition in real-world images," in *Proceedings of ACCV*, 2011.

[12] J. Matas, O. Chum, M. Urban, and T. Pajdla, "Robust wide baseline stereo from maximally stable extremal regions," in *Proceedings of the British Machine Vision Conference*, 2002.

[13] Z. Cai, C. Wei, and Y. Yuan, "An efficient method for electric meter readings automatic location and recognition," *Procedia Engineering*, vol. 23, no. 0, pp. 565 – 571, 2011.

[14] 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.

[15] M. Grafmller and J. Beyerer, "Performance improvement of character recognition in industrial applications using prior knowledge for more reliable segmentation," *Expert Systems with Applications*, vol. 40, no. 17, pp. 6955 – 6963, 2013.

[16] M. Chouiten and P. Schaeffer, "Vision based mobile gas-meter reading machine learning method and application on real cases," in *Proceedings of Scientific Cooperations International Workshops on Electrical and Computer Engineering Subfields*, 2014.

[17] H. Puttnies, V. Altmann, F. Golatowski, and D. Timmermann, "Cost-efficient universal approach for remote meter reading using web services and computer vision," in *Proceedings of IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks*, 2015.

[18] M. Riedmiller and H. Braun, "A direct adaptive method for faster backpropagation learning: the rprop algorithm," in *Proceedings of IEEE International Conference on Neural Networks*, 1993.

[19] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *Proceedings of International Conference on Computer Vision & Pattern Recognition*, 2005.

[20] S. Suzuki and K. Abe, "Topological structural analysis of digitized binary images by border following," *Computer Vision, Graphics, and Image Processing*, vol. 30, no. 1, pp. 32–46, 1985.

[21] P. Mukhopadhyay, *An Introduction to Estimating Functions*. Alpha Science International, 2004.