

InstructOCR: Instruction Boosting Scene Text Spotting

Chen Duan¹, Qianyi Jiang¹, Pei Fu^{1*}, Jiamin Chen², Shengxi Li¹, Zining Wang¹, Shan Guo¹, Junfeng Luo¹

¹Meituan

²Xi'an Jiaotong University

{duanchen02,jiangqianyi02,fupeilishengxi,wangzining03,guoshan05,luojunfeng}@meituan.com
{chenjm}@stu.xjtu.edu.cn

Abstract

In the field of scene text spotting, previous OCR methods primarily relied on image encoders and pre-trained text information, but they often overlooked the advantages of incorporating human language instructions. To address this gap, we propose InstructOCR, an innovative instruction-based scene text spotting model that leverages human language instructions to enhance the understanding of text within images. Our framework employs both text and image encoders during training and inference, along with instructions meticulously designed based on text attributes. This approach enables the model to interpret text more accurately and flexibly. Extensive experiments demonstrate the effectiveness of our model and we achieve state-of-the-art results on widely used benchmarks. Furthermore, the proposed framework can be seamlessly applied to scene text VQA tasks. By leveraging instruction strategies during pre-training, the performance on downstream VQA tasks can be significantly improved, with a 2.6% increase on the TextVQA dataset and a 2.1% increase on the ST-VQA dataset. These experimental results provide insights into the benefits of incorporating human language instructions for OCR-related tasks.

Code — <https://github.com/ChenD-VL/InstructOCR>

Introduction

Scene text spotting technology aims to detect and recognize characters directly within natural scene images. Recently, significant advancements in integrating vision and text have been made across various visual-language tasks, leading to the emergence of innovative instruction-based models. Some studies (Liu, Zeng et al. 2023; Geng et al. 2023; Brooks, Holynski, and Efros 2023; Kirillov, Mintun et al. 2023) have validated that incorporating human language instructions can enable models to comprehend the content of images more accurately. Drawing on the insights gained from these studies, we pose the following question: **For scene text images, which inherently involve visual text, wouldn't the incorporation of human language instructions be more beneficial?**

*Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



Figure 1: Examples of text recognition results generated using various instructions. This illustrates how different instructions can influence the output of the text recognition process.

In OCR tasks, several works, including oCLIP (Xue, Zhang et al. 2022) and ODM (Duan, Fu et al. 2024), have validated the effectiveness of pretraining the image encoder with relevant text information. However, these methods only employ the image encoder for text spotting tasks without utilizing the text encoder. TCM (Yu, Liu et al. 2023) and FastTCM (Yu, Liu et al. 2024) extract image and text-based prior knowledge through visual prompt learning and cross-attention within CLIP (Radford et al. 2021). Nonetheless, the text encoder used in their works is from CLIP and is frozen during their training process. STEP(Garcia-Bordils, Karatzas et al. 2024) attempts to control recognition through regular expressions. However, this method can only process numbers or letters.

We believe that aligning human language instructions with visual text will undoubtedly be advantageous for scene text spotting tasks. To substantiate this point, we propose a novel framework named InstructOCR, an instruction-based scene text spotting model that facilitates interaction through human language. Our objective is to leverage both text and image encoders during training and inference via instruc-

Condition	Templates
Recognition of all text in the image	<i>Recognize all text</i>
Recognition of text by specific length	<i>Text of more than $\{N\}$ characters</i> <i>Text of less than $\{N\}$ characters</i> <i>Text of $\{N\}$ characters</i>
Recognition of specific text	<i>Text consisting of $\{Word\}$</i>
Recognition of text containing specific characters	<i>Text starts with the character $\{C\}$</i> <i>Text ends with the character $\{C\}$</i> <i>Text contains the character $\{C\}$</i>
Recognition of numbers in the image	<i>Numerical text</i>
Recognition of alphabetic text in the image	<i>Alphabetical text</i>

Table 1: Different templates used in InstructOCR for scene text spotting allow N to be any integer, C to be any letter, and $Word$ to be either a single word or multiple words in the image, resulting in highly diverse instructions.

tions. To this end, we design a text encoder that extracts linguistic features from instructions to enhance the model’s understanding of text within images. Figure 1 illustrates the flexibility of InstructOCR in producing distinct detection and recognition results corresponding to varying input instructions.

The foundation of InstructOCR is built upon a series of instructions derived from annotated information in existing datasets, which are related to text attributes. This is because text attributes, such as text length, are crucial for enhancing the performance of models in scene text recognition tasks (Xie, Huang et al. 2019; Du, Chen et al. 2023). Building on this insight, we have designed ten templates based on these text attributes. As shown in Table 1, these templates can be used to randomly generate a wide variety of training instructions. This diversity is vital as it provides clear semantic directives, enabling the model to perform tasks aligned with human intent. Moreover, this approach fully leverages existing annotated information, thereby eliminating the need for additional annotation costs.

The main contributions of this work can be summarized as follows:

- We propose InstructOCR, an end-to-end instruction-based scene text spotting model that significantly enhances the model’s understanding of visual text through the use of instructions. To the best of our knowledge, this is the first work that innovatively integrates human language instructions into the domain of scene text spotting. Additionally, with the introduction of instructions, the framework can also be seamlessly extended to VQA tasks.

- We have meticulously designed a range of instructions specifically tailored for scene text domain, which can leverage existing publicly available scene text datasets at no additional cost. By providing precise and diverse textual annotations, our instructions enhance the model’s ability to accurately and efficiently interpret the complex variations of text in images, thereby further improving performance metrics in both scene text spotting and VQA tasks.
- We conduct extensive experiments to evaluate the effectiveness of our proposed method, and the results show that InstructOCR provides excellent performance on a range of scene text spotting datasets, achieves state-of-the-art (SOTA).

Related Work

None Sequence-based Method

Most previous scene text spotting methods have treated detection and recognition as two separate tasks (Borisyuk, Gordo et al. 2018; Liao, Shi et al. 2017; Huang, Liu et al. 2022; Liu, Shen et al. 2021). MANGO (Qiao, Chen et al. 2021) proposes a one-stage text spotting framework with Mask Attention Guidance, allowing for the direct recognition without the need for RoI operations. The Mask TextSpotter series (Lyu, Liao et al. 2018, 2021; Liao, Pang et al. 2020) leverages the advantages of character-level annotations to achieve character segmentation and can recognize scene text of arbitrary shapes. PGNet (Wang, Zhang et al. 2021) introduces a graph refinement module to optimize coarse recognition and enhance end-to-end performance. Additionally, there is also a parallel mode of detection and recognition. TTS (Kittenplon, Lavi et al. 2022) proposes a weakly supervised learning method and employs a shared query mechanism for its detector and recognizer. Estextspotter (Huang, Zhang et al. 2023) and SRSTS (Wu, Lyu et al. 2022) process text detection and recognition in parallel, thereby decoupling text recognition from dependency on detection. Inspired by the DETR (Carion, Massa et al. 2020) family models. TESTR (Zhang, Su et al. 2022) proposes a framework free from Region-of-Interest operations and heuristics-driven post-processing procedures. DeepSolo (Ye, Zhang et al. 2023) is a DETR-like model that employs a query form with explicit points sampled from the Bezier center curve of text instance lines to efficiently encode the text’s position, shape, and semantics.

Similar to scene text spotting methods, early VQA approaches also followed a two-stage task solutions (Singh, Natarajan et al. 2019; Hu, Singh et al. 2020; Yang, Lu et al. 2021), OCR results and pre-computed features are fed into a vision-and-language model. These approaches lack an interaction between text being recognized and the representation of its context.

Sequence-based Method

Inspired by the immense success in natural language processing, computer vision tasks are converging to Transformers. These approaches extract the image feature along with

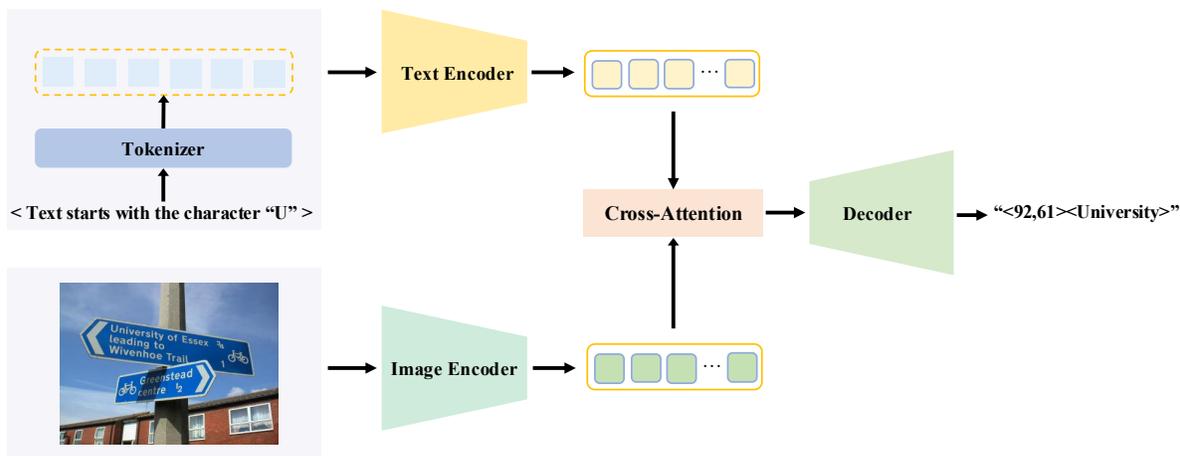


Figure 2: Main framework of InstructOCR. InstructOCR is an encoder-decoder architecture, with input branches consisting of an image encoder and a text encoder that handle visual and textual features separately.

the transformer decoder to handle various tasks by predicting sequence. Pix2seq V1 and V2 (Chen, Saxena et al. 2021, 2022) show that it is possible to output boxes and labels as a sequence of discrete tokens, allowing them to have a similar training and decoding interface as language models. (Kim, Hong et al. 2021) pioneered in employing a prompt-based sequence generation framework for document understanding.

In the end-to-end scene text spotting methods, the SPTS series (Peng, Wang et al. 2022; Liu, Zhang et al. 2023) use the central point of text regions to represent the location and auto-regressively predict coordinate tokens and word transcription tokens. Instead of using separate losses for detection and recognition, (Kil, Kim et al. 2023) unifies various detection formats, including single point, quadrilaterals, and polygons to a sequence generation paradigm. OmniParser (Wan, Song et al. 2024) sets itself apart from other text spotting methods by employing a unified architecture capable of simultaneously addressing text spotting, key information extraction, and table recognition using a single model. These methods either only utilize image features or do not provide an explicit learnable text encoder, resulting in a lack of understanding of human language instructions and failing to fully leverage human language instructions to enhance the model’s understanding of text within images.

Similar to the prevailing trend in scene text spotting tasks, current VQA tasks also primarily utilize sequence-based methods. Particularly, extremely large OCR-free image-text models have shown promising results on VQA tasks (e.g., InternVL (Chen, Wu et al. 2024), Monkey (Li et al. 2024) and Flamingo (Alayrac, Donahue et al. 2022)).

Method

We introduce InstructOCR, an end-to-end scene text spotting method that leverages instructions to guide the model’s output and enhance its understanding of the text. This enables the model to generate recognition results based on the provided instructions. The complete process is illustrated in

Figure 2, the instructions are input into the model alongside the image. In the following sections, we will detail the structure of our model.

InstructOCR Architecture

InstructOCR is an encoder-decoder architecture that encompasses a text encoder, an image encoder, and a decoder, designed to synergistically process and interpret textual information within visual contexts.

Image Encoder. The image encoder utilizes the ResNet50 architecture (He, Zhang et al. 2016) to extract features from the input image.

Text Encoder. The text encoder adopts BERT (Devlin, Chang et al. 2018), a 12-layer transformer specifically designed to extract features from instructions. By applying cross-attention between the extracted visual and textual features, we obtain instruction-based encoded features. These features are further processed by the decoder to generate the output sequence.

Decoder. The decoder of InstructOCR following the approach of SPTS (Peng, Wang et al. 2022), utilizes an auto-regressive Transformer to generate long sequences for all text instances. Each text instance is represented by a sequence composed of three parts: $[x, y, t]$, where (x, y) represents the center point coordinates and t represents the transcription text. To simplify the representation, the coordinates are uniformly discretized into integers ranging from 1 to 1000. The text is either padded or truncated to a fixed length of 25, and the maximum number of text instances in an image is set to 60. In the sequence representation, the $\langle PAD \rangle$ token is used to fill in the gaps for shorter text instances. Additionally, $\langle SOS \rangle$ and $\langle EOS \rangle$ tokens are inserted at the beginning and end of the sequence, respectively, to indicate the start and end of the sequence. For the VQA task, the question is encoded by the text encoder, while the answer is treated as a direct sequence, as illustrated in Figure 3. The entire sequence length is set to 256.

Instructions Generation

Existing OCR annotations typically consist of locations and recognition results. To eliminate the need for human-annotated instructions, we have developed an automated data processing method for generating instructions. Each word can be divided into different attributes, such as whether the initial character is capitalized, whether it contains a certain character, etc. We have designed ten types of templates, as detailed in Table 1, each of which can generate the final instruction through random conditions.

It should be noted that, unlike some methods that use indicators (Kil, Kim et al. 2023; Peng, Yang et al. 2023) or specific task prompts (Lv, Huang et al. 2023; Kim, Kim et al. 2023; Ye, Hu et al. 2023) to define different downstream tasks, we use human language as instructions. There are three reasons for this: (1) Our instructions incorporate random numbers and letters, making it impossible to enumerate the types of tasks. (2) Using human language allows us to better leverage the BERT weights pre-trained on large-scale corpora, thereby enabling the model to better understand the semantics of instructions. (3) Pre-training scene text spotting with human language instructions makes it easier to extend its use to other downstream tasks, such as Visual Question Answering (VQA).

During the training process, instructions are selected randomly. The training set is denoted as $\{x_i\}$, where each training datum x_i can be represented in the form of $\{c_i, s_i, t_i\}$. Here, c_i denotes the control instruction, s_i represents the source image, and t_i represents the target text sequence. For a given image, we have a set of annotated text instances denoted as $t' = \{t'_0, t'_1, \dots, t'_{n-1}\}$. Based on different instructions, we filter out qualifying annotated text instances, resulting in a set denoted as $t = \{t_0, t_1, \dots, t_{m-1}\}$. The objective of our method is to generate the target text sequence t_i given an input source image s_i and the corresponding instruction c_i . For the prediction phase, the instruction is set to “<Recognize all text>”.

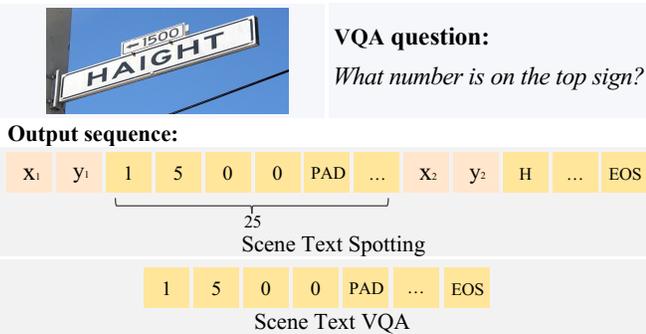


Figure 3: Examples of decoder output for scene text spotting and VQA tasks.

Loss Function

In InstructOCR, the training objective is to predict tokens, and we utilize the standard cross-entropy loss for model training. This loss function aims to maximize the likelihood

of the correct tokens during training. The mathematical expression of the cross-entropy loss is as follows:

$$\mathcal{L}_{seq} = \text{maximize} \sum_{i=1}^L w_i \log P(\tilde{s}_i | I, s_{1:i}) \quad (1)$$

where I is the input image, s is the input sequence, \tilde{s} is the output sequence, L is the length of the sequence, and w_i is the weight of the likelihood of the i -th token, which is empirically set to 1.

Experiments

Datasets

Scene Text Spotting. In our experiments, we evaluate our method on Total-Text (Ch’ng and Chan 2017), ICDAR2015 (Karatzas et al. 2015), and ICDAR2013 (Karatzas, Shafait et al. 2013). Total-Text is an arbitrarily shaped word-level scene text benchmark, with 1,255 training images and 300 testing images. ICDAR2015 contains 1,000 training images and 500 testing images for quadrilateral scene text. ICDAR2013 contains 229 training images and 233 testing images with horizontal text.

VQA for Scene Text. Scene text VQA involves answering questions about the natural scene images or reasoning about the scene text. TextVQA (Singh, Natarajan et al. 2019) contains 45,336 questions on 28,408 images that require reasoning about text to answer. ST-VQA comprises 23, 038 images sourced from a combination of public datasets (Biten, Tito et al. 2019)

Implementation Details

The Transformer encoder and decoder consist of 6 layers with 8 heads. The max length of recognition queries is 25 and the maximum number of objects is 60. The entire model is distributively trained on 32 NVIDIA A100-80G GPUs. We pretrain the model on a combination dataset that includes ICDAR2013, ICDAR2015, Total-Text, Curved Synthetic Dataset 150k (Liu, Chen et al. 2020), and ICDAR2017 MLT (Nayef, Yin et al. 2017).

We use a batch size of 320, and the pretrain model is trained for 200 epochs, with an initial 5-epoch warm-up phase. We use AdamW optimizer with a learning rate of 4.4×10^{-4} . The input image’s short size is randomly resized to a range from 704 to 1024 (intervals of 32), the maximum length of image is set as 1024. Subsequently, the model is trained for another 40 epochs, with a fixed learning rate of 1×10^{-4} , and the maximum length of image is set as 1600. Then, instructions are added, and the model is further trained for another 50 epochs. For the scene text spotting task, the model is fine-tuned on the corresponding real datasets for another 140 epochs, with a fixed learning rate of 1×10^{-5} . For the scene text VQA task, the model is fine-tuned on the TextVQA and ST-VQA datasets for another 120 epochs. At the inference stage, we resize the image’s maximum length shorter than 1920 pixels.

Comparison with Scene Text Spotting Methods

We evaluate the model using the point-based metric proposed in SPTS (Peng, Wang et al. 2022). Notably, our model

Methods	Total-Text		ICDAR2015			ICDAR2013		
	None	Full	S	W	G	S	W	G
Bounding Box-based methods								
Mask TextSpotter (Lyu, Liao et al. 2018)	65.3	77.4	83.0	77.7	73.5	93.3	91.3	88.2
FOTS(Liu et al. 2018)	-	-	83.6	79.1	65.3	92.0	90.1	84.8
Boundary TextSpotter (Wang, Lu et al. 2020)	65.0	76.1	79.7	75.2	64.1	88.2	87.7	84.1
PGNet (Wang, Zhang et al. 2021)	63.1	-	83.3	78.3	63.5	-	-	-
MANGO (Qiao, Chen et al. 2021)	72.9	83.6	81.8	78.9	67.3	92.9	92.7	88.3
ABCNet v2 (Liu, Shen et al. 2021)	70.4	78.1	82.7	78.5	73.0	-	-	-
SwinTextSpotter(Huang, Liu et al. 2022)	74.3	84.1	83.9	77.3	70.5	-	-	-
SRSTS(Wu, Lyu et al. 2022)	78.8	86.3	85.6	81.7	74.5	-	-	-
GLASS(Ronen, Tsiper et al. 2022)	79.9	86.2	84.7	80.1	76.3	-	-	-
TESTR(Zhang, Su et al. 2022)	73.3	83.9	85.2	79.4	73.6	-	-	-
TTS(Kittenplon, Lavi et al. 2022)	78.2	86.3	85.2	81.7	77.4	-	-	-
DeepSolo(Ye, Zhang et al. 2023)	82.5	88.7	88.0	83.5	79.1	-	-	-
ESTextSpotter(Huang, Zhang et al. 2023)	80.8	87.1	87.5	83.0	78.1	-	-	-
UNITS-Swin(Kil, Kim et al. 2023)	78.7	86.0	89.0	84.1	80.3	-	-	-
DNTextSpotter(Qiao, Xie et al. 2024)	84.5	89.8	88.7	84.3	79.9	-	-	-
Point-based methods								
TOSS(Tang, Qiao et al. 2022)	65.1	74.8	65.9	59.6	52.4	-	-	-
SPTS(Peng, Wang et al. 2022)	74.2	82.4	77.5	70.2	65.8	93.3	91.7	88.5
SPTS-v2(Liu, Zhang et al. 2023)	75.5	84.0	82.3	77.7	72.6	93.9	91.8	88.6
InstructOCR	77.1	84.1	82.5	77.1	72.1	93.3	92.4	88.8
InstructOCR†	83.4	88.3	87.5	84.2	80.6	94.9	94.1	91.7

Table 2: Scene text spotting results on Total-Text, ICDAR2015, and ICDAR2013. ‘None’ means lexicon-free. ‘Full’ indicates that we use all the words that appeared in the test set. ‘S’, ‘W’, and ‘G’ represent recognition with ‘Strong’, ‘Weak’, and ‘Generic’ lexicons, respectively. † denotes the incorporation of TextOCR and HierText datasets in the training data.

adeptly outputs the coordinates of single-point and has been compared with other point-based methods. We have also listed methods based on bounding boxes for comparison. However, some of these methods use additional datasets, such as UNITS (Kil, Kim et al. 2023), which additionally utilizes the TextOCR (Singh, Pang et al. 2021) and HierText (Long, Qin et al. 2022) datasets. To ensure a fair performance comparison, we have augmented the pre-training with these two additional datasets. Subsequent ablation studies also include these two datasets. Figure 4 displays some exemplary visualization results for scene text spotting.

Total-Text: Arbitrarily-Shaped Text. To validate the generalization ability of our method for arbitrarily shaped scene text spotting, we tested our approach on Total-Text. The scene text spotting results are shown in Table 2, where InstructOCR significantly surpasses SPTS-V2 on TotalText, by 1.6% without a dictionary and by 0.1% with a “full” dictionary.

ICDAR2015: Multi-oriented Text. To evaluate the robustness of our method for multi-oriented text, we conduct experiments on ICDAR2015, with the results shown in Table 2. Our method outperforms the previous single-point methods across all dictionary settings. Notably, in the strong dictionary setting, InstructOCR achieves an Hmean of 82.5%, and after adding more training data, it reaches 87.5%.

ICDAR2013: Horizontal text. To further compare with

point-based methods, we conduct experimental comparisons on ICDAR2013, which already has high baseline metrics, resulting in relatively smaller improvements. In the weak dictionary setting, InstructOCR achieves an Hmean of 92.4%, which is 0.6% higher than SPTS-v2.

Discussion. InstructOCR endows the model with the ability to recognize scene text and seamlessly integrate it with the visual context. It is evident that the incorporation of human language instructions indeed significantly aids the model’s understanding of text within images. As a result, compared to the SPTS-V1, there has been a substantial improvement in metrics without a dictionary. Specifically, the improvements are: +2.9% on the Total-Text dataset, +0.3% on the ICDAR2013 dataset, and +6.3% on the ICDAR2015 dataset. This also validates our hypothesis that aligning human language instructions with visual text enhances accuracy.

Applicability to Scene-Text VQA

In this section, we further explore other scene-text related domains (Table 3). We show that InstructOCR is also applicable on VQA tasks with a considerable accuracy of 42.0% on TextVQA and 45.8% on ST-VQA. While apples-to-apples comparison is difficult due to different data and parameter sizes, we emphasize the applicability to the VQA task. Specifically, most of the recent works utilize strong



Figure 4: Visual results on Total-Text, ICDAR2015 and ICDAR2013. Our model can effectively handle curved, distorted, and blurred scene text.

Methods	Param	Data	TextVQA	ST-VQA
TAP	160M	1.5M	54.7	59.8
GIT _{Large}	347M	20M	37.5	44.6
PreSTU	278M	13M	54.5	62.6
WSCOB	202M	12M	56.2	62.6
LLaVA1.5	7B	1.2M	38.7	38.1
InternVL	8B	4.98B	59.8	62.2
Monkey	8B	1.44M	64.3	54.7
Qwen-VL	8B	1.4B	63.8	55.9
InstructOCR	78M	0.2M	42.0	45.8

Table 3: The public benchmark of TAP (Yang, Lu et al. 2021), GIT_{Large} (Wang, Yang et al. 2022), PreSTU (Kil et al. 2023) WSCOB (Kim, Kim et al. 2023), LLaVA1.5 (Liu, Li et al. 2023), InternVL (Chen, Wu et al. 2024), Monkey (Li et al. 2024) and Qwen-VL (Bai, Bai et al. 2023) on TextVQA (acc) (Singh, Natarajan et al. 2019) and ST-VQA (ANLS) (Biten, Tito et al. 2019) for scene text VQA.

backbones such as ViT (Dosovitskiy, Beyer et al. 2020) and large language models such as T5_{large} (Raffel, Shazeer et al. 2020), while ours adopt ResNet-50 and BERT. Our model has, to the best of our knowledge, the least number of parameters (78M) among the similar levels of VQA performance, whereas other approaches range from hundreds of millions to even billions of parameters. It is well known that the VQA performance goes up as more pre-training data is included. The aforementioned methods employ up to millions of text-image pairs, not to mention the multimodal large language models that utilize various forms of large-scale data. Our model is not specifically tailored for VQA tasks and is trained only on 0.2M scene text images. These promising results show the utility and applicability of InstructOCR on image understanding. Our research further extends the boundaries of small-scale models in the VQA task.

Ablation Studies

We conduct a thorough analysis to understand the individual contributions of each component in our framework, with a particular focus on the effectiveness of the VQA tasks.

Text Encoder		INS	ICDAR2015			
W	T		S	W	G	None
✗	✗	✗	86.5	82.7	79.1	77.1
✓	✗	✗	86.9	82.8	79.2	77.3
✓	✓	✗	87.1	83.4	80.6	78.8
✓	✓	✓	87.5	84.2	80.6	78.9
			(+1.0)	(+1.5)	(+1.5)	(+1.8)

Table 4: Ablation study of our proposed components on ICDAR2015, ‘W’, ‘T’, and ‘INS’ refer to including a text encoder in end-to-end training, training the text encoder specifically for end-to-end training, and incorporating instructions into the end-to-end training, respectively. ‘S’, ‘W’, ‘G’, and ‘None’ represent recognition with ‘Strong’, ‘Weak’, and ‘Generic’ lexicons and without lexicons, respectively

Impact of Module Integration In this section, we perform a comparative analysis to evaluate the effects of several proposed modules. Prior to validating the effectiveness of instructions, we assess the impact of incorporating a text encoder into the model. We compare the model’s performance without a text encoder to its performance with an included and frozen text encoder. We then evaluate the performance metrics when the text encoder is integrated and trained throughout the entire process. The results indicate that the inclusion and active training of the text encoder significantly enhance the capabilities of InstructOCR. Finally, we incorporate instructions, observing further improvements in performance metrics. As shown in Table 4, in the strong dictionary setting, the inclusion of text encoder results in a 1.0% improvement, while the addition of instructions leads to a further increase of 0.4%. The experimental results demonstrate that incorporating a text encoder enables the model to more effectively comprehend text within images, and the addition of instructions can further facilitate this. Figure 5 presents some examples where the results became correct after incorporating human language instructions.

The performance of InstructOCR on the VQA task In this section, we conduct experiments to verify the effectiveness of our model in VQA tasks. Since our model is pre-trained for scene text spotting, we choose the ST-VQA and

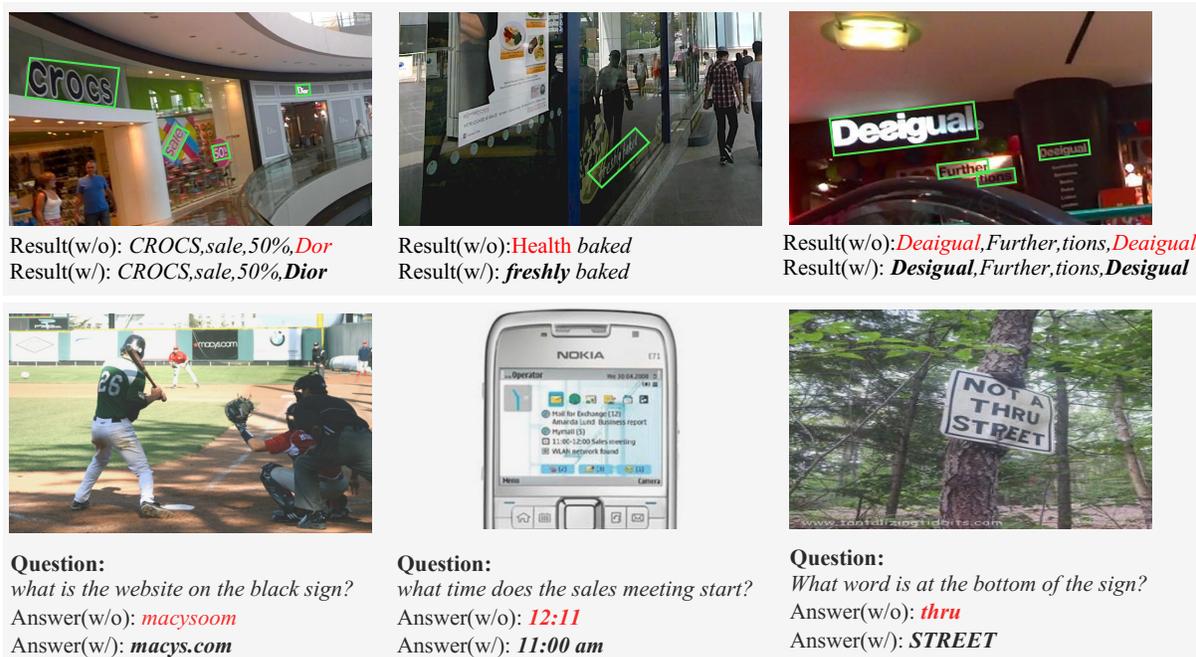


Figure 5: Visualization of correct recognition results on the scene text spotting and VQA datasets after incorporating instructions. The first row corresponds to the scene text spotting dataset, while the second row corresponds to the VQA dataset. “W/” and “W/O” indicate whether OCR training was conducted with or without instructions, respectively.

TextVQA datasets, which relate to natural scene images. We hypothesize that our framework can be transferred to VQA tasks and that pretraining on text spotting is beneficial for downstream VQA tasks. Table 5 supports this argument. Significant improvements are generally observed in VQA tasks when pretraining on text spotting is applied. This is likely because accurate character recognition is a prerequisite for better text understanding, and VQA is closely related to OCR as most answers exist in images containing text. Moreover, the integration of our proposed instructions leads to enhanced image and text understanding, resulting in an improvement of 2.1% on ST-VQA datasets and 2.6% on TextVQA datasets. We think that human language instructions during the pre-training stage are beneficial to understanding capability, thereby enabling it to better adapt to more complex understanding tasks in downstream applications. Figure 5 demonstrates the impact of with instructions and without instructions on the output results.

Limitation

Owing to the VQA task requiring a large amount of data for pre-training and instruction fine-tuning, and different methods using inconsistent data amounts that make it impossible to align data volume. We only use data from the scene text spotting task for pre-training and only use the ST-VQA and TextVQA datasets for instruction fine-tuning. We do not add a large amount of data to demonstrate the best performance of InstructOCR on the VQA task.

Exp	W/O	W	INS	ST-VQA	TextVQA
1	✓			6.7	5.6
2		✓		43.7	39.4
3		✓	✓	45.8 (+2.1)	42.0 (+2.6)

Table 5: Test results for different methods on ST-VQA (ANLS) and TextVQA (Acc.) datasets. “W/O”, “W”, and “INS” refer to training the model without 0.2M OCR data pre-training, training the model with 0.2M OCR data pre-training, training the model with instructions, respectively.

Conclusion

In this paper, we propose InstructOCR, a novel instruction-based scene text spotting model that leverages instructions to enhance the understanding of text within images. The model integrates a text encoder that processes instructions and a visual-text fusion module that combines image features with these instructions to guide the decoding process. With the introduction of human language instructions and our meticulously designed instruction set based on text attributes, InstructOCR demonstrates an extraordinary ability to comprehend and process textual information within natural scenes, indicating the benefits of aligning human language instructions with visual text for OCR tasks. Future research will explore increasing training data and incorporating more instructions, thereby enabling the model to achieve higher metrics in VQA tasks and address a broader range of OCR challenges.

References

- Alayrac, J.-B.; Donahue, J.; et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35: 23716–23736.
- Bai, J.; Bai, S.; et al. 2023. Qwen-VL: A Versatile Vision-Language Model for Understanding, Localization, Text Reading, and Beyond. *arXiv preprint arXiv:2308.12966*.
- Biten, A. F.; Tito, R.; et al. 2019. Scene text visual question answering. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4291–4301.
- Borisyuk, F.; Gordo, A.; et al. 2018. Rosetta: Large scale system for text detection and recognition in images. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*, 71–79.
- Brooks, T.; Holynski, A.; and Efros, A. A. 2023. Instruct-pix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18392–18402.
- Carion, N.; Massa, F.; et al. 2020. End-to-end object detection with transformers. In *European conference on computer vision*, 213–229. Springer.
- Chen, T.; Saxena, S.; et al. 2021. Pix2seq: A language modeling framework for object detection. *arXiv preprint arXiv:2109.10852*.
- Chen, T.; Saxena, S.; et al. 2022. A unified sequence interface for vision tasks. *Advances in Neural Information Processing Systems*, 35: 31333–31346.
- Chen, Z.; Wu, J.; et al. 2024. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 24185–24198.
- Ch’ng, C. K.; and Chan, C. S. 2017. Total-text: A comprehensive dataset for scene text detection and recognition. In *2017 14th IAPR international conference on document analysis and recognition (ICDAR)*, volume 1, 935–942. IEEE.
- Devlin, J.; Chang, M.-W.; et al. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Dosovitskiy, A.; Beyer, L.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Du, Y.; Chen, Z.; et al. 2023. Context Perception Parallel Decoder for Scene Text Recognition. *arXiv preprint arXiv:2307.12270*.
- Duan, C.; Fu, P.; et al. 2024. ODM: A Text-Image Further Alignment Pre-training Approach for Scene Text Detection and Spotting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15587–15597.
- Garcia-Bordils, S.; Karatzas, D.; et al. 2024. STEP-Towards Structured Scene-Text Spotting. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 883–892.
- Geng, Z.; Yang, B.; Hang, T.; et al. 2023. Instructdiffusion: A generalist modeling interface for vision tasks. *arXiv preprint arXiv:2309.03895*.
- He, K.; Zhang, X.; et al. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hu, R.; Singh, A.; et al. 2020. Iterative answer prediction with pointer-augmented multimodal transformers for textvqa. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9992–10002.
- Huang, M.; Liu, Y.; et al. 2022. Swintextspotter: Scene text spotting via better synergy between text detection and text recognition. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4593–4603.
- Huang, M.; Zhang, J.; et al. 2023. Estextspotter: Towards better scene text spotting with explicit synergy in transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 19495–19505.
- Karatzas, D.; Shafait, F.; et al. 2013. ICDAR 2013 robust reading competition. In *12th international conference on document analysis and recognition*, 1484–1493. IEEE.
- Karatzas, D.; et al. 2015. ICDAR 2015 competition on robust reading. In *13th international conference on document analysis and recognition*, 1156–1160. IEEE.
- Kil, J.; et al. 2023. Prestu: Pre-training for scene-text understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15270–15280.
- Kil, T.; Kim, S.; et al. 2023. Towards Unified Scene Text Spotting based on Sequence Generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15223–15232.
- Kim, D.; Kim, Y.; et al. 2023. SCOB: Universal Text Understanding via Character-wise Supervised Contrastive Learning with Online Text Rendering for Bridging Domain Gap. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 19562–19573.
- Kim, G.; Hong, T.; et al. 2021. Donut: Document understanding transformer without ocr. *arXiv preprint arXiv:2111.15664*, 7: 15.
- Kirillov, A.; Mintun, E.; et al. 2023. Segment anything. *arXiv preprint arXiv:2304.02643*.
- Kittenplon, Y.; Lavi, I.; et al. 2022. Towards weakly-supervised text spotting using a multi-task transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4604–4613.
- Li, Z.; Yang, B.; Liu, Q.; et al. 2024. Monkey: Image resolution and text label are important things for large multi-modal models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 26763–26773.
- Liao, M.; Pang, G.; et al. 2020. Mask textspotter v3: Segmentation proposal network for robust scene text spotting. In *European conference on computer vision*, 706–722.
- Liao, M.; Shi, B.; et al. 2017. Textboxes: A fast text detector with a single deep neural network. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31.
- Liu, H.; Li, C.; et al. 2023. Visual Instruction Tuning.
- Liu, S.; Zeng, Z.; et al. 2023. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*.

- Liu, X.; et al. 2018. Fots: Fast oriented text spotting with a unified network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5676–5685.
- Liu, Y.; Chen, H.; et al. 2020. Abcnet: Real-time scene text spotting with adaptive bezier-curve network. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9809–9818.
- Liu, Y.; Shen, C.; et al. 2021. Abcnet v2: Adaptive bezier-curve network for real-time end-to-end text spotting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11): 8048–8064.
- Liu, Y.; Zhang, J.; et al. 2023. SPTS v2: single-point scene text spotting. *arXiv preprint arXiv:2301.01635*.
- Long, S.; Qin, S.; et al. 2022. Towards end-to-end unified scene text detection and layout analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1049–1059.
- Lv, T.; Huang, Y.; et al. 2023. Kosmos-2.5: A multimodal literate model. *arXiv preprint arXiv:2309.11419*.
- Lyu, P.; Liao, M.; et al. 2018. Mask textspotter: An end-to-end trainable neural network for spotting text with arbitrary shapes. In *Proceedings of the European conference on computer vision (ECCV)*, 67–83.
- Lyu, P.; Liao, M.; et al. 2021. Mask textspotter: An end-to-end trainable neural network for spotting text with arbitrary shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(2): 532–548.
- Nayef, N.; Yin, F.; et al. 2017. Icdar2017 robust reading challenge on multi-lingual scene text detection and script identification-rrc-mlt. In *2017 14th IAPR international conference on document analysis and recognition (ICDAR)*, volume 1, 1454–1459. IEEE.
- Peng, D.; Wang, X.; et al. 2022. SPTS: single-point text spotting. In *Proceedings of the 30th ACM International Conference on Multimedia*, 4272–4281.
- Peng, D.; Yang, Z.; et al. 2023. Upocr: Towards unified pixel-level ocr interface. *arXiv preprint arXiv:2312.02694*.
- Qiao, L.; Chen, Y.; et al. 2021. Mango: A mask attention guided one-stage scene text spotter. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2467–2476.
- Qiao, Q.; Xie, Y.; et al. 2024. DNTextSpotter: Arbitrary-shaped scene text spotting via improved denoising training. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 10134–10143.
- Radford, A.; Kim, J. W.; Hallacy, C.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Raffel, C.; Shazeer, N.; et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140): 1–67.
- Ronen, R.; Tsiper, S.; et al. 2022. Glass: Global to local attention for scene-text spotting. In *European Conference on Computer Vision*, 249–266. Springer.
- Singh, A.; Natarajan, V.; et al. 2019. Towards vqa models that can read. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8317–8326.
- Singh, A.; Pang, G.; et al. 2021. Textocr: Towards large-scale end-to-end reasoning for arbitrary-shaped scene text. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8802–8812.
- Tang, J.; Qiao, S.; et al. 2022. You can even annotate text with voice: Transcription-only-supervised text spotting. In *Proceedings of the 30th ACM International Conference on Multimedia*, 4154–4163.
- Wan, J.; Song, S.; et al. 2024. OmniParser: A Unified Framework for Text Spotting, Key Information Extraction and Table Recognition. *arXiv preprint arXiv:2403.19128*.
- Wang, H.; Lu, P.; et al. 2020. All you need is boundary: Toward arbitrary-shaped text spotting. In *Proceedings of the AAAI conference on artificial intelligence*, 12160–12167.
- Wang, J.; Yang, Z.; et al. 2022. Git: A generative image-to-text transformer for vision and language. *arXiv preprint arXiv:2205.14100*.
- Wang, P.; Zhang, C.; et al. 2021. Pgnnet: Real-time arbitrarily-shaped text spotting with point gathering network. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, 2782–2790.
- Wu, J.; Lyu, P.; et al. 2022. Decoupling recognition from detection: Single shot self-reliant scene text spotter. In *Proceedings of the 30th ACM International Conference on Multimedia*, 1319–1328.
- Xie, Z.; Huang, Y.; et al. 2019. Aggregation cross-entropy for sequence recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6538–6547.
- Xue, C.; Zhang, W.; et al. 2022. Language matters: A weakly supervised vision-language pre-training approach for scene text detection and spotting. In *European Conference on Computer Vision*, 284–302. Springer.
- Yang, Z.; Lu, Y.; et al. 2021. Tap: Text-aware pre-training for text-vqa and text-caption. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8751–8761.
- Ye, J.; Hu, A.; et al. 2023. Ureader: Universal ocr-free visually-situated language understanding with multimodal large language model. *arXiv preprint arXiv:2310.05126*.
- Ye, M.; Zhang, J.; et al. 2023. DeepSolo: Let transformer decoder with explicit points solo for text spotting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19348–19357.
- Yu, W.; Liu, Y.; et al. 2023. Turning a clip model into a scene text detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6978–6988.
- Yu, W.; Liu, Y.; et al. 2024. Turning a clip model into a scene text spotter. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Zhang, X.; Su, Y.; et al. 2022. Text spotting transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9519–9528.