

# An Empirical Study of Scaling Law for OCR

Miao Rang\* Zhenni Bi\* Chuanjian Liu Yunhe Wang† Kai Han†  
Huawei Noah’s Ark Lab  
{rangmiaol,bizhenni,liuchuanjian,  
yunhe.wang,kai.han}@huawei.com

## Abstract

The laws of model size, data volume, computation and model performance have been extensively studied in the field of Natural Language Processing (NLP). However, the scaling laws in Optical Character Recognition (OCR) have not yet been investigated. To address this, we conducted comprehensive studies that involved examining the correlations between performance and the scale of models, data volume and computation in the field of text recognition. Conclusively, the study demonstrates smooth power laws between performance and model size, as well as training data volume, when other influencing factors are held constant. Additionally, we have constructed a large-scale dataset called **REBU-Syn**, which comprises 6 million real samples and 18 million synthetic samples. Based on our scaling law and new dataset, we have successfully trained a scene text recognition model, achieving a new state-of-the-art on 6 common test benchmarks with a top-1 average accuracy of 97.42%.

## 1. Introduction

Optical Character Recognition (OCR) is a technology designed to detect and interpret textual content within images, such as scanned documents and photographs. The field of OCR has undergone significant evolution in recent years, largely driven by advancements in deep learning. These advancements have enabled OCR system to be successfully deployed in a variety of real-world applications, including the automatic recognition of license plates, official documents and bills. A typical OCR system operates in two primary phases: text detection and text recognition. In this context, our focus is on the text recognition phase, which involves identifying and extracting text from predefined bounding boxes.

With the introduction of large-scale models in deep

\*These authors contributed equally to this work.

†Corresponding Author.




Figure 1. Mean word accuracy vs Parameters on the 6 common test benchmarks. P-Ti, P-S and P-B refer to PARSeq-Ti, PARSeq-S and PARSeq-B, respectively. \* denotes our result based on the power law.

learning, an increasing number of academics are focusing on the potential and growth trends of these models, hoping that they will contribute to the development and design of future models. In the field of NLP [6, 11], numerous experiments have been carried out to investigate scaling model laws [17, 23, 31, 94]. The results show that the larger the volume of data fed into the neural network, the better its performance. Therefore, large language models trained on vast amounts of data have dominated the field of NLP. However, in the OCR domain, research predominantly focuses on enhancing model performance using fixed data sizes and model sizes [4, 42, 70]. Studies specifically addressing the scaling laws in large OCR models are noticeably sparse, which casts uncertainty on the potential impact of large-scale models and substantial data volumes in OCR. Transformer-based models have achieved state-of-the-art performance in various text recognition tasks and challenges [42, 63, 68, 78].

In this paper, we explore the scaling laws of text recognition using transformer-based models. Our focus is on unraveling the relationships between model size, data volume and computation in relation to model performance. Our experimental framework encompasses a wide range of models, with parameter counts ranging from 50 million to 1 billion, and datasets that vary from 1 million to 1 billion training samples. Additionally, we extend our exploration to computational durations ranging from 100 to 1000 hours. This comprehensive analysis allows us to draw insightful conclusions about the scaling law in text recognition. Furthermore, we introduce a novel dataset called **REBU-Syn**, which combines real-world and synthetic data. This dataset has been meticulously compiled from existing public datasets, providing a valuable resource for further research in this field.

Throughout our research, we develop an advanced method for large-scale training. This method involves a comprehensive examination of various strategies, such as optimizing training hyperparameters, analyzing data distributions and utilizing pre-training techniques. Our objective is to create a model characterized by exceptional precision and accuracy. The culmination of these efforts is the training of CLIP4STR-L using REBU-Syn. This approach results in achieving a groundbreaking state-of-the-art performance of 97.42% on the test benchmark (see Fig 1). The following is a compilation of additional scaling laws for OCR observations:

- The scaling law holds in the field of OCR. There exist smooth power laws between the size of the model, the volume of data, computation and performance.
- Large-scale models make better use of samples than small-scale models which means that large models achieve lower error rate with fixed amount of data.
- The proportion of training data from different sources is crucial for model training.
- Models pre-trained on OCR-related data are more effective in OCR tasks than models pretrained on general images like ImageNet.

## 2. Related Work

**Model Scale** Recent research has extensively explored the scaling laws for Transformer language models, particularly in the field of NLP [31, 46]. These studies have established a set of universal principles for modeling scale. However, research specifically addressing OCR remains scarce. Transformer-based methods, known for their higher tolerance to increased model depth and width, have been applied in various fields [13, 16, 19, 40, 62]. This study leverages these methods, with a specific focus on their application in OCR, to provide guidance on making effective adjustments in model size.

**Data Scale** In the domain of image recognition, the scale of data plays a critical role. The performance of various models is significantly influenced by the size of the datasets used [7, 64, 66]. While different model types require varying data volumes, some previous methods [3] explored the impact of OCR recognition tasks on different data scales, but their main focus was on CNN-based or attention-based approaches [10, 12, 37, 47, 70, 74], and they focus solely on reducing the data scale. Furthermore, the availability of public datasets has facilitated extensive research and experimentation in this field [7, 64, 66]. This paper aims to build upon these foundations by conducting a comprehensive investigation into the effects of data scale, both at the lower and upper limits, as well as the distribution of data in OCR tasks. Additionally, this study offers new insights into the alignment of real and synthetic data during the training of optimal models, filling a gap in current research.

**Scaling Laws** The rapid advancement of Large Language Models (LLMs) like ChatGPT [57] and GPT-4 [58] has sparked research into universal scaling laws [31, 40] in deep learning. These studies explore the relationship between model size, data volume, computation and performance, providing training principles for large models in NLP. [27] describes laws for autoregressive generative modeling. Similar scaling theories have also emerged in the field of computer vision [94], as demonstrated by the training of ViT-G with 2B parameters [94]. Furthermore, recent work has been done on the scaling law of CLIP [61] has revealed task- and dataset-dependent scaling behaviors [43]. Building upon these foundational insights, this study represents a unique exploration of scaling laws within the context of OCR. Specifically, it explores the allocation of parameters and the internal structure of the transformer model, with the aim of optimizing performance for text recognition. This investigation makes a unique contribution to the expanding body of research on scaling laws, particularly in the underexplored domain of OCR.

## 3. Method Details

In this paper, our primary focus is to explore the scaling laws for the transfer performance of Transformer-based models in text recognition tasks. Concurrently, we have amalgamated all publicly available datasets to construct the REBU-Syn dataset. This paper also includes a thorough analysis of the data proportions obtained from various sources. Finally, we will provide a detailed overview of the training parameter settings used in our study.

### 3.1. Model Scaling

**TrOCR** TrOCR [42] is an end-to-end OCR model that utilizes the Transformer architecture. It integrates pre-trained Computer Vision (CV) and NLP models. And it

is the first work that jointly leverages image Transformer and text Transformer for text recognition tasks in OCR. The scaling laws of Transformer language models [31] and Vision Transformers [94] have been studied, the scaling laws for models in the OCR field have not yet been explored. Based on this, we scaled the TrOCR model size and attempted to analyze the accuracy change curves of models with varying sizes.

In TrOCR, the encoder and decoder parts utilize pre-trained image Transformer and text Transformer models, respectively. These pre-trained models utilize large-scale unlabeled data for image understanding and language modeling. As a result, TrOCR does not require additional language models for post-processing, and the model outperforms current state-of-the-art models in tasks related to printing and handwriting recognition. In order to continue benefiting from pre-training for related tasks, we select the most suitable combination of encoder and decoder in TrOCR for scaling.

For TrOCR-S, we use DeiT<sub>SMALL</sub> [79] to initialize the encoder and MiniLM [85] to initialize the decoder. TrOCR-B uses BEiT<sub>BASE</sub> [6] to initialize the encoder and RoBERTa<sub>LARGE</sub> [49] to initialize the decoder. TrOCR-L and TrOCR-H utilize BEiT<sub>LARGE</sub> to initialize the encoder and RoBERTa<sub>LARGE</sub> to initialize the decoder. The model’s parameters range from 43.09 million to 1 billion, and the details of the parameters are shown in Table 1.

Model	Encoder			FLOPs (G)	Params (M)
	layers	hidden sizes	heads		
TROCR-S	12	384	6	13.31	43.09
TROCR-B	12	768	12	62.01	281.87
TROCR-L	24	1024	16	191.00	505.50
TROCR-H	48	1200	16	497.91	1037.61

Table 1. Architecture specifications of TrOCR variants.

**PARSeq** PARSeq [8] follows an encoder-decoder architecture. The encoder part utilizes the Vision Transformer (ViT) model to extract image features, while the decoder follows the same architecture as the pre-LayerNorm [83] Transformer decoder in this study utilizes twice the number of attention heads, where  $n_{head} = d_{model}/32$ . In contrast to the standard ViT, the encoder removes the  $[class]$  token and inputs all the output tokens into the decoder.

PARSeq has two models in the original paper, PARSeq-Ti and PARSeq-S. In order to investigate the law of large models in the field of OCR, the scaling law of the ViT model was demonstrated. [94]. Based on this, we scaled PARSeq to 4 different sizes. On the basis of the original paper PARSeq-S, the model was expanded to 3 sizes: PARSeq-B, PARSeq-L, and PARSeq-H. The scale of the model was also expanded from 22 million to 0.6 billion. The configurations with different scale PARSeq models can be seen in Table 2.

Model	Encoder			FLOPs (G)	Params (M)
	layers	hidden sizes	heads		
ParSeq-S	12	384	6	2.76	22.51
ParSeq-B	12	768	12	17.20	104.01
ParSeq-L	24	1024	16	49.90	335.92
ParSeq-H	32	1280	16	98.10	682.14

Table 2. Architecture specifications of PARSeq variants.

### 3.2. Dataset

**Training Dataset** The training datasets for text recognition are typically categorized into synthetic and real data. Historically, scene text recognition models primarily relied on synthetic data due to the scarcity of real-world data. However, the recent increase in the availability of real data has shifted this trend. It has been observed that models trained on real data tend to be more sample-efficient compared to those trained on synthetic data. In light of this, we meticulously collected both synthetic and real data, employing various strategies to construct the REBU-Syn dataset. This dataset comprises approximately 6M real data samples and 18M public synthetic data samples, as detailed in Table 3. The ratio of synthetic to real data in REBU-Syn is 3:1. Furthermore, we utilized synthesis technology to generate an additional 60M data samples, similar to MJST, termed *MJST+*.

**Real Dataset** We gather real images from 4 widely-accessible datasets to assemble the REBU. The **R** component consists of commonly used real data [8], including COCO-Text (COCO) [80], RCTW17 [72], Uber-Text (Uber) [95], ArT [15], LSVT [77], MLT19 [56], ReCTS [48], TextOCR [76] and OpenVINO [35]. A detailed analysis of these datasets is presented in [8]. **U**, another segment of REBU, includes 4 million labeled images across 14 datasets, collectively referred to as Union14M-L [8]. **B** represents the training data from benchmark sources, encompassing datasets such as IIIT 5k-word (IIIT5k) [54], Street View Text (SVT) [82], ICDAR13 [32] and ICDAR15 [33]. Furthermore, **E** is composed of images from two commonly used real datasets in text detection tasks, namely Total Text [14] and CTW1500 [93]. This inclusion significantly expands the range of real data in our collection.

**Public Synthetic Dataset** MJSynth (MJ) [29] and SynthText (ST) [25] are two widely-utilized synthetic datasets in the field of scene text recognition, containing 8.9M million and 5.5M million data samples respectively. Additionally, we incorporated two other composite datasets into our study. Curved SynthText (CST) and SynthAdd (SA) [28]. CST is specifically designed for text detection tasks, primarily comprising curved text data. SA, generated with the SynthText engine, is aimed at synthesizing less common characters, including punctuation marks.

**Generate Synthetic Dataset** To closely align with the MJ and ST datasets, we created *MJST+* using two data

generation tools: TextRecognitionDataGenerator<sup>1</sup> and SynthText<sup>2</sup>. The TextRecognitionDataGenerator is adept at producing data that mimics complex scenes, encompassing effects such as blurring, tilting and distortion. SynthText, on the other hand, specializes in synthesizing data akin to ST, resulting in samples that blend more seamlessly with natural scenes.

To augment the diversity of the generated corpus, we sourced 700,000 corpus entries from the most extensively utilized English corpus website globally<sup>3</sup>. For the background selection in our synthesized images, we employed natural scene pictures provided by SynthText as the backdrop. Utilizing these two synthesis methods, we successfully synthesized a total of 60M data samples.

Source	Dataset	Instances
Public Real	Real (R)	3.3M
Public Real	Extra Real Data (E)	15k
Public Real	BenchMark (B)	7.5K
Public Real	Union14M (U)	3.1M
Public Synthetic	MJ	5.5M
Public Synthetic	ST	8.9M
Public Synthetic	CST	1.8M
Public Synthetic	SA	1.2M
Generate Synthetic	<i>MJST</i> <sup>+</sup>	60M

Table 3. Statistics of REBU-Syn datasets, including Public Real and Public Synthetic. Generate Synthetic can be used additionally.

**Test Dataset** To assess the performance of our model, we utilized 6 publicly available real scene text datasets: IIIT5k-Words (IIIT5k) [54], Street View Text (SVT) [82], ICDAR 2013 (IC13) [32], ICDAR 2015 (IC15) [33], SVT-Perspective (SVTP) [59] and CUTE80 (CUTE) [65]. Both the IC13 and IC15 test sets have various subdivisions. We follow the division proposed by Yu *et al* [91], using a version of the IC15 test set containing 1,811 images, and the IC13 test set comprising 857 images.

However, to address challenges posed by differing annotation formats and the presence of duplicate, non-Latin, and damaged samples, we employed the following data fusion strategy:

- **Polygonal Text** We sourced synthesized data from datasets used in text detection tasks with polygonal annotation boxes, such as Curved SyntheText, SyntheAdd and STR Benchmark. To adapt these polygonal texts for use, we improved upon the method proposed in [8]. Our approach involves identifying the minimum bounding box of the polygon and applying a perspective transformation, avoiding direct clipping using maximum and minimum coordinates. This method retains challenging samples, as suggested in [8], while minimizing background interference, thus enabling the recognizer to focus on pertinent areas.

<sup>1</sup><https://github.com/Belval/TextRecognitionDataGenerator>

<sup>2</sup><https://github.com/ankush-me/SynthText>

<sup>3</sup><https://www.english-corpora.org/corpora.asp>

- **Remove invalid chars and samples** Focusing on Latin characters, which have extensive data availability, we retained samples composed only of letters and symbols. Samples not in our predefined dictionary were discarded.
- **Remove duplicate data** As we integrated multiple datasets, some of which overlapped, we meticulously removed any duplicate entries.

### 3.3. Experiment Settings

We utilized the publicly available implementations of TrOCR and ParSeq as baseline models. To achieve optimal performance, we tailored the number of training epochs and adjusted the learning rates. The specific implementation details are as follows:

**Hyper-Parameters** For our experiments, we use V100 GPUs equipped with 32GB of memory to train all models. The learning rates are set differently for various models. Specifically, TrOCR-S is trained with a batch size of 1024 and a learning rate of  $4e-4$ . TrOCR-B employs a batch size of 256 with a learning rate of  $1e-4$ , and TrOCR-L operates with a batch size of 128 and a learning rate of  $4e-5$ . We use BPE [67] of Fairseq and SentencePiece [36] for tokenizing text lines into word pieces. For PARSeq models, a consistent learning rate of  $7e-4$  is used, with the batch size adjusted to be as close to 1024 as possible.

**Evaluation Metrics** Word accuracy was the primary metric for evaluating the datasets of scene text. In this work, we standardized the final output string to match the commonly used 36-character set (lowercase alphanumeric) to ensure a fair comparison across different models and datasets.

## 4. Results and Analysis

### 4.1. Smooth Power Laws

Model performance is primarily influenced by three variables: the number of model parameters  $N$ , the volume of the training data  $D$ , and the computation of the model  $C$ . In this section, we explore the power laws among these influential factors with model performance  $E$ . To effectively characterize the scaling of OCR models, we have conducted training with a variety of models, including TrOCR and PARSeq.

Model	Avg	Regular Text			Irregular Text		
		IC13	IIIT5k	SVT	CUTE80	IC15	SVTP
		857	3,000	647	288	1,811	645
TrOCR-S	81.93	90.65	85.60	85.94	74.31	72.73	78.44
TrOCR-B	88.56	96.14	92.00	91.56	80.56	81.14	83.91
TrOCR-L	89.84	96.50	92.90	92.81	84.38	82.52	86.72
TrOCR-H	90.94	97.31	93.57	94.22	87.50	83.79	88.59

Table 4. Word accuracy with different TrOCR model sizes. *Train data*: Synthetic datasets with MJ and ST.






Figure 2. Improvement in TrOCR model performance with increasing model size, data volume, and training computation. Model performance is measured by calculating the average word error rate on 6 common test benchmarks **Left**: Evaluation of model performance with changing model sizes. **Center**: Evaluation of model performance with varying data volumes. **Right**: Performance variations with different data sizes under varying computational resources. For optimal performance, all three factors must be scaled up in tandem. Empirical performance exhibits a power-law relationship with each individual factor when it is not constrained by the other two factors.

---

#### Algorithm 1: the power-law function

---

**input** : x-axis data for data volume, model size or compute time  $X$ , word error rate  $E$ .  
**output**:  $a_0, a_1$  are the coefficients of the power law function  $E(\cdot) = (a_0 * X)^{a_1}$ ,  $v$  is used to determine whether the power law holds.

- 1  $X' \leftarrow \log X, E' \leftarrow \log E$ ;
- 2 define  $LineFunc(X', E') = k * X' + b$ ;
- 3 **for**  $i \leftarrow 1$  to  $t - 1$  **do**
- 4     Use the first  $t-1$  points to fit the straight line equation  $LineFunc(X', E')$  and obtain the coefficients,  $k$  and  $b$ .
- 5 **end**
- 6 // Replace  $(X', E')$  in the straight line formula  $LineFunc$  with  $(X, E)$  to obtain the coefficients  $(a_0, a_1)$  of the power law function  $E(\cdot) = (a_0 * X)^{a_1}$ .
- 7  $(a_0, a_1) \leftarrow \log E = k * \log X + b$
- 8 // Verify that  $(X_t, E_t)$  is on the equation of the power law function  $E(\cdot) = (a_0 * X)^{a_1}$ .
- 9  $E_t^{pred} \leftarrow (a_0 * X_t)^{a_1}$ ;
- 10  $dev \leftarrow E_t^{pred} - E_t$ ;
- 11 **if**  $dev < 0.1$  **then**  $v \leftarrow 1$ ;
- 12 **else**  $v \leftarrow 0$ ;

---

#### 4.1.1 The power law of model when data is fixed.

- **Scaling TrOCR Models** We trained 4 different scales (ranging in size from 43.09M to 1B) of TrOCR models on MJ and ST. The experimental results on 6 common test benchmarks are shown in Table 4. As shown in Fig 2a,

our analysis reveals a linear relationship on the log-log plot between the parameter count  $N$  and modeling performance. This relationship can be described by a power-law equation ( $E = aC^b$ ). Employing Algorithm 1, we utilized the first three models (TrOCR-S, TrOCR-B and TrOCR-L) to obtain the power function equation  $E(\cdot)$ . The last model (TrOCR-H) accurately aligns with the fitted straight line, demonstrating the effectiveness of the power law. The power law of the TrOCR model is as follows.

$$E(N) = (1.97 * 10^4 / N)^{0.223} \quad (1)$$

- **Scaling PARSeq Models** To further validate the power law in relation to model parameters, we trained PARSeq models across 4 different scales with sizes ranging from 22M to 0.6B parameters, using the **REBU-Syn** dataset. The results of these experiments on 6 common test benchmarks are detailed in Table 5. As shown in Fig 3, the PARSeq demonstrates a similar trend to that observed with TrOCR, reinforcing the existence of the power law on model size. The power law of the PARSeq model is as follows.

$$E(N) = (6.316 * 10^{-74} / N)^{0.018} \quad (2)$$

Model	Avg	Regular Text			Irregular Text		
		IC13	IIIT5k	SVT	CUTE80	IC15	SVTP
PARSeq-S	96.85	857	3,000	647	288	1,811	645
PARSeq-B	98.72	98.63	98.45	99.65	92.27	95.97	
PARSeq-L	99.07	98.53	98.76	99.31	92.44	96.74	
PARSeq-H	99.30	98.63	98.61	99.31	92.32	97.21	
PARSeq-H	97.06	99.11	98.93	98.45	99.65	91.66	97.67

Table 5. Word accuracy with different model size of PARSeq. Train data: REBU-Syn.




Figure 3. The average word error rate on 6 common test benchmarks was calculated using the PARSeq model size. The solid line represents the fitted power law  $E(\cdot)$ , and the points on the dotted line correspond to the power law equation.

#### 4.1.2 The power-law of data when model size is fixed.

- **Scaling Data Volume on TrOCR** In order to explore the impact of data volume on model performance. We use MJ+ST and  $MJST^+$  for training TrOCR-B. We randomly sampled data at various scales, with sizes ranging from 0.75M to 75M. The experimental results of TrOCR-B based on data of different scales are compiled in Table 6. We used various levels of data volume (solid blue line) to fit the power function (Eq. 3) as shown by the solid grey line in Fig. 2b. The remaining portion of the data size (represented by the dashed blue line) still closely follows the power function, providing further evidence that the data volume adheres to the power function.

$$E(D) = (1.84 * 10^5 / D)^{-0.3271} \quad (3)$$

Data Volume	Volume	Avg	Regular Text			Irregular Text		
			IC13 857	IIIT5k 3,000	SVT 647	CUTE80 288	IC15 1,811	SVTP 645
5%	0.75M	50.35	64.53	57.77	47.60	29.51	39.98	38.14
10%	1.50M	52.61	64.18	59.43	55.63	33.68	42.13	42.33
25%	3.75M	74.86	86.11	79.30	80.22	60.07	63.83	71.47
50%	7.50M	76.47	86.35	79.73	80.83	60.07	64.88	72.40
100%	15.00M	88.56	96.14	92.00	91.56	80.56	81.14	83.91
500%	75.00M	93.09	97.32	93.51	96.33	89.93	86.47	91.47

Table 6. TrOCR-B average accuracy in different percent of training data.

- **Scaling Data Volume on PARSeq** Based on the power law of data volume, we utilize REBU-Syn in ParSeq-S training. By gradually expanding the data samples, the accuracy of PARSeq-S has been significantly improved in the Table 7.

Data	Volume	Avg	Regular Text			Irregular Text		
			IC13 857	IIIT5k 3,000	SVT 647	CUTE80 288	IC15 1,811	SVTP 645
R	3.30M	95.57	97.32	97.87	97.37	97.22	90.34	94.73
R+B+E	3.32M	95.63	97.43	97.97	97.84	98.96	90.28	93.64
R+B+E+U	6.42M	96.12	99.53	97.93	97.53	98.96	91.39	95.66
R+B+E+U+MJST	20.82M	96.45	98.48	98.40	97.84	98.61	91.44	96.43
R+B+E+U+MJST+Syn	23.82M	96.85	98.93	98.63	98.61	99.31	92.32	97.21

Table 7. PARSeq-S average accuracy in different percent of training data.




Figure 4. Average word error rate on 6 common test benchmarks, with respect to images seen (batch size times number of steps) during PARSeq training stage of different sizes.

#### 4.1.3 The power law of computation

With the power laws of model size and data volume separately, we infer that the error rate and the compute budget can also be fit with the power law. We perform the study on TrOCR model. The outcome is depicted as the gray line on the plot on the right-hand side of Fig. 2c. It can be fitted with the power formula as in Eq. 4.

$$E(C) = (4.45 * 10^4 / C)^{-0.3271} \quad (4)$$

#### 4.2. Other Observations

**Large-scale models make better use of samples.** As we continue to enlarge the model size, the model accuracy improves consistently. The phenomenon can be observed in Table 4 and 5. To improve training results, we can modify recognition models built on the Transformer architecture by utilizing the scaling laws of the vision Transformer. As shown in Figure 4 with respect to the total number of images “seen” during PARSeq training stage of different sizes, it is clear that larger models utilize samples more effectively than their smaller models. When PARSeq models of different sizes are trained with an equal number of samples, smaller models exhibit a higher error rate compared to larger models. Furthermore, we observed that larger models tend to require fewer epochs to converge. For instance, PARSeq-S reached optimal accuracy in 32 epochs, whereas PARSeq-B needed only 14 epochs, and PARSeq-L just 5 epochs. These findings suggest that with adequate training resources, it is more beneficial to train a larger model for fewer steps. This mirrors similar findings in language modelling [31] and machine translation [22]. However, when training time is a limiting factor, opting for a smaller model may be more practical.

**The proportion of training data from various sources is crucial for model training.** The REBU-Syn comprises both real and synthetic data. According to prior studies [8, 96], real data typically outperforms synthetic data in

training efficiency, though synthetic data still plays a valuable role. Due to the high costs associated with obtaining and labeling real data, which often do not meet the volume requirements for model training, reliance on synthetic data is necessary. However, the effectiveness of synthetic data raises a question: Does more synthetic data always equate to better performance? Our findings suggest that an optimal ratio between real and synthetic data is crucial for enhanced model performance.

To achieve this objective, we conducted an experiment to investigate the proportional relationship between data obtained from various sources and determine the most efficient utilization of synthetic data. Synthetic data can be primarily categorized into two types: MJ+ST and Syn (CST+SA). MJ+ST, characterized by its large volume but homogeneous nature (consisting mostly of straight and clear samples), contrasts with SynText, which has a smaller volume (only one-fifth of MJ+ST) and primarily consists of curved texts. To evaluate the impact of different synthetic data sources on model accuracy, we trained PARSeq using a combination of real data and these synthetic datasets. The results, as presented in Table 8, are revealing. The accuracy achieved using real data combined with MJ+ST is 96.24%, only marginally higher by 0.05% compared to using real data with Syn. Given that the volume of MJ+ST is five times that of Syn, this implies that complex synthetic data is more sample-efficient. By simultaneously utilizing synthetic data from both MJ+ST and SynText along with real data, we observed a substantial enhancement in the accuracy of PARSeq, elevating it to state-of-the-art levels. This combination of diverse synthetic data styles, when integrated with real data, expands the range of the training data distribution. Such comprehensive coverage effectively enhances the overall quality and performance of the model.

Real DataSet	Syn DataSet	Data Ratio	Word Acc
R+E+B+U	Syn	1:0.5	96.19
R+E+B+U	MJ+ST	1:2.5	96.24
R+E+B+U	MJ+ST+Syn	1:3	<b>96.85</b>

Table 8. PARSeq-S average accuracy of integrating diverse synthetic and real data types.

Additionally, we investigated the effect of different synthetic-to-real data ratios on the accuracy of PARSeq-S. We maintained a constant amount of real data and progressively increased the volume of synthetic data. The ratio of synthetic to real data varied from 0.5 to 5 times. These varying proportions were achieved through random sampling. To augment the total volume of synthetic data, we randomly selected 18M samples from  $MJST^+$  and combined them with the synthetic data in REBU-Syn, culminating in a total of 36M synthetic data samples.

While synthetic data proves effective, it requires careful balancing with real data. As shown in Table 9, a gradual increase in synthetic data led to some improvement in

Data Ratio	Word Acc
Real:Syn=1:0.5	96.32
Real:Syn=1:1	96.50
Real:Syn=1:2	96.59
Real:Syn=1:3	<b>96.85</b>
Real:Syn=1:4	96.76
Real:Syn=1:5	95.70

Table 9. PARSeq-S average accuracy on 6 common test benchmarks with varying ratios of synthetic and real data.

accuracy. Notably, the highest accuracy of 96.85% was achieved with a synthetic-to-real data ratio of 1:3. Beyond this ratio, the accuracy began to decline, likely due to the data distribution becoming overly skewed towards synthetic data, which can adversely affect model performance. Therefore, we recommend a synthetic-to-real data ratio of 1:3. This balance offers the most significant improvement in accuracy without incurring excessive training costs.

Pretrain	Dataset	Backbone	Word Acc
Scratch	R+E+B+U	vit-s	96.12
Scratch	R+E+B+U+Syn	vit-s	96.85
R+E+B+U+Syn	R+E+B+U	vit-s	<b>96.96</b>
Scratch	R+E+B+U+Syn	vit-L	97.03
ImageNet-21k	R+E+B+U+Syn	vit-L	96.74

Table 10. Average accuracy achieved by using visual task pre-training and OCR task pre-training on 6 common test benchmarks.

**Task-related pre-trained models are more effective.** The utility of pretrained models in low-level vision tasks is well-known, but their applicability in OCR tasks warrants investigation. To address this, we experimented with various pretrained models, some trained on ImageNet and others specifically for OCR recognition tasks. In the last two rows of Table 10, we maintained consistent training schedules, learning rates and epochs as used for PARSeq. Intriguingly, the ImageNet-21k pre-trained models underperformed compared to those trained from scratch, a trend observed in both PARSeq and CLIP4STR models. This suggests that pretraining on non-OCR-specific tasks might not be beneficial, and can even be detrimental to OCR performance. The OCR task necessitates a connection between visual and textual elements, akin to the CLIP experiment’s original purpose, whereas purely visual tasks focus more on high-level semantics and lack the textual nuances critical for OCR.

Additionally, when we trained PARSeq-S using the REBU-Syn dataset, it achieved a higher accuracy of 96.85% compared to training solely on the real data REBU. Further fine-tuning the 96.85% model with REBU led to an increased accuracy of 97.01%, indicating an improvement. This demonstrates the efficacy of task-related pretrained models in OCR tasks. To attain higher accuracy, a recommended approach is training on all data first and then fine-tune on real data.

Method	Training data	Regular Text			Irregular Text			Avg
		IC13 857	IIT5k 3,000	SVT 647	CUTE80 288	IC15 1,811	SVTP 645	
ViTSTR-B [2]	MJ+ST	92.40	88.40	87.70	81.30	72.60	81.80	85.46
SE-ASTER [60]	MJ+ST	92.80	93.80	89.60	83.60	80.0	81.40	88.35
TRBA [4]	MJ+ST	93.4	92.10	88.90	89.20	77.4	89.30	90.07
SRN [90]	MJ+ST	95.5	94.80	91.50	87.80	82.70	85.10	90.42
TextScanner [81]	MJ+ST	94.90	95.70	92.70	91.60	83.50	84.80	91.15
VisionLAN [86]	MJ+ST	95.7	95.80	91.70	88.50	83.7	86.00	91.23
PETR [87]	MJ+ST	97.00	95.80	92.40	89.90	83.30	86.20	91.42
ABINet [21]	MJ+ST	97.4	96.20	93.50	89.20	86.00	89.30	92.66
MaskOCR(ViT-B) [51]	MJ+ST	98.1	95.8	94.7	89.2	87.3	89.9	93.1
PARSeq <sub>A</sub> [8]	MJ+ST	96.20	97.00	93.60	92.20	82.90	88.90	93.16
MATR <sub>N</sub> [55]	MJ+ST	95.80	96.60	95.00	93.50	82.80	90.60	93.20
TrOCR <sub>Large</sub> [42]	MJ+ST+B	97.30	94.10	96.10	95.10	84.10	93.00	93.23
DiG-ViT-B [88]	MJ+ST	96.90	96.70	94.60	91.30	87.10	91.00	93.41
MaskOCR(ViT-L) [51]	MJ+ST	98.2	98.0	96.9	95.8	90.1	94.6	93.78
ViTSTR-S [2]	Real	97.80	97.90	96.00	96.20	89.00	91.50	94.85
DiG-ViT-B [88]	Real	97.60	97.60	96.50	96.50	88.90	92.90	94.86
PARSeq <sub>A</sub> <sup>‡</sup> [8]	Real	97.32	97.87	97.37	97.22	90.34	94.73	95.57
ABINet [21]	Real	98.00	98.60	98.20	97.20	90.50	94.10	96.01
MAERec-B [30]	Union14M-L	98.10	98.50	97.80	98.60	90.7	94.40	96.15
CLIP4STR-B [96]	Real	98.36	98.73	97.68	98.96	91.39	96.74	96.89
CLIP4STR-L [96]	Real	98.48	<b>99.43</b>	98.15	98.96	91.66	97.36	97.06
<b>CLIP4STR-B*</b>	REBU-Syn	99.29	98.96	<b>98.76</b>	99.65	92.27	97.83	97.25
<b>CLIP4STR-L*</b>	REBU-Syn	<b>99.42</b>	99.13	98.61	<b>99.65</b>	<b>92.6</b>	<b>98.13</b>	<b>97.42</b>

Table 11. Word accuracy on 6 common test benchmarks, \* indicates experimental results based on scaling law, Avg is the weighted average result on 6 common test benchmarks. ‡ indicates reproduced by us.

### 4.3. Comparison with SOTA Methods

Recently, the remarkable performance of CLIP4STR across multiple benchmarks prompted us to conduct further experiments, guided by our scaling law. Initially, we focused on data composition, employing a 3:1 ratio of synthetic to real data for training the model, in conjunction with fine-tuning using a pre-trained model for related tasks. Our reproducible results led to a notable improvement in CLIP4STR-B, enhancing its accuracy from 96.54% to 97.25%, an increase of 0.65%. This achievement represents the best result to date in the text recognition task.

To further delve into the impact of larger models, we replicated this experiment on CLIP4STR-L. This model achieved a new state-of-the-art, recording a top-1 average accuracy of 97.42% on 6 common test benchmarks, as detailed in Table 9. These findings underscore the significant role of large models in advancing the field of OCR.

## 5. Discussion and Conclusion

In this paper, we establish the presence of smooth power laws within the OCR field, demonstrating that model performance predictably improves with sufficient increases in model size, data volume and computational resources. Additionally, we identified several key principles crucial for ef-

fective model training in OCR: 1) Large-scale models more efficiently utilize samples. 2) The proportion of training data from various sources is crucial for model training; 3) Task-related pre-trained models enhance effectiveness. Beyond identifying these guiding principles, we compiled a large-scale dataset to improve the performance of the OCR model. Leveraging these rules, we successfully trained a model that achieved a new state-of-the-art average accuracy of 97.42% on the test benchmark.

We conduct extensive experiments on both model scaling and data scaling, successfully demonstrating the existence of scaling laws in OCR. Furthermore, we observe that data scaling is a particularly advantageous approach as it enhances model accuracy without incurring additional costs during training or inference. However, challenges persist in the realm of model scaling. While large-scale models exhibit superior performance with substantial data, their training is considerably more costly. Adjusting each parameter can be extremely expensive, with each training iteration potentially costing millions of dollars. To optimize the performance of these models, careful selection of the best hyperparameters during training is essential. We hope that our work can attract more researchers' attention to reduce the training cost of large-scale models.



# An Empirical Study of Scaling Law for OCR

## Supplementary Material

### A. More Experiment Analysis

#### A.1. The impact of model training details

Regarding how to train the optimal model, we conduct relevant research on multiple dimensions such as batch size and depth used during training.

**BatchSize** we focused on examining the impact of various batch sizes on the accuracy of the PARSeq-B model. This investigation was integral to determining the model’s optimal training conditions. The findings, as presented in Table 1, reveal that the model reaches its optimal performance, with an accuracy of 96.35%, when the batch size is configured to 1024. This result corroborates the conclusions from the CLIP4STR[96]. It is underscoring the significant role that larger batch sizes play in enhancing model accuracy. Notably, it was also observed that an excessively large batch size leads to a reduction in accuracy, indicating a critical balance in batch size selection for optimal model training.

Model	Backbone	Batch	Word Acc
PARSeq	ViT-B	1344	96.28
PARSeq	ViT-B	1024	<b>96.35</b>
PARSeq	ViT-B	896	96.33
PARSeq	ViT-B	448	96.3

Table 1. Average accuracy using different batch size on common benchmarks, training model in real dataset.

**Depth** PARSeq is divided into encoder and decoder. The encoder leverages the widely-recognized Vision Transformer (ViT) series, specifically employing the ViT-S variant. Conversely, the decoder is subject to more intricate fine-tuning, particularly concerning its depth. This aspect of the model architecture was a focal point of our research.

Our empirical investigations, as detailed in Table 2, centered on the interplay between the encoder’s ViT-S configuration and varying depths of the decoder. The experimental findings were revealing. With the encoder consistently utilizing ViT-S, we observed that setting the decoder’s depth to 1 layers resulted in optimal model accuracy. This suggests a significant relationship between decoder depth and model performance, underlining the importance of carefully calibrated model architecture in achieving high OCR accuracy. Our results contribute to a deeper understanding of the architectural nuances in Transformer-based OCR models and their impact on performance.

#### A.2. Benefits of pretraining in different languages

In this supplementary section, we conduct a thorough examination of the impact of language-specific pretraining on

Model	Encoder	Decoder-Depth	Word Acc
PARSeq	ViT-S	1	<b>95.56</b>
PARSeq	ViT-S	2	95.31
PARSeq	ViT-S	3	94.77
PARSeq	ViT-S	4	94.50
PARSeq	ViT-S	5	93.77

Table 2. Average accuracy using different depth for decoder on benchmark test set, training model in real dataset.

OCR models, with a particular focus on fine-tuning for English datasets. Our approach involved utilizing models pretrained in Arabic, Latin, and a hybrid of Chinese-English, each trained on a dataset comprising 300,000 entries drawn from private sources. The core architecture for these models is based on the CMT-S framework, as detailed in Guo et al. (2022) [24]. Subsequent secondary training was conducted on the REB dataset, utilizing REBU-Syn, wherein different language-specific pretrained models were employed. Notably, the final classification layer’s parameters were not loaded from these pretrained models to ensure a fair comparison.

As illustrated in Table 4, our results reveal pronounced improvements in models pretrained in Latin, Chinese, and English, with Latin demonstrating the most substantial enhancement. This improvement is likely due to the visual congruence between Latin and English scripts, emphasizing the OCR models’ dependency on visual features for effective recognition. Meanwhile, the performance of models pretrained in Chinese and English, though slightly lower by a margin of 0.01% compared to the Latin model, indicates a potential bias introduced by the inclusion of Chinese data in the pretraining phase.

Intriguingly, models pretrained in Arabic did not exhibit significant benefits over their non-pretrained counterparts. This can be attributed to the stark visual differences between Arabic and English scripts, reinforcing the notion that visual similarity plays a crucial role in the efficacy of pretraining for OCR tasks. Collectively, these findings suggest that pretraining OCR models with languages visually akin to the target language offers enhanced benefits. Conversely, a pronounced visual dissimilarity between the scripts negates the advantages of pretraining, a critical consideration for the development of effective OCR systems.

### B. Comparisons on Union14M benchmark

To evaluate the generalization capabilities of our model, we conducted an extensive assessment using the Union14M benchmark dataset [30]. This benchmark is particularly comprehensive, encompassing a vast array of real-world

Method	Training data	Artistic	Contextless	Curve	General	Multi-Oriented	Multi-Words	Salient	Avg
CRNN [69]	MJ+ST	20.7	25.6	7.5	32.0	0.9	25.6	13.9	18.0
SVTR [20]	MJ+ST	37.9	44.2	63.0	52.8	32.1	49.1	67.5	49.5
MORAN [50]	MJ+ST	29.4	20.7	8.9	35.2	0.7	23.8	17.9	19.5
ASTER [71]	MJ+ST	27.7	33.0	34.0	39.8	10.2	27.6	48.2	31.5
NRTR [68]	MJ+ST	36.6	37.3	31.7	48.0	4.4	54.9	30.6	34.8
SAR [41]	MJ+ST	42.6	44.2	44.3	50.5	7.7	51.2	44.0	40.6
DAN [84]	MJ+ST	35.0	40.3	26.7	42.1	1.5	42.2	36.5	32.0
SATRN [38]	MJ+ST	48.0	45.3	51.1	58.5	15.8	52.5	62.7	47.7
RobustScanner [92]	MJ+ST	41.2	42.6	43.6	39.5	7.9	46.9	44.9	38.1
SRN [90]	MJ+ST	34.1	28.7	63.4	46.3	25.3	26.7	56.5	40.1
ABINet [21]	MJ+ST	43.3	38.3	59.5	55.6	12.7	50.8	62.0	46.0
VisionLAN [86]	MJ+ST	47.8	48.0	57.7	52.1	14.2	47.9	64.0	47.4
MATRN [55]	MJ+ST	43.8	41.9	63.1	57.0	13.4	53.2	66.4	48.4
CRNN [69]	Union14M	31.9	39.3	18.9	58.1	4.3	21.5	15.1	27.0
SVTR [20]	Union14M	50.2	63.0	70.5	74.7	66.6	42.6	71.4	62.7
MORAN [50]	Union14M	44.3	51.1	42.4	42.9	12.4	36.8	41.0	38.7
ASTER [71]	Union14M	39.2	47.9	37.4	64.4	12.5	34.5	30.2	38.0
NRTR [68]	Union14M	51.8	65.1	47.9	72.9	39.1	51.4	40.1	52.6
SAR [41]	Union14M	58.0	69.0	66.9	73.7	54.7	51.2	57.0	61.5
DAN [84]	Union14M	47.0	56.6	44.6	66.7	22.1	39.8	41.5	45.5
SATRN [38]	Union14M	64.3	71.1	73.0	78.8	64.7	47.4	69.2	66.9
RobustScanner [92]	Union14M	58.7	72.7	64.2	73.5	52.8	47.8	56.9	60.9
SRN [90]	Union14M	47.6	57.9	48.7	60.7	20.0	27.9	41.6	43.5
ABINet [21]	Union14M	62.2	66.3	73.0	75.6	59.6	43.1	69.5	64.2
VisionLAN [86]	Union14M	54.4	60.1	68.8	72.1	55.2	37.9	64.7	59.0
MATRN [55]	Union14M	67.3	71.0	79.3	78.4	66.0	53.8	74.9	70.0
MAERec-S [30]	Union14M-L	68.9	77.8	79.3	80.4	69.5	51.9	75.1	71.8
MAERec-B [30]	Union14M-L	75.9	80.7	86.6	83.8	82.1	56.2	82.2	78.2
CLIP4STR-B [96]	R	86.5	<b>92.2</b>	96.3	89.9	96.1	88.9	91.2	91.6
CLIP4STR-L [96]	R	87.2	91.0	<b>97.0</b>	<b>90.3</b>	96.6	89.9	91.5	91.9
<b>CLIP4STR-B*</b>	REBU-Syn	88.6	90.1	96.4	89.1	96.3	<b>92.2</b>	91.9	92.1
<b>CLIP4STR-L*</b>	REBU-Syn	<b>88.6</b>	90.4	96.4	89.3	<b>97.2</b>	90.7	<b>92.7</b>	<b>92.2</b>

Table 3. Word accuracy on Union14M benchmark, \* indicates experimental results based on scaling law.

Pretrain	Model	Datasets	Word Acc
From Scratch	PARSeq	REB	95.60
Arabic	PARSeq	REB	95.62
Cn-En	PARSeq	REB	95.81
Latin	PARSeq	REB	<b>95.82</b>

Table 4. Average accuracy using language-specific pretraining on benchmark test set, training model in real dataset of REB.

textual data, systematically categorized into seven distinct subsets: Artistic, Contextless, Curve, General, Multi-Oriented, Multi-Words and Salient. The results of this evaluation, presented in Table 3, demonstrate the model’s robust and consistent performance across a range of scenarios. Notably, In comparative evaluations against standard benchmarks and the multifaceted Union14M dataset, the CLIP4STR-L\* model emerges as a standout performer. This model demonstrates exceptional accuracy across the majority of datasets. Its ability to consistently deliver high-quality results, particularly in the context of the challenging Union14M benchmark, underscores its robustness and versatility. Such performance highlights the efficacy of the CLIP4STR-L\* architecture in handling a diverse range of

textual data scenarios, making it a benchmark in the field.

### C. Visualization Analysis

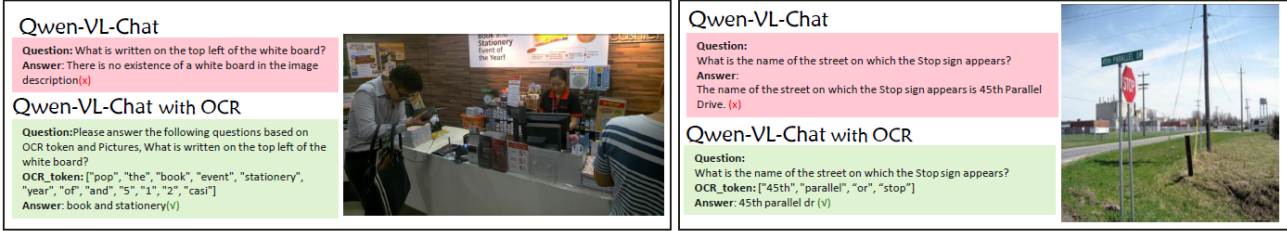
In Fig 1, we present a visualization of our model’s performance across the seven major categories of the Union14M benchmark. The results demonstrate that our model outperforms in the majority of datasets. However, a slight dip in effectiveness is noted in the Contextless dataset. This can be attributed to the limitations of the text encoder in processing texts lacking semantic information.

Despite this, our model distinguishes itself from other contemporary OCR systems through its enhanced ability to accurately interpret and navigate a diverse range of complex real-world scenarios. This advancement significantly bolsters the robustness of OCR models, enabling them to operate with greater reliability in varied and challenging environments. The enhanced robustness of our model not only showcases its technical excellence but also emphasizes its practical applicability in real-world settings characterized by high variability and complexity.

Artist	GT	cup	finert	dellarte	cupcake	style	arrow	imagine	howardjohnson
	CLIP4STR-B	curp	finerit	deliarte	cupcale	stule	annow	magine	howardjohnson
	CLIP4STR-B*	cup	finert	dellarte	cupcake	style	arrow	magine	howardjohnson
	CLIP4STR-L*	cup	finert	dellarte	cupcake	style	arrow	imagine	howardjohnson
Contextless									
	GT	ice2or0665	xingzhengxue	xinfengru	amarsex	iapp	beiteru	3333363	md04aca4004tb
	CLIP4STR-B	ice2or0665	xingzhengxuf	ainfengru	amalsex	iapp	betteru	3333363	mdo4aca4004tb
	CLIP4STR-B*	ice2qr0665	xingzhengxue	ainfengku	amalsex	lapp	beiteru	3333363	md04aca4004tb
CLIP4STR-L*	ice2qr0665	xingzhengxue	xinfengru	amarsex	lapp	beiteru	3333363	md04aca4004tb	
Curve									
	GT	smile	wood	ourense	chez	candidate	boudin	eexpert	pendergrass
	CLIP4STR-B	smiley	woods	durense	2002	canddate	buudin	expert	pendercrass
	CLIP4STR-B*	smiley	woods	durense	2002	canddate	buudin	expert	pendercrass
CLIP4STR-L*	smile	wood	ourense	chez	candidate	boudin	eexpert	pendergrass	
Multi-Oriented									
	GT	kannstmitheffen	nationaux	tian	defnsio	crowd	la	imagine	woo
	CLIP4STR-B	jannstmitheffen	nationalix	man	deensio	around	le	magine	com
	CLIP4STR-B*	fannstmitheffen	nationaux	tian	defnsio	crowd	la	imagine	woo
CLIP4STR-L*	kannstmitheffen	nationaux	tian	defnsio	crowd	la	imagine	woo	
Multi-Words									
	GT	itsreallygood	musicstudio	communityservice	semiskimmed	shanghongpharmacy	officebuilding	magicbox	
	CLIP4STR-B	ztsreallygood	musiostudio	communityservice	semiskimmer	siangbongpharmacy	officebuilleling	maglcbx	
	CLIP4STR-B*	ztsreallygood	nationaux	communityservice	semiskimmed	shanghongpharmacy	officebuilding	magicbox	
CLIP4STR-L*	itsreallygood	musicstudio	communityservice	semiskimmed	shanghongpharmacy	officebuilding	magicbox		
Salient									
	GT	abolish	ceremonies	cceso	xploratory	chocolat	chick	complainin	sterilgard
	CLIP4STR-B	abglish	geremonies	accesso	exploratory	chocolate	chickli	complaining	sterifgard
	CLIP4STR-B*	abolish	ceremonies	cceso	exploratory	chocolate	chickli	complaining	sterilgard
CLIP4STR-L*	abolish	ceremonies	cceso	xploratory	chocolat	chick	complainin	sterilgard	
General									
	GT	taivan	movilida	delville	alone	10	snickers	pittsdurgh	61091679taobaocon
	CLIP4STR-B	talvan	movilidad	belville	abme	to	snicklies	pittsburgh	61091679farbapcom
	CLIP4STR-B*	taivan	movilidad	belville	abre	10	snickers	pittsburgh	61091679taobaycon
CLIP4STR-L*	taivan	movilida	delville	alone	10	snickers	pittsdurgh	61091679taobaocon	

Figure 1. Error analysis of the Union14M benchmark. We select three representative models and show their prediction results (Text in black represents correct prediction and red text vice versa).

## STVQA



## TextVQA




Figure 2. Visual answer comparison for Qwen-VL-Chat with or without OCR in natural scenes VQA.

## D. OCR Enhanced LMM

In the realm of large-scale models, we observe a distinct bifurcation into two primary categories: Large Language Models (LLMs) and Large Multimodal Models (LMMs). It is crucial to acknowledge that while LLMs are devoid of a visual component, LMMs’ visual branches demonstrate room for enhancement in terms of OCR capabilities[73]. This observation underscores the relative underdevelopment of OCR proficiency within large-scale models. Optical Character Recognition (OCR) tasks, however, offer a promising avenue to address this shortfall, thereby motivating our investigation into the benefits of integrating OCR with these models.

**Dataset and Metric** Our analysis utilized a diverse range of tasks from the Visual Question Answering (VQA) series, specifically STVQA [9], TextVQA [75], DocVQA [52] and InfoVQA [53]. While STVQA and TextVQA are geared towards natural scenes, DocVQA and InfoVQA focus on general document contexts. Here are some details of evaluation dataset:

- **STVQA** contains 31K questions that require understanding the scene text, based on 23K images from : IC-DAR2013 and ICDAR2015, ImageNet [18], VizWiz [26], IIIT Scene Text Retrieval, Visual Genome [34] and COCO-Text.
- **TextVQA** contains 45K questions that need to read and reason the text in images, based on 28K images from natural images.
- **DocVQA** contains 50K questions and 12K images from industry documents.
- **InfoVQA** contains 30K questions that require understanding the document text, based on 5.4K images combining textual, graphical and visual elements from InfoGraphics.

We employed the Average Normalized Levenshtein Similarity (ANLS) as our evaluation metric, a standard in the VQA domain.

**Experiment Setting** For the large-scale model, we selected the recently unveiled Qwen-VL-chat [5], a state-of-the-art multimodal model. In terms of OCR, we utilized Rosetta [10] for detection, and CLIP4STR-L\* for recognition. We began by concatenating the text recognized through coordinate information to generate OCR tokens. These tokens, combined with the question, formed our prompts. The prompt format was meticulously refined to: 'OCR token: {*ocrtokens*}, please answer the following questions based on OCR token and pictures, {*question*}'. This approach involved inputting both the prompt and images into the large-scale model.

Model	STVQA	TextVQA	DocVQA	InfoVQA
BLIP-2 [1]	21.7	32.2	4.9	-
LLaVAR [45]	39.2	48.5	11.6	-
InstructBLIP [44]	-	50.7	-	38.3
LLaMA-7B [89]	-	52.6	62.2	38.2
Pix2Struct-base [39]	-	-	72.1	38.2
Qwen-VL-Chat	50.25	61.5	63.41	31.7
Qwen-VL-Chat with OCR	<b>70.32</b>	<b>69.64</b>	<b>73.44</b>	<b>38.48</b>

Table 5. Result on benchmarks for VQA tasks using LMM models with or without OCR, all result are ANLS on the val split.

**Result and Analyze** We performed a detailed comparative analysis to assess the accuracy of the Qwen-VL-chat model, examining its performance with and without OCR integration, as delineated in Table 5. Our results reveal a significant improvement in the accuracy of the model for scene-based VQA tasks upon the integration of OCR. Additionally, there is a noticeable enhancement in document-based VQA tasks. These findings suggest that the incorporation of OCR not only enhances the model’s accuracy but also extends its generalization capabilities across diverse



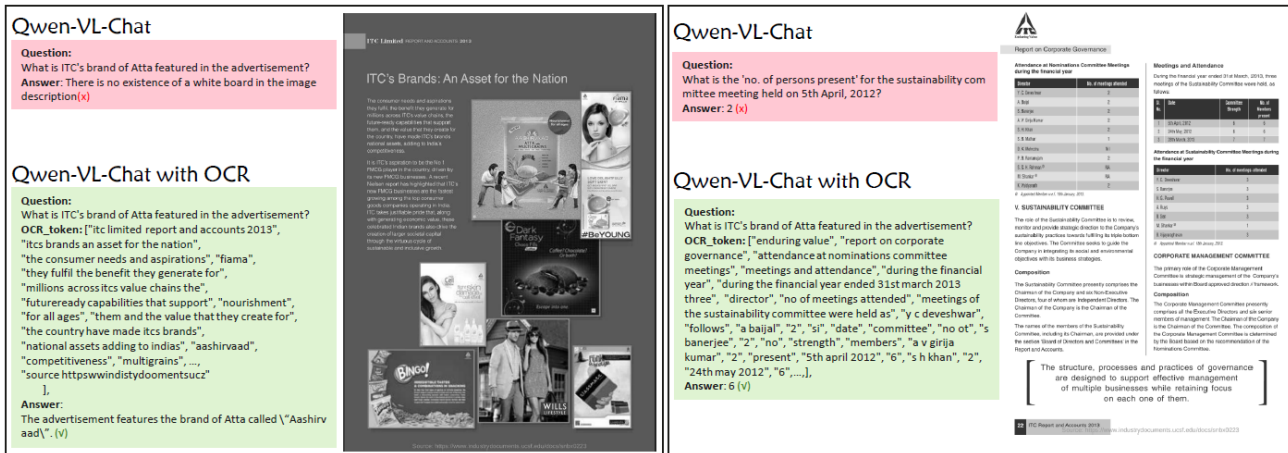


Figure 3. Visual answer comparison for Qwen-VL-Chat with or without OCR in Document VQA.

VQA scenarios. This evidence distinctly highlights the vital role that OCR inputs play in augmenting the performance of LVLM for downstream tasks. Furthermore, the improved accuracy with OCR integration underscores the model’s enhanced ability to interpret and analyze combined visual and textual data, thereby validating the efficacy of multimodal approaches in tackling complex analytical challenges.

**VQA Visualization Analysis** Our visual analysis of Qwen-VL-Chat, with and without the OCR module, across varied datasets offers critical insights. In natural scene Visual Question Answering (VQA) tasks, Qwen-VL-Chat encounters difficulties in detecting small text in images. The upper left corner of Fig 2, the model overlooks pertinent content, erroneously indicating its absence. Moreover, its tendency to inaccurately complete blurred text stems from its sophisticated linguistic abilities. This is evident in the upper right corner of Fig 2, where ‘dr’ in ‘45th parallel dr’ is incorrectly expanded to ‘drive’. Notably, the model’s text misidentification issues, such as converting ‘honghe’ to ‘Hongte’ on a cigarette pack as depicted in the lower left corner of Fig 2 (mistaking the second ‘h’ for a ‘t’), are significantly mitigated by integrating our OCR module.

In general document scenarios involving dense textual information, the performance of Qwen-VL-Chat remains suboptimal. In the left of Fig 3, when tasked with identifying brands in advertisements amidst extensive text, the model erroneously generates non-existent words from the image. Incorporating OCR crucially directs the model towards accurate text recognition. This pattern is consistent in table-based VQA Tasks in the right of Fig 3, where the model frequently errs in its responses. The inclusion of OCR proves instrumental in steering the model towards correct answers. This comprehensive analysis highlights the pivotal role of OCR in augmenting LMM models’ comprehension and recognition capabilities within intricate visual-

textual contexts.

**References**

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022. 4
- [2] Rowel Atienza. Vision transformer for fast and efficient scene text recognition. 2021. 8
- [3] Jeonghun Baek, Geewook Kim, Junyeop Lee, Sungrae Park, Dongyoon Han, Sangdoon Yun, Seong Joon Oh, and Hwal-suk Lee. What is wrong with scene text recognition model comparisons? dataset and model analysis. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4715–4723, 2019. 2
- [4] Jeonghun Baek, Yusuke Matsui, and Kiyoharu Aizawa. What if we only use real datasets for scene text recognition? toward scene text recognition with fewer labels. 2021. 1, 8
- [5] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*, 2023. 4
- [6] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021. 1, 3
- [7] Andrei Barbu, David Mayo, Julian Alverio, William Luo, Christopher Wang, Dan Gutfreund, Josh Tenenbaum, and Boris Katz. Objectnet: A large-scale bias-controlled dataset for pushing the limits of object recognition models. *Advances in neural information processing systems*, 32, 2019. 2
- [8] Darwin Bautista and Rowel Atienza. Scene text recognition with permuted autoregressive sequence models. 2022. 3, 4, 6, 8
- [9] Ali Furkan Biten, Ruben Tito, Andres Mafra, Lluís Gomez,

- Marçal Rusinol, Ernest Valveny, CV Jawahar, and Dimosthenis Karatzas. Scene text visual question answering. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4291–4301, 2019. 4
- [10] Fedor Borisyyuk, Albert Gordo, and Viswanath Sivakumar. 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*, pages 71–79, 2018. 2, 4
- [11] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. 1
- [12] Zhanzhan Cheng, Fan Bai, Yunlu Xu, Gang Zheng, Shiliang Pu, and Shuigeng Zhou. Focusing attention: Towards accurate text recognition in natural images. In *Proceedings of the IEEE international conference on computer vision*, pages 5076–5084, 2017. 2
- [13] Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*, 2019. 2
- [14] Chee Kheng Ch’ng and Chee Seng Chan. Total-text: A comprehensive dataset for scene text detection and recognition. In *2017 14th IAPR international conference on document analysis and recognition (ICDAR)*, pages 935–942. IEEE, 2017. 3
- [15] Chee-Kheng Chng, Yuliang Liu, Yipeng Sun, Chun Chet Ng, Canjie Luo, Zihan Ni, ChuanMing Fang, Shuaitao Zhang, Junyu Han, Errui Ding, Jingtuo Liu, Dimosthenis Karatzas, Chee Seng Chan, and Lianwen Jin. Icdar2019 robust reading challenge on arbitrary-shaped text (rrc-art). 2019. 3
- [16] Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Łukasz Kaiser. Universal transformers. *arXiv preprint arXiv:1807.03819*, 2018. 2
- [17] Mostafa Dehghani, Josip Djolonga, Basil Mustafa, Piotr Padlewski, Jonathan Heek, Justin Gilmer, Andreas Steiner, Mathilde Caron, Robert Geirhos, Ibrahim Alabdulmohsin, et al. Scaling vision transformers to 22 billion parameters. *arXiv preprint arXiv:2302.05442*, 2023. 1
- [18] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 4
- [19] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2
- [20] Yongkun Du, Zhineng Chen, Caiyan Jia, Xiaoting Yin, Tianlun Zheng, Chenxia Li, Yuning Du, and Yu-Gang Jiang. Svtr: Scene text recognition with a single visual model. *arXiv preprint arXiv:2205.00159*, 2022. 2
- [21] Shancheng Fang, Hongtao Xie, Yuxin Wang, Zhendong Mao, and Yongdong Zhang. Read like humans: Autonomous, bidirectional and iterative language modeling for scene text recognition. 2021. 8, 2
- [22] Patrick Fernandes, Behrooz Ghorbani, Xavier Garcia, Markus Freitag, and Orhan Firat. Scaling laws for multi-lingual neural machine translation. 2023. 6
- [23] Behrooz Ghorbani, Orhan Firat, Markus Freitag, Ankur Bapna, Maxim Krikun, Xavier Garcia, Ciprian Chelba, and Colin Cherry. Scaling laws for neural machine translation. *arXiv preprint arXiv:2109.07740*, 2021. 1
- [24] Jianyuan Guo, Kai Han, Han Wu, Yehui Tang, Xinghao Chen, Yunhe Wang, and Chang Xu. Cmt: Convolutional neural networks meet vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12175–12185, 2022. 1
- [25] Ankush Gupta, Andrea Vedaldi, and Andrew Zisserman. Synthetic data for text localisation in natural images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2315–2324, 2016. 3
- [26] Danna Gurari, Qing Li, Abigale J. Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P. Bigham. Vizwiz grand challenge: Answering visual questions from blind people, 2018. 4
- [27] Tom Henighan, Jared Kaplan, Mor Katz, Mark Chen, Christopher Hesse, Jacob Jackson, Heewoo Jun, Tom B Brown, Prafulla Dhariwal, Scott Gray, et al. Scaling laws for autoregressive generative modeling. *arXiv preprint arXiv:2010.14701*, 2020. 2
- [28] Mingxin Huang, Yuliang Liu, Zhenghao Peng, Chongyu Liu, Dahua Lin, Shenggao Zhu, Nicholas Yuan, Kai Ding, and Lianwen Jin. 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*, pages 4593–4603, 2022. 3
- [29] Max Jaderberg, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Synthetic data and artificial neural networks for natural scene text recognition. *arXiv preprint arXiv:1406.2227*, 2014. 3
- [30] Qing Jiang, Jiapeng Wang, Dezhi Peng, Chongyu Liu, and Lianwen Jin. Revisiting scene text recognition: A data perspective. 2023. 8, 1, 2
- [31] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020. 1, 2, 3, 6
- [32] Dimosthenis Karatzas, Faisal Shafait, Seiichi Uchida, Masakazu Iwamura, Lluís Gomez i Bigorda, Sergi Robles Mestre, Joan Mas, David Fernandez Mota, Jon Almazan Almazan, and Lluís Pere De Las Heras. Icdar 2013 robust reading competition. pages 1484–1493, 2013. 3, 4
- [33] Dimosthenis Karatzas, Lluís Gomez-Bigorda, Angelos Nicolaou, Suman Ghosh, Andrew Bagdanov, Masakazu Iwamura, Jiri Matas, Lukas Neumann, Vijay Ramaseshan Chandrasekhar, Shijian Lu, et al. Icdar 2015 competition on robust reading. pages 1156–1160, 2015. 3, 4
- [34] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Fei-Fei Li. Visual genome: Connecting language and vision using crowdsourced dense image annotations, 2016. 4

- [35] Ilya Krylov, Sergei Nosov, and Vladislav Sovrasov. Open images v5 text annotation and yet another mask text spotter. In *Asian Conference on Machine Learning*, pages 379–389. PMLR, 2021. 3
- [36] Taku Kudo and John Richardson. Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing. 2018. 4
- [37] Chen-Yu Lee and Simon Osindero. Recursive recurrent nets with attention modeling for ocr in the wild. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2231–2239, 2016. 2
- [38] Junyeop Lee, Sungrae Park, Jeonghun Baek, Seong Joon Oh, Seonghyeon Kim, and Hwalsuk Lee. On recognizing texts of arbitrary shapes with 2d self-attention. 2019. 2
- [39] Kenton Lee, Mandar Joshi, Iulia Raluca Turc, Hexiang Hu, Fangyu Liu, Julian Martin Eisenschlos, Urvashi Khandelwal, Peter Shaw, Ming-Wei Chang, and Kristina Toutanova. Pix2struct: Screenshot parsing as pretraining for visual language understanding. In *International Conference on Machine Learning*, pages 18893–18912. PMLR, 2023. 4
- [40] Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. Gshard: Scaling giant models with conditional computation and automatic sharding. 2020. 2
- [41] Hui Li, Peng Wang, Chunhua Shen, and Guyu Zhang. Show, attend and read: A simple and strong baseline for irregular text recognition. 2019. 2
- [42] Minghao Li, Tengchao Lv, Jingye Chen, Lei Cui, Yijuan Lu, Dinei Florencio, Cha Zhang, Zhoujun Li, and Furu Wei. Trocr: Transformer-based optical character recognition with pre-trained models. *arXiv preprint arXiv:2109.10282*, 2021. 1, 2, 8
- [43] Xianhang Li, Zeyu Wang, and Cihang Xie. An inverse scaling law for clip training. *arXiv preprint arXiv:2305.07017*, 2023. 2
- [44] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. *arXiv preprint arXiv:2310.03744*, 2023. 4
- [45] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023. 4
- [46] Peter J Liu, Mohammad Saleh, Etienne Pot, Ben Goodrich, Ryan Sepassi, Lukasz Kaiser, and Noam Shazeer. Generating wikipedia by summarizing long sequences. *arXiv preprint arXiv:1801.10198*, 2018. 2
- [47] Wei Liu, Chaofeng Chen, and Kwan-Yee Wong. Char-net: A character-aware neural network for distorted scene text recognition. In *Proceedings of the AAAI conference on artificial intelligence*, 2018. 2
- [48] Xi Liu, Rui Zhang, Yongsheng Zhou, Qianyi Jiang, Qi Song, Nan Li, Kai Zhou, Lei Wang, Dong Wang, Minghui Liao, Mingkun Yang, Xiang Bai, Baoguang Shi, Dimosthenis Karatzas, Shijian Lu, and C. V. Jawahar. Icdar 2019 robust reading challenge on reading chinese text on signboard. 2019. 3
- [49] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019. 3
- [50] Canjie Luo, Lianwen Jin, and Zenghui Sun. A multi-object rectified attention network for scene text recognition. 2019. 2
- [51] Pengyuan Lyu, Chengquan Zhang, Shanshan Liu, Meina Qiao, Yangliu Xu, Liang Wu, Kun Yao, Junyu Han, Er-rui Ding, and Jingdong Wang. Maskocr: text recognition with masked encoder-decoder pretraining. *arXiv preprint arXiv:2206.00311*, 2022. 8
- [52] Minesh Mathew, Dimosthenis Karatzas, and CV Jawahar. Docvqa: A dataset for vqa on document images. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2200–2209, 2021. 4
- [53] Minesh Mathew, Viraj Bagal, Rubèn Tito, Dimosthenis Karatzas, Ernest Valveny, and CV Jawahar. Infographicvqa. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1697–1706, 2022. 4
- [54] Anand Mishra, Karteek Alahari, and CV Jawahar. Scene text recognition using higher order language priors. *BMVC-British machine vision conference*, 2012. 3, 4
- [55] Byeonghu Na, Yoonsik Kim, and Sungrae Park. Multimodal text recognition networks: Interactive enhancements between visual and semantic features. 2022. 8, 2
- [56] Nibal Nayef, Yash Patel, Michal Busta, Pinaki Nath Chowdhury, Dimosthenis Karatzas, Wafa Khelif, Jiri Matas, Umada Pal, Jean-Christophe Burie, Cheng lin Liu, and Jean-Marc Ogier. Icdar2019 robust reading challenge on multilingual scene text detection and recognition – rrc-mlt-2019. 2019. 3
- [57] OpenAI. Chatgpt. <https://openai.com/blog/chatgpt/>, 2023. 2
- [58] OpenAI. Gpt-4 technical report. 2023. 2
- [59] Trung Quy Phan, Palaiahnakote Shivakumara, Shangxuan Tian, and Chew Lim Tan. Recognizing text with perspective distortion in natural scenes. pages 569–576, 2013. 4
- [60] Zhi Qiao, Yu Zhou, Dongbao Yang, Yucan Zhou, and Weiping Wang. Seed: Semantics enhanced encoder-decoder framework for scene text recognition. 2020. 8
- [61] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2
- [62] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551, 2020. 2
- [63] Zobeir Raisi, Mohamed A Naiel, Georges Younes, Steven Wardell, and John S Zelek. Transformer-based text detection in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3162–3171, 2021. 1

- [64] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do imagenet classifiers generalize to imagenet? In *International conference on machine learning*, pages 5389–5400. PMLR, 2019. 2
- [65] Anhar Risnumawan, Palaiahankote Shivakumara, Chee Seng Chan, and Chew Lim Tan. A robust arbitrary text detection system for natural scene images. *Expert Systems with Applications*, 41(18):8027–8048, 2014. 4
- [66] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115:211–252, 2015. 2
- [67] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. 2016. 4
- [68] Fenfen Sheng, Zhineng Chen, and Bo Xu. Nrtr: A no-recurrence sequence-to-sequence model for scene text recognition. In *2019 International conference on document analysis and recognition (ICDAR)*, pages 781–786. IEEE, 2019. 1, 2
- [69] Baoguang Shi, Xiang Bai, and Cong Yao. An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. 2015. 2
- [70] Baoguang Shi, Xiang Bai, and Cong Yao. An end-to-end trainable neural network for image-based sequence recognition and its application to scene text recognition. *IEEE transactions on pattern analysis and machine intelligence*, 39(11):2298–2304, 2016. 1, 2
- [71] Baoguang Shi, Mingkun Yang, Xinggang Wang, Pengyuan Lyu, Cong Yao, and Xiang Bai. Aster: An attentional scene text recognizer with flexible rectification. *IEEE transactions on pattern analysis and machine intelligence*, 41(9):2035–2048, 2018. 2
- [72] Baoguang Shi, Cong Yao, Minghui Liao, Mingkun Yang, Pei Xu, Linyan Cui, Serge Belongie, Shijian Lu, and Xiang Bai. Icdar2017 competition on reading chinese text in the wild (rctw-17). 2018. 3
- [73] Yongxin Shi, Dezhi Peng, Wenhui Liao, Zening Lin, Xinhong Chen, Chongyu Liu, Yuyi Zhang, and Lianwen Jin. Exploring ocr capabilities of gpt-4v (ision): A quantitative and in-depth evaluation. *arXiv preprint arXiv:2310.16809*, 2023. 4
- [74] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 2
- [75] Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8317–8326, 2019. 4
- [76] Amanpreet Singh, Guan Pang, Mandy Toh, Jing Huang, Wojciech Galuba, and Tal Hassner. Textocr: Towards large-scale end-to-end reasoning for arbitrary-shaped scene text. 2021. 3
- [77] Yipeng Sun, Jiaming Liu, Wei Liu, Junyu Han, Errui Ding, and Jingtuo Liu. Chinese street view text: Large-scale chinese text reading with partially supervised learning. 2020. 3
- [78] Yew Lee Tan, Adams Wai-Kin Kong, and Jung-Jae Kim. Pure transformer with integrated experts for scene text recognition. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVIII*, pages 481–497. Springer, 2022. 1
- [79] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. pages 10347–10357, 2021. 3
- [80] Andreas Veit, Tomas Matera, Lukas Neumann, Jiri Matas, and Serge Belongie. Coco-text: Dataset and benchmark for text detection and recognition in natural images. 2016. 3
- [81] Zhaoyi Wan, Minghang He, Haoran Chen, Xiang Bai, and Cong Yao. Textscanner: Reading characters in order for robust scene text recognition. *CoRR*, abs/1912.12422, 2019. 8
- [82] Kai Wang, Boris Babenko, and Serge Belongie. End-to-end scene text recognition. pages 1457–1464, 2011. 3, 4
- [83] Qiang Wang, Bei Li, Tong Xiao, Jingbo Zhu, Changliang Li, Derek F Wong, and Lidia S Chao. Learning deep transformer models for machine translation. *arXiv preprint arXiv:1906.01787*, 2019. 3
- [84] Tianwei Wang, Yuanzhi Zhu, Lianwen Jin, Canjie Luo, Xiaoxue Chen, Yaqiang Wu, Qianying Wang, and Mingxiang Cai. Decoupled attention network for text recognition. 2019. 2
- [85] Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. *Advances in Neural Information Processing Systems*, 33:5776–5788, 2020. 3
- [86] Yuxin Wang, Hongtao Xie, Shancheng Fang, Jing Wang, Shenggao Zhu, and Yongdong Zhang. From two to one: A new scene text recognizer with visual language modeling network. 2021. 8, 2
- [87] Yuxin Wang, Hongtao Xie, Shancheng Fang, Mengting Xing, Jing Wang, Shenggao Zhu, and Yongdong Zhang. Petr: Rethinking the capability of transformer-based language model in scene text recognition. *IEEE Transactions on Image Processing*, 31:5585–5598, 2022. 8
- [88] Mingkun Yang, Minghui Liao, Pu Lu, Jing Wang, Shenggao Zhu, Hualin Luo, Qi Tian, and Xiang Bai. Reading and writing: Discriminative and generative modeling for self-supervised text recognition. 2023. 8
- [89] Jiabo Ye, Anwen Hu, Haiyang Xu, Qinghao Ye, Ming Yan, Yuhao Dan, Chenlin Zhao, Guohai Xu, Chenliang Li, Junfeng Tian, et al. mplug-docowl: Modularized multimodal large language model for document understanding. *arXiv preprint arXiv:2307.02499*, 2023. 4
- [90] Deli Yu, Xuan Li, Chengquan Zhang, Junyu Han, Jingtuo Liu, and Errui Ding. Towards accurate scene text recognition with semantic reasoning networks. *CoRR*, abs/2003.12294, 2020. 8, 2
- [91] Deli Yu, Xuan Li, Chengquan Zhang, Tao Liu, Junyu Han, Jingtuo Liu, and Errui Ding. Towards accurate scene text



- recognition with semantic reasoning networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12113–12122, 2020. 4
- [92] Xiaoyu Yue, Zhanghui Kuang, Chenhao Lin, Hongbin Sun, and Wayne Zhang. Robustscanner: Dynamically enhancing positional clues for robust text recognition. 2020. 2
- [93] Liu Yuliang, Jin Lianwen, Zhang Shuaitao, and Zhang Sheng. Detecting curve text in the wild: New dataset and new solution. *arXiv preprint arXiv:1712.02170*, 2017. 3
- [94] Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. Scaling vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12104–12113, 2022. 1, 2, 3
- [95] Y Zhang, L Gueguen, I Zharkov, P Zhang, K Seifert, and B Kadlec. Uber-text: A large-scale dataset for optical character recognition from street-level imagery. In *Proc. Scene Understand. Workshop*, pages 1–2, 2017. 3
- [96] Shuai Zhao, Xiaohan Wang, Linchao Zhu, Ruijie Quan, and Yi Yang. Clip4str: A simple baseline for scene text recognition with pre-trained vision-language model. 2023. 6, 8, 1, 2