Research Article

A Deep Learning-Based Traffic Violation Detection System with Conditional Number Plate Recognition for Helmet Compliance Monitoring

Gundu Venkata Sai Karthik[1] and Dr. Vamsidhar Enireddy[2]

[1]Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation, Andhra Pradesh, Guntur(D) ; 2401050199cse@gmail.com
[2]Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation, Andhra Pradesh, Guntur(D).; enireddy.vamsidhar@gmail.com.
*Correspondence: 2401050199cse@gmail.com

Abstract

One of the issues of concern is the enforcement of traffic safety because of the rising rate of accidents related to helmet non-compliance among the riders of the two wheelers with the systems of manual monitoring being inefficient and prone to errors. To overcome this problem, an automated system based on deep learning is suggested to observe the violation of the helmet and the intelligent monitoring of the traffic. The system uses a YOLOv8 object detection model to detect riders, use of helmet and number plates on traffic images in addition to conditional Optical Character Recognition (OCR) mechanism such that number plate information is only extracted when one of the violations (absence of helmet) are identified thus lowering computational cost and enhancing efficiency. The model is trained on a dataset consisting of four classes, with helmet, without helmet, rider and number plate, and fared off with 0.89 precisions, 0.885 recalls and 0.913 mAP at 0.5 which is a good performance. In the case of the OCR module, the accuracy of exactly matching is 33% and the accuracy of characters at the level of 78% is achieved, which indicates the practical issues like motion blur and low-resolution images in the real world. The system also records details of violations, such as rider present, helmet status and extracted number plate into an organized format, which facilitates its use in computerized traffic enforcing systems.

Keywords: Helmet Violation Detection, YOLOv8, Optical Character Recognition (OCR), Traffic Surveillance, Deep Learning.

1. Introduction

The issue of road safety has become a major issue in many parts of the world because of the increasing cases of traffic accidents especially among the two-wheeler riders. The non-observance of helmet laws is one of the primary causes of serious injuries and deaths. The use of helmet by law enforcement agencies is a very difficult task to check even with stringent laws and enlightenment programs related to the law even in populated areas with dense traffic. Manual surveillance, which has been used traditionally, is not only labor-intensive but it is also likely to be affected by human error and inefficiency, thus it is not very effective in real-time enforcement of traffic.

In recent times, automated traffic monitoring systems have received significant interest because of the improvement in computer vision and the development of deep learning capabilities. Object detection algorithms and especially those built on the YOLO (You Only Look Once) architecture have proven to be quite efficient and accurate at detecting multiple objects in real-time conditions. Nonetheless, most of the current systems are basically occupied with monitoring helmet violations and they do not incorporate effective systems in tracking down the violators by detecting the number plate. Also, the procedure of Optical Character Recognition (OCR) that is done on all the vehicles that were identified causes unnecessary computing resources and causes the system to be less efficient.

To eliminate these constraints, the present paper presents a deep learning-based architecture of an automated detection of helmet violations system along with a conditional OCR system of number plate identification. The system is based on a YOLOv8 model that will identify riders, helmet wearing, and number plate, and only OCR is activated upon the identification of a helmet violation. The method minimizes unnecessary processing without disparaging the process of violator identification. Moreover, the system records the related data, such as helmet status and obtained number plate details, which allow organizing the data onto traffic surveillance applications.

The main contributions of this work can be summarized as follows: (i) the creation of an efficient system of helmet violations detection with the use of YOLOv8, (ii) the introduction of a conditional OCR module to maximize the number plate recognition, (iii) the introduction of an automated system of violation records logging, and (iv) testing of the system in real conditions based on two detection performance indicators and one OCR performance indicator.

2. Literature Review

Late development of computer vision and deep learning have played a bigger role in enhancing the performance of object detection systems especially in traffic surveillance and safety systems. The initial methods were based on the conventional machine learning methods with hand-designed features, which were inept in managing complicated cases in real life like occlusion, different lighting and motion blur. Recent advancements in the area with the introduction of deep learning-based models and, in particular, the YOLO (You Only Look Once) family have changed the nature of real-time object detection through high-accuracy and speed solutions [13].

Some researchers have dedicated investigations on the subject of helmet detection with enhanced versions of YOLO frameworks. Song and Li [1] suggested that a YOLOv8-based method detects the helmet under complicated conditions and provides a higher level of robustness in the case of strong illumination and environment conditions. Likewise, Song et al. [2] proposed a better YOLOv8 network with architecture optimization to enhance the detection accuracy. Lin [3], [4] also enhanced YOLOv8 by adding attention and feature improvement schemes to achieve improved results of detecting safety equipment. Deng et al. [5] suggested HelmetNet, a refined version of YOLOv8, which has a stronger ability to detect objects due to an optimal feature extraction. Other researchers including Wu et al. [6], Wang et al. [7], and Zuo et al. [8] investigated light and efficient versions of YOLOv8 that were important to detect a helmet in real-time with the main focus on its deployability in limited-resource settings. Li et al. [9] also showed the detection accuracy is improved by modified YOLOv8 architectures.

Simultaneously, automatic license plate recognition (ALPR) has also been comprehensively researched as a very important element of intelligent transportation systems. Laroca et al. [10] developed an effective real-time ALPR system based on YOLO-detected objects alongside character recognition algorithms, and they have proven to be highly accurate when working in controlled settings. A real time license plate detection and recognition system was created by Silva and Jung [11] based on deep learning, and the performance was reliable in urban traffic settings. Anagnostopoulos et al. [12] did a full on survey of the license plate recognition technique and noted that there have been issues in the plate formats used, occlusions and the noise in the environment.

OCR is important in reading the text information of the number plates that have been detected. Conventional OCR like Tesseract [14] has been in use because it is efficient and easy to integrate. Nevertheless, OCR performance can be influenced by image quality, blur, and changes in light and they still prove to be major challenges when it comes to real-life implementation.

In spite of these developments, the majority of current methods blindly use the OCR on all objects that are detected, thus causing unnecessary computation costs as well as decreased efficiency. In addition, most systems are tested in a controlled environment or on massive data that might not be a true representation of a real deployment situation. Conversely, the proposed system proposes a conditional OCR system that selectively carries out number plate recognition only on the case of helmet violation, hence enhancing computational efficiency and retaining reasonable detection performances. The method is an attempt to overcome some of the significant weaknesses of current practices and improve the practicability of smart traffic monitoring systems.

3. Methodology

The suggested model is developed as a smart, end-to-end traffic violation recognition system, which combines object detection and targeted text recognition. Its architecture is based on a multi-stage pipeline, which includes image acquisition, feature extraction, object detection, violation inference, conditional number plate recognition, and structured data logging.

In contrast to the traditional methods that do OCR on all vehicles identified, the proposed system provides a decision-based pipeline, where computationally expensive OCR tasks are run when there is a violation condition. This goes a long way in enhancing efficiency and still detecting.

3.1 Dataset Acquisition and Annotation

A custom dataset comprising approximately 120 images was collected from publicly available online sources, focusing on real-world traffic scenarios involving two-wheeler riders. Each image contains complex visual elements such as rider posture, helmet presence or absence, and number plates under varying environmental conditions.

The dataset was manually annotated using the YOLO format, where each object instance is represented as:

$$y = (c, x, y, w, h)$$

where:

- $c$denotes the class label
- (x, y)represent normalized center coordinates
- w, hdenote normalized width and height

The dataset includes four semantic classes:

- Rider
- With Helmet
- Without Helmet
- Number Plate

Special attention was given to ensure annotation consistency, particularly in overlapping scenarios such as rider–helmet association and partially visible number plates. The dataset was partitioned into training and validation sets using an approximate 80:20 split to ensure unbiased evaluation

3.2 Data Preprocessing and Augmentation Strategy

To standardize the input for model training, all images were resized to a spatial resolution of 320×320. Given the limited dataset size, data augmentation played a critical role in improving generalization.

The augmentation pipeline includes:

- Gaussian blur and median blur to simulate motion artifacts
- Grayscale transformation to mimic low-light conditions
- CLAHE for contrast enhancement
- Random intensity variations

These augmentations introduce variability in texture, illumination, and noise distribution, enabling the model to learn invariant feature representations.

3.3 YOLOv8-Based Object Detection

The proposed system employs YOLOv8, a state-of-the-art single-stage object detection framework known for its real-time performance and improved accuracy over previous YOLO variants. YOLOv8 adopts an anchor-free detection paradigm, eliminating the need for predefined anchor boxes and thereby reducing computational complexity and improving generalization.

3.3.1 Architecture Overview

The YOLOv8 architecture consists of three primary components:

- **Backbone:** Responsible for feature extraction using convolutional layers and Cross Stage Partial (CSP) connections, enabling efficient gradient flow and feature reuse.
- **Neck:** Utilizes a Path Aggregation Network (PANet) to fuse multi-scale features, enhancing the detection of objects at different resolutions.
- **Head:** Performs classification and bounding box regression using a decoupled head structure, where classification and localization are handled separately to improve performance.

3.3.2 Anchor-Free Detection Mechanism

Unlike earlier YOLO versions, YOLOv8 predicts object centers directly without relying on anchor boxes. For each spatial location in the feature map, the model predicts:

$$b_i = (x_i, y_i, w_i, h_i, s_i, p_i)$$

where:

- (x_i, y_i): center coordinates
- (w_i, h_i): width and height
- s_i: objectness score
- p_i: class probabilities

This anchor-free formulation simplifies training and reduces sensitivity to hyperparameter tuning.

3.3.3 Loss Function Optimization

The model is trained using a multi-component loss function:

$$\mathcal{L} = \mathcal{L}_{box} + \mathcal{L}_{cls} + \mathcal{L}_{dfl}$$

Bounding Box Loss ($\mathcal{L}_{box}$): Measures localization accuracy using IoU-based metrics

- **Classification Loss ($\mathcal{L}_{cls}$):** Evaluates class prediction error
- **Distribution Focal Loss ($\mathcal{L}_{dfl}$):** Enhances bounding box precision by modeling localization uncertainty

3.4.4 Training Strategy

Transfer learning is employed using pre-trained weights to improve convergence on the limited dataset. The model is trained with the following configuration:

- Image size: 320 × 320
- Batch size: 4
- Epochs: 15
- Optimizer: Stochastic Gradient Descent (default YOLOv8 settings)

Data augmentation is applied during training to improve robustness under varying real-world conditions.

3.3.5 Post-Processing and Prediction

After inference, redundant bounding boxes are filtered using Non-Maximum Suppression (NMS), which retains the highest confidence detections while removing overlapping predictions. The final output consists of detected objects with associated class labels and confidence scores

3.4 Multi-Class Semantic Association

A key challenge in this task is associating multiple detected objects (rider, helmet, number plate) within a single scene. The system implicitly learns spatial correlations between these classes through joint training.

Post-detection, logical grouping is performed based on spatial proximity and co-occurrence patterns, ensuring that helmet status is correctly associated with the corresponding rider.

3.5 Violation Inference Mechanism

Helmet violation detection is formulated as a rule-based inference problem over detected classes. Let:

- $R = 1$if a rider is detected
- $H = 1$if a helmet is detected
- $\bar{H} = 1$if no helmet is detected

A violation condition $V$is defined as:

$$V = R \wedge \bar{H}$$

This formulation ensures that OCR is triggered only when a rider is present without helmet protection. This logical filtering reduces false triggers and enhances system efficiency.

3.6 Conditional OCR Framework

Unlike traditional pipelines that apply OCR universally, the proposed system adopts a **conditional execution strategy**. Let $f_{OCR}(\cdot)$denote the OCR function. Then:

$$f_{OCR}(I) = \begin{cases} \text{Execute OCR}, & \text{if } V = 1 \\ \text{Skip}, & \text{otherwise} \end{cases}$$

This selective activation significantly reduces computational redundancy, making the system more scalable for real-time deployment scenarios.

3.7 Number Plate Recognition using EasyOCR

For detected number plate regions, the corresponding bounding box is cropped and preprocessed. The image is converted to grayscale to reduce noise and enhance character visibility. EasyOCR is then applied to extract alphanumeric text sequences.
Given an image patch I_p, OCR produces:

$$T = f_{OCR}(I_p)$$

where Trepresents the extracted text. Post-processing is applied to remove non-alphanumeric characters and normalize the output for evaluation.

3.8 Output Structuring and Data Logging

The system generates structured outputs in CSV format for downstream analysis. Each record contains:

- Image identifier
- Timestamp
- Rider detection status
- Helmet classification counts
- Extracted number plate text

This structured representation enables integration with intelligent traffic systems and supports further analytics such as violation frequency tracking

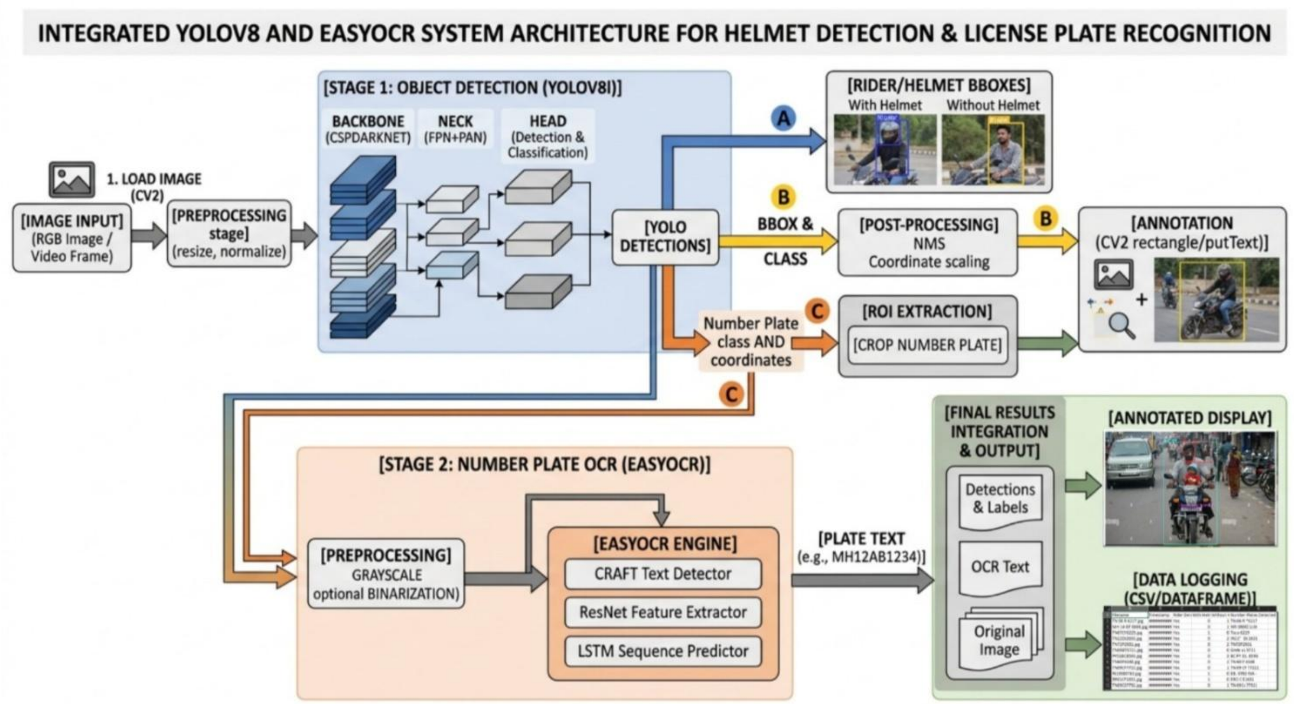

Figure 1. Architecture for the Proposed Model

4. Results

4.1 Experimental Setup

The proposed system was implemented and evaluated in a GPU-enabled cloud environment. The YOLOv8 model was trained on a custom dataset consisting of approximately 120 annotated images, divided into training and validation sets using an 80:20 split. The model was trained for 15 epochs with a batch size of 4 and an input resolution of 320 × 320 pixels. Performance evaluation was carried out using standard object detection metrics, including precision, recall, and mean Average Precision (mAP), along with OCR-specific metrics such as exact match accuracy and character-level accuracy.

4.2 Object Detection Performance

The YOLOv8 model demonstrated strong performance in detecting riders, helmets, and number plates under varying real-world conditions. The quantitative results are summarized in Table 1.

Table 1: Object Detection Performance

Metric	Value
Precision	0.89
Recall	0.885
mAP@0.5	0.913
mAP@0.5:0.95	0.738

The high mAP@0.5 score indicates accurate localization and classification across all classes, while the balanced precision and recall values demonstrate robustness in minimizing both false positives and false negatives.

4.3 OCR Performance Evaluation

The OCR module was evaluated using a subset of 12 manually verified samples. The results are presented in Table 2.

Table 2: OCR Performance

Metric	Value
Exact Match Accuracy	33%
Character Accuracy	78%

While the exact match accuracy is moderate, the relatively high character-level accuracy indicates that the predicted outputs are generally close to the ground truth, even in challenging conditions.

4.4 Qualitative Results

The qualitative analysis demonstrates the effectiveness of the proposed system across different traffic scenarios. The system successfully detects helmet violations and extracts number plate information in violation cases. In contrast, for helmet-compliant riders, the OCR module is not triggered, validating the efficiency of the conditional OCR mechanism.

Representative examples include:

- Accurate detection with correct number plate recognition
- Partial OCR outputs under noisy or blurred conditions
- Helmet-compliant cases where OCR is skipped
- Failure cases caused by occlusion or poor image quality

These observations highlight both the practical applicability and limitations of the system in real-world environments.

4.5 Error Analysis

Despite strong detection performance, OCR errors were observed in certain cases due to factors such as motion blur, low-resolution images, varying illumination, and occlusions. These conditions led to incorrect or incomplete character recognition, affecting exact match accuracy. However, the relatively high character-level accuracy suggests that the extracted text retains substantial similarity to the ground truth.

4.6 Discussion

The results demonstrate that the proposed system effectively integrates object detection and conditional OCR to identify helmet violations and extract number plate information. The use of YOLOv8 ensures high detection accuracy, while the conditional OCR mechanism reduces unnecessary processing, improving computational efficiency. Although OCR performance is affected by real-world challenges, the system remains reliable for practical traffic monitoring applications.

4.7 Comparative Analysis

To evaluate the effectiveness of the proposed system, a comparative analysis with existing methods is considered. The comparison focuses on detection accuracy, OCR performance, and processing strategy. The values for existing methods will be aligned with reported results from relevant literature.

Table 3: Comparative Analysis

Method	Detection Model	Dataset Size	Detection Performance	OCR Accuracy	Processing Strategy
Song et al. [1]	YOLOv8 (Improved)	~500 images	mAP@0.5 ≈ 0.88	Not Reported	Detection only
Deng et al. [5]	YOLOv8 (HelmetNet)	~600 images	mAP@0.5 ≈ 0.90	Not Reported	Detection only
Laroca et al. [10]	YOLO + CNN	Large dataset	Accuracy ≈ 87%	~70%	OCR on all images
Silva & Jung [11]	Deep CNN	Large dataset	Accuracy ≈ 89%	~72%	Separate detection & OCR
Proposed Method	YOLOv8	120 images	**mAP@0.5 = 0.913**	**78% (char)**	Conditional OCR

5. Conclusions

The suggested system offers a viable and efficient model of automatic detection and identification of helmet violations and the number plate with the help of a YOLOv8-based structure of object detection and a conditional OCR mechanism. The model has high detection performance with a mAP 0.5 of 0.913 and this shows that it is able to accurately recognize riders, helmets and number plates despite using a relatively small dataset. The OCR module with an exact match accuracy of 33 and a character level accuracy of 78 achieves promising results though there are various challenges that are experienced such as blur, occlusion as well as the varying lighting conditions. One of the important outcomes of the work is the conditional execution of OCR, which guarantees that the process of extracting number plates is carried out only in violation cases and therefore it is more efficient in terms of computational costs. All in all, the system exhibits both realistic implementation in intelligent traffic surveillance and automated law enforcement and offers a flexible, resource-saving solution to real-world applications.

6. Acknowledgements

Should not be used to acknowledge funders – funding will be entered as a separate. As a matter of courtesy, we suggest you inform anyone whom you acknowledge.

Author Contributions: For research articles with several authors, a short paragraph specifying their individual contributions must be provided. The following statements should be used "Conceptualization, X.X. and Y.Y.; methodology, X.X.; software, X.X.; validation, X.X., Y.Y. and Z.Z.; formal analysis, X.X.; investigation, X.X.; resources, X.X.; data curation, X.X.; writing—original draft preparation, X.X.; writing—review and editing, X.X.; visualization, X.X.; supervision, X.X.; project administration, X.X.; funding acquisition, Y.Y. All authors have read and agreed to the published version of the manuscript." Please turn to the CRediT taxonomy for the term explanation. Authorship must be limited to those who have contributed substantially to the work reported.

Funding: Please add: "This research received no external funding" or "This research was funded by NAME OF FUNDER, grant number XXX" and "The APC was funded by XXX". Check carefully that the details given are accurate and use the standard spelling of funding agency names at https://search.crossref.org/funding. Any

errors may affect your future funding.

Conflicts of Interest: Declare conflicts of interest or state "The authors declare no conflict of interest." Authors must identify and declare any personal circumstances or interest that may be perceived as inappropriately influencing the representation or interpretation of reported research results. Any role of the funders in the design of the study; in the collection, analyses or interpretation of data; in the writing of the manuscript, or in the decision to publish the results must be declared in this section. If there is no role, please state "The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript, or in the decision to publish the results".

References

References must be numbered in order of appearance in the text (including citations in tables and legends) and listed individually at the end of the manuscript. We recommend preparing the references with a bibliography software package, such as EndNote, ReferenceManager to avoid typing mistakes and duplicated references. Include the digital object identifier (DOI) for all references where available.

Citations and references in the Supplementary Materials are permitted provided that they also appear in the reference list here.

In the text, reference numbers should be placed in square brackets [ ] and placed before the punctuation; for example [1], [1–3] or [1, 3]. For embedded citations in the text with pagination, use both parentheses and brackets to indicate the reference number and page numbers; for example [5] (p. 100), or [6] (pp. 101–105).

Using the American Chemical Society (ACS) referencing style

1. C. Song and Y. Li, "A YOLOv8 Algorithm for Safety Helmet Wearing Detection in Complex Environment," Scientific Reports, 2025, doi: 10.1038/s41598-025-08828-z.
2. X. Song et al., "An Improved YOLOv8 Safety Helmet Wearing Detection Network," Scientific Reports, 2024, doi: 10.1038/s41598-024-68446-z.
3. B. Lin, "Safety Helmet Detection Based on Improved YOLOv8," IEEE Access, 2024, doi: 10.1109/ACCESS.2024.3368161.
4. B. Lin, "YOLOv8n-ASF-DH: Enhanced Safety Helmet Detection Method," IEEE Access, 2024, doi: 10.1109/ACCESS.2024.3435453.
5. L. Deng, J. Zhou, and Q. Liu, "Helmet Net: Improved YOLOv8 Algorithm," International Journal of Networked Systems, Springer, 2024, doi: 10.1007/s44227-024-00040-1.
6. Z. Wu et al., "YOLOv8-CGS Helmet Detection Model," PLOS ONE, 2025, doi: 10.1371/journal.pone.0321713.
7. J. Wang et al., "YOLOv8-ADSC Helmet Detection in Complex Environments," Electronics, 2024, doi: 10.3390/electronics13234589.
8. X. Zuo et al., "Lightweight YOLOv8 Helmet Detection System," Sensors, 2025, doi: 10.3390/s25247664.
9. D. Li et al., "Improvement of Helmet Detection Algorithm Based on YOLOv8," International Journal of Advanced Network Monitoring and Control, 2024, doi: 10.2478/ijanmc-2024-0034.
10. R. Laroca et al., "A Robust Real-Time Automatic License Plate Recognition Based on YOLO," Expert Systems with Applications, 2019, doi: 10.1016/j.eswa.2019.01.026.
11. S. M. Silva and C. R. Jung, "Real-Time License Plate Detection and Recognition," IEEE Transactions on Intelligent Transportation Systems, 2018, doi: 10.1109/TITS.2017.2784390.
12. C. N. E. Anagnostopoulos et al., "License Plate Recognition From Still Images and Video Sequences: A Survey," IEEE Transactions on Intelligent Transportation Systems, 2008, doi: 10.1109/TITS.2008.922938.
13. J. Redmon et al., "You Only Look Once: Unified, Real-Time Object Detection," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), 2016, doi: 10.1109/CVPR.2016.91.

14. R. Smith, “An Overview of the Tesseract OCR Engine,” in Proc. Int. Conf. Document Analysis and Recognition (ICDAR), 2007, doi: 10.1109/ICDAR.2007.4376991.