

mPLUG-DocOwl 1.5: Unified Structure Learning for OCR-free Document Understanding

Anwen Hu¹, Haiyang Xu^{1*}, Jiabo Ye¹, Ming Yan^{1*}
Liang Zhang², Bo Zhang¹, Chen Li¹, Ji Zhang¹, Qin Jin², Fei Huang¹, Jingren Zhou¹

¹Alibaba Group

²Renmin University of China

{huanwen.haw, shuofeng.xhy, ym119608}@alibaba-inc.com

Abstract

Structure information is critical for understanding the semantics of text-rich images, such as documents, tables, and charts. Existing Multi-modal Large Language Models (MLLMs) for Visual Document Understanding are equipped with text recognition ability but lack general structure understanding abilities for text-rich document images. In this work, we emphasize the importance of structure information in Visual Document Understanding and propose Unified Structure Learning to boost the performance of MLLMs. Based on publicly available text-rich images, we build a comprehensive training set DocStruct4M to support structure-aware parsing tasks and multi-grained text localization tasks across 5 domains: document, webpage, table, chart, and natural image. To better encode structure information, we design a simple and effective vision-to-text module H-Reducer, which can not only maintain the layout information but also reduce the length of visual features by merging horizontal adjacent patches through convolution, enabling the LLM to understand high-resolution images more efficiently. Our model DocOwl 1.5 achieves state-of-the-art performance on 10 visual document understanding benchmarks. All codes, models, and datasets are publicly available at <https://github.com/X-PLUG/mPLUG-DocOwl/tree/main/DocOwl1.5>.

1 Introduction

Textual information in images manifests with a multitude of visual structures, spanning the simplicity of plain text to the systematic grid layouts of tables and incorporating a spectrum of graphical representations such as pie, line, and bar charts. These elements may appear in isolation or be intricately interwoven within the framework of documents and webpages, reflecting a rich diversity of informational architecture across posters, invoices,

infographics, scientific reports, academic and news websites, etc. As shown in Fig. 1, besides the basic textual content, structure information also plays a big role in Visual Document Understanding (Huang et al., 2022; Tang et al., 2023b; Lee et al., 2023).

With basic abilities to understand general images and comprehend structured texts through the LLM decoder, MLLMs (Ye et al., 2023c,d; Liu et al., 2023b,a; Zhu et al., 2023; Li et al., 2023a) have the potential to achieve unified structure learning on text-rich images. However, due to the visual encoder and vision-to-text module being trained on general image-text pairs and not specifically optimized to represent the textual and structural information in text-rich images, they still face great challenges with images with rich text information. For better Visual Document Understanding with MLLMs, some works (Ye et al., 2023a,b; Bai et al., 2023b; Feng et al., 2023; Wei et al., 2023) attempt to design text-reading tasks to strengthen the text recognition ability, but either ignore the structure comprehension or only cover limited domains, such as just webpages (Lee et al., 2023) or documents (Feng et al., 2023). In this work, we first perform unified structure learning on text-rich images for MLLMs across 5 domains: document, webpage, table, chart, and natural image.

For better structural understanding, we first design a simple and effective vision-to-text module, namely H-Reducer. Unlike the Resampler (Alayrac et al., 2022) or Q-former (Li et al., 2023a) which fuses visual features with learnable queries but affects spatial information, the H-Reducer accumulates neighborhood visual features through convolution to keep the relative positional relationships. Compared with V2T modules with only linear layers (Liu et al., 2023b,a), it produces much fewer visual features, which is more efficient for LLM to understand high-resolution document images. Considering texts in document images are most organized from left to right, H-Reducer merges visual

*Corresponding authors

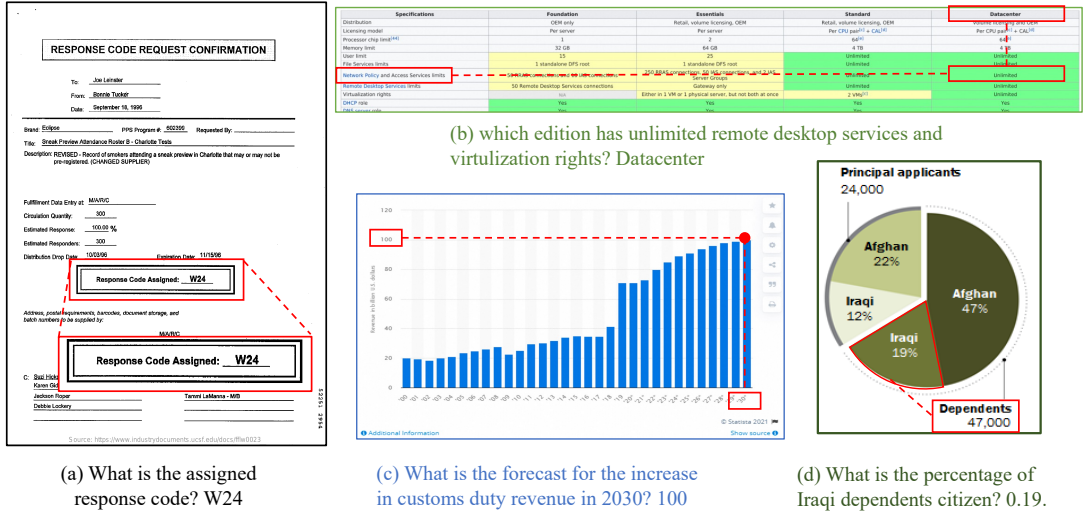


Figure 1: Illustrations of the importance of structure information in Visual Document Understanding.

features at the horizontal level.

Our Unified Structure Learning comprises structure-aware parsing tasks and multi-grained text localization tasks. To learn the organization of text contents, the former mainly teaches the model to parse the texts in the image in a structure-aware style, such as using line feeds and spaces to represent the structure of documents or webpages, and using extended Markdown syntax to represent the structure of tables and charts. Multi-grained text localization tasks further enhance the ability to correlate visually situated texts and concrete positions in the image. To support unified structure learning, based on publicly available datasets, we carefully build a comprehensive training set DocStruct4M by constructing structure-aware sequences and multi-grained pairs of text and bounding boxes. The DocOwl 1.5 is first trained with the Unified Structure Learning and then followed by the Multi-task Tuning among downstream tasks.

Our contributions in this work are three-fold:

- We first propose Unified Structure Learning on text-rich images for MLLMs and design both structure-aware parsing tasks and multi-grained text localization tasks across 5 domains. A comprehensive dataset DocStruct4M is carefully built to support Unified Structure Learning.
- We design a simple and effective vision-to-text module for structure learning and validate its effectiveness with extensive experiments.
- DocOwl 1.5 achieves state-of-the-art OCR-free performance on 10 Visual Document Un-

derstanding tasks.

2 Related Work

Visual Document Understanding (VDU) (Lee et al., 2023; Ye et al., 2023b), aims to comprehend images with rich text information. Such images range from documents (Mathew et al., 2021, 2022; Svetlichnaya, 2020; Stanislawek et al., 2021; Zhang et al., 2023a), tables (Pasapat and Liang, 2015; Chen et al., 2020; Zhong et al., 2020), charts (Masry et al., 2022; Kafle et al., 2018; Methani et al., 2020; Kantharaj et al., 2022; Tang et al., 2023a; Hu et al., 2023; Han et al., 2023; Meng et al., 2024), natural images (Sidorov et al., 2020; Singh et al., 2019; Hu et al., 2021) to webpage screenshots (Tanaka et al., 2021; Chen et al., 2021), where diverse composition of text and visual objects contains a wealth of information. According to whether relying on an off-the-shelf OCR system to recognize texts in the image, models for Visual Document Understanding can be categorized into OCR-dependent models (Tang et al., 2023b; Xu et al., 2021; Huang et al., 2022; Yang et al., 2021) and OCR-free ones (Kim et al., 2022; Lee et al., 2023). To leverage recognized texts from an OCR system, OCR-dependent models are always trained to align textual and visual inputs, such as recovering masked text and layout information given image and retained text as inputs (Tang et al., 2023b). As for OCR-free methods, training with tasks about text recognition is indispensable. Dount (Kim et al., 2022) designs the text reading task to output continuous text sequences that ignore structure information. To leverage structure

information, Pix2Struct (Lee et al., 2023) designs a Screenshot Parsing Task to generate the HTML DOM tree for webpage screenshots but is hard to apply to other types of images. In this work, we first propose Unified Structure Learning for all image types and carefully build a comprehensive dataset to support layout learning.

Multimodal Large Language Models (MLLM) have shown strong vision understanding and open-ended conversation abilities (Ye et al., 2023c,d; Zhu et al., 2023; Dai et al., 2023; Bai et al., 2023b; Hong et al., 2023; Zhang et al., 2024) for natural images. They follow the architecture paradigm of connecting a vision encoder, e.g. ViT (Dosovitskiy et al., 2021; Radford et al., 2021), with a Large Language Model (LLM) (Touvron et al., 2023; Vicuna, 2023; Bai et al., 2023a) by a Vision-to-Text (V2T) module, such as simple linear layers (Liu et al., 2023b,a) or a Q-Former (Li et al., 2023a)/Resampler (Alayrac et al., 2022)/Abstractor (Ye et al., 2023c,d) with learnable queries. To enable MLLMs to comprehend images with rich texts, there are major two challenges: how to encode high-resolution images and how to understand visually-situated texts. To tackle high-resolution images, most works choose to further train (Bai et al., 2023b; Feng et al., 2023) or extraly add a high-resolution vision encoder (Hong et al., 2023). UReader (Ye et al., 2023b) first proposes to keep the low-resolution vision encoder and use a shape-adaptive cropping module to crop raw images into multiple sub-images with low resolution. To enhance the visually-situated text understanding, some work design tasks of reading texts from top-left to bottom-right without taking into account the importance of structure (Ye et al., 2023b; Bai et al., 2023b). CogAgent (Hong et al., 2023) and DocPedia (Feng et al., 2023) further try strengthening the layout understanding for documents, webpages, and natural images with text grounding tasks. However, the comprehension of the overall structure is ignored, and tables and charts are not covered. In this work, to strengthen structure understanding, we design structure-aware parsing and multi-grained text localization tasks for all types of images, covering documents, tables, charts, webpages, and natural images. Besides, we propose a vision-to-text architecture H-Reducer to better maintain spatial information of visual features by convolution. Compared with MiniGPT-v2 (Chen et al., 2023), H-Reducer merges adjacent tokens with a convolution layer while MiniGPT-v2 con-

catenates 4 tokens and merges them with a linear layer. The convolution layer is more convenient for adjusting the merging scope, enabling us to perform an ablation study of different window sizes to determine the most suitable one for document understanding. The C-Abstractor (Cha et al., 2024) comprises 18 convolution layers with 1x1 sliding window and 1 mean pooling layer, showing that it doesn't reduce tokens with the convolution, but average pooling. However, H-Reducer just applies 1 convolution layer with a 1x4 window size to reduce vision tokens, with much fewer parameters and achieving comparable performance.

3 DocOwl 1.5

DocOwl 1.5 consists of a visual encoder, a large language model, and a vision-to-text module, namely H-Reducer to better keep the visual text and layout information by ensembling horizontal visual features. It's first trained with Unified Structure Learning to enhance text recognition and structure understanding and then tuned on multiple downstream tasks of Visual Document understanding.

3.1 Model Architecture

High-resolution Image Encoding. Encoding high-resolution images is critical to ensure that the decoder can use rich text information from document images (Kim et al., 2022; Lee et al., 2023; Ye et al., 2023b). As shown in Fig. 2, we utilize a parameter-free Shape-adaptive Cropping Module (Ye et al., 2023b) to crop a shape-variable high-resolution image I into multiple fixed-size sub-images (I_1, I_2, \dots, I_C), where C is the number of crops. To keep the overall layout information, the raw image is also resized to a low-resolution one as the global image I_0 . Then, each image I_i in (I_0, I_1, \dots, I_C) is independently encoded to a sequence of visual features $V_i = (v_i^1, v_i^2, \dots, v_i^L)$, $0 \leq i \leq C$ by a transformer-based Visual Encoder, where $v_i^j, 1 \leq j \leq L$ is a D -dimension vector, L is the length of visual features for each image.

Spatial-aware V2T Module: H-Reducer. There are two kinds of popular vision-to-text modules for MLLMs: a MLP (Liu et al., 2023b,a; Zhu et al., 2023) or a cross-attention module with learnable queries (Ye et al., 2023c; Bai et al., 2023b; Alayrac et al., 2022; Li et al., 2023a). Both two are not quite suitable for representing high-resolution text-rich images. The former projects complete visual features into the language embedding space. It

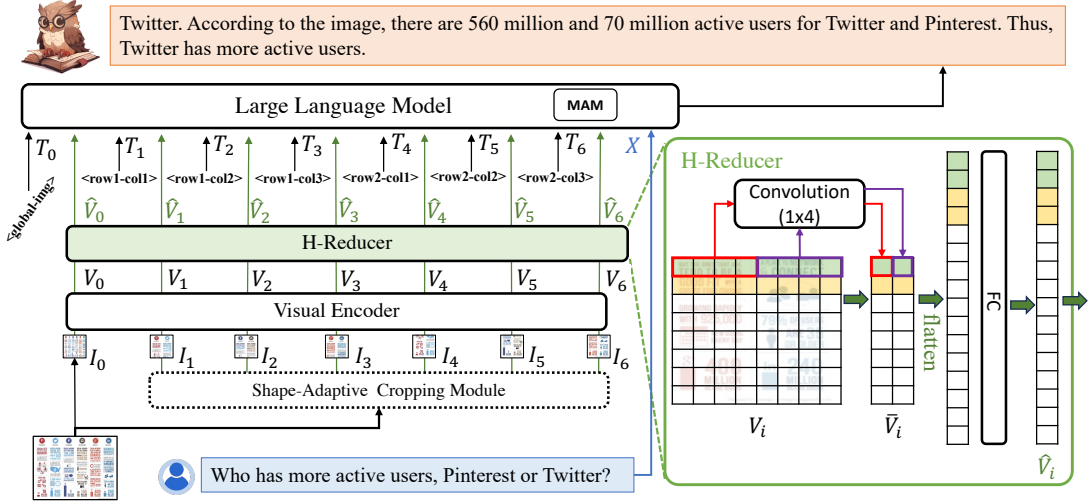


Figure 2: The overall architecture of DocOwl 1.5. The global image and cropped images are processed independently by the Visual Encoder and H-Reducer. $\langle \text{row}x\text{-col}y \rangle$ is the special textual token to indicate that the position of the cropped image in the original image is the x^{th} row and y^{th} column.

maintains all spatial information in the document image but keeps the sequence length of raw visual features, which is too long when processing high-resolution images. For example, encoding a $1,344 \times 1,344$ image with the ViT/L-14 results in 9,216 visual tokens. The cross-attention module could greatly reduce the length of the visual sequence to the number of learnable queries, but may lose spatial information during semantic fusion.

In this work, we design a more suitable vision-to-text module for Visual Document Understanding, namely H-Reducer, which not only reduces visual sequence length but also keeps the spatial information. As shown in Fig. 2, the H-Reducer is comprised of a convolution layer to reduce sequence length and a fully-connected layer to project visual features to language embedding space. Since most textual information in document images is arranged from left to right, the horizontal text information is usually semantically coherent. Thus, the kernel size and stride size in the convolution layer are set as 1×4 to ensemble horizontal 4 visual features. The output channel is set equal to the input channel D . The convolution calculation is as follows:

$$\bar{v}_i^j = f(v_i^{4j-3}, v_i^{4j-2}, v_i^{4j-1}, v_i^{4j}), 1 \leq j \leq L/4, \quad (1)$$

$$\bar{V}_i = (\bar{v}_i^1, \bar{v}_i^2, \dots, \bar{v}_i^{L/4}), \quad (2)$$

where f represents the dot product with kernel weights on multiple channels. After the convolution layer, the visual features V_i are converted to the \bar{V}_i . Then, with a fully connected layer to align

visual features to the language embedding space, the \bar{V}_i are transferred to $\hat{V}_i = (\hat{v}_i^1, \hat{v}_i^2, \dots, \hat{v}_i^{L/4})$.

Multimodal Modeling with LLM. To better distinguish visual and textual inputs, we insert the Modality-adaptive Module (MAM) (Ye et al., 2023d) into the LLM, which utilizes two sets of linear projection layers to separately perform the key/value projection for visual features and textual features. To help the LLM correlate multiple cropped sub-images, UReader (Ye et al., 2023b) designs learnable crop position embeddings to denote the row and column position in the raw image. In this work, we simply add special textual tokens ' $\langle \text{row}x\text{-col}y \rangle$ ' before the visual features of each cropped image, where x and y refer to the row and column index respectively. For the global image, the textual indicator token is ' $\langle \text{global_img} \rangle$ '. This design eliminates the need to introduce additional parameters and is more friendly to the LLM decoder. Our experiments validate that it achieves comparable effects as the crop position embedding. Overall, the decoding of the LLM is as follows:

$$Y = \text{LLM}([T_0; \hat{V}_0, T_1; \hat{V}_1, \dots, T_C; \hat{V}_C; X]) \quad (3)$$

where $[\cdot]$ means the concatenation operation, C is the crop number of the image, $T_j, 0 \leq j \leq C$ is the textual embeddings of the special textual indicator for the global image or positions of cropped images, \hat{V}_j is the visual features of a global or cropped image, X is the textual embeddings of the instruction, Y is the predicted answer.

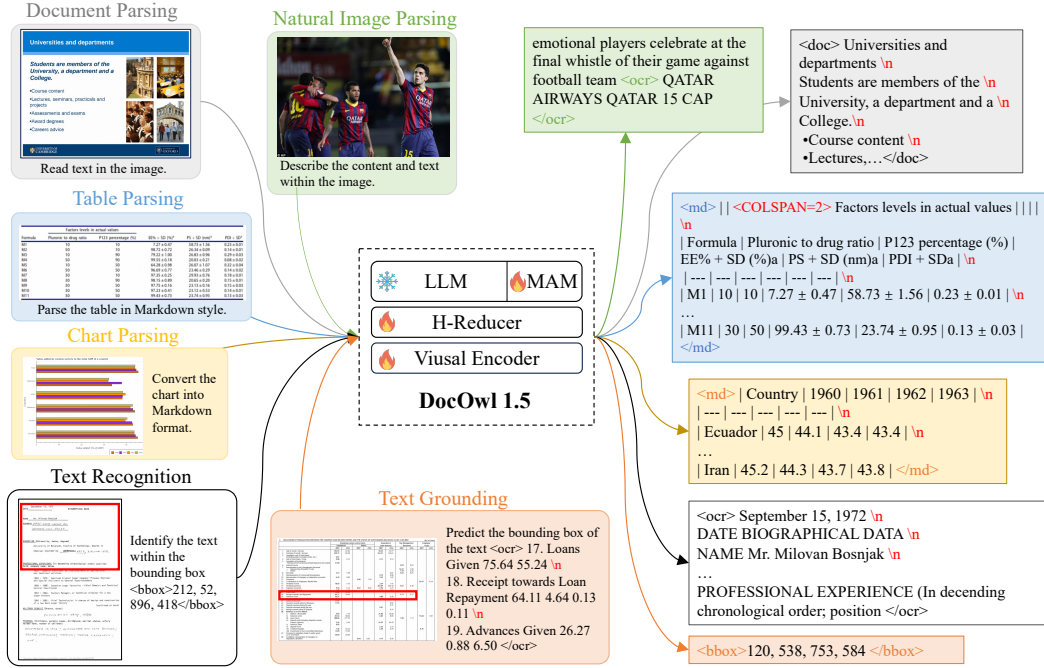


Figure 3: The illustration of Unified Structure Learning of DocOwl 1.5.

3.2 Unified Structure Learning

For comprehensive document understanding, we design a Unified Structure Learning across 5 domains, including natural images, documents, tables, charts, and webpages. It involves both structure-aware parsing tasks and multi-grained text localization tasks, as shown in Fig. 3.

Document/Webpage Parsing. To represent the structure information, Pix2Struct (Lee et al., 2023) parses webpage screenshots with condensed HTML DOM trees based on the HTML source codes but not feasible for other formats of documents or webpage screenshots, e.g. PDF. In documents or webpages, horizontal and vertical distances between texts form the main layout information. Therefore, to make the structure-aware parsing task applicable to most documents and webpage screenshots, we choose to add extra line feeds(‘\n’) and spaces into the text sequence to denote different lines and horizontal distances. The greater the horizontal distance, the more space characters.

Table Parsing. Different from documents or webpages, tables are structured in a more standardized way, where row and column correspondences represent key-value pairs. HTML and Markdown codes are mainly two kinds of text sequences used to represent a table. HTML codes can represent all kinds of tables, with or without cells spanning multiple rows and grids, but they contain too many paired

labels (e.g. ‘<tr></tr>’ and ‘<td></td>’), causing text sequences to be too long. Markdown codes can represent a table with concise text sequence, but they cannot represent cells spanning multiple rows and columns. To represent all tables with concise text sequence, we follow the main grammar of Markdown to represent table structure with ‘|’ and line feeds(‘\n’). To represent cells spanning multiple rows and columns, we add special text tokens ‘<COLSPAN=x>’ and ‘<ROWSPAN=y>’ before the value, as shown in Fig. 3.

Chart Parsing. Unlike documents and tables, organizing texts in reading order cannot represent the structure of charts. Considering that the chart is a visualization form of the table, parsing charts to tables could best maintain the mathematical characteristics of the chart. This requires the model to understand the structure of the chart and the alignment of the x/y axis. Besides, to keep consistent with the Table Parsing task, we also use Markdown codes to represent the data tables of charts, as shown in Fig. 3.

Natural Image Parsing. Semantics of natural images is a combination of natural objects and scene texts. Thus, parsing natural images is necessary to organize scene texts and mention the main image content. Manually annotating captions to describe the relationship between objects and scene texts is labour- and financial-intensive. Like TAP (Yang et al., 2021), we concatenate the general caption

with OCR texts to form the target parsing sequence. **Multi-grained Text Localization.** For Visual Document Understanding, structure-aware parsing tasks mainly focus on organizing texts according to the overall structure, while neglecting the correspondence between specific texts and local positions. Correlating texts with the concrete position in images is another basic structure understanding ability for visual documents. To support text position learning, we design two symmetrical tasks, namely Multi-grained Text Grounding and Multi-grained Text Recognition. The former aims to predict the bounding box given the visually-situated texts, while the latter does the opposite. We set four granularities of texts for these two tasks: word, phrase, line, and block, with progressively increasing text lengths. More details about such four granularities can be found in Appendix A.5.

Overall, to support the unified structure learning for text-rich images, we build a DocStruct4M dataset with 4M samples by ensembling training sets of 17 publicly available datasets and constructing structure-aware text sequences or text-position pairs as the targets. The form of instructions for each task is very diverse for developing the general instruction-following ability of the model. The statistics of DocStruct4M and details of target construction can be found in Appendix A.

3.3 Multi-task Fine-tuning

Through Unified Structure Learning, models could well understand the structure of diverse document images but cannot follow users' instructions to do different types of tasks, such as information extraction or image captioning. So, we further perform multi-task fine-tuning to train a generalist of visual document understanding as UReader.

3.4 Training Paradigm

DocOwl 1.5 is trained in a two-stage framework. Considering the LLM has strong comprehension abilities for structured text (Wang et al., 2023b; Zhang et al., 2023b), we argue that the main limitation of MLLM in visual document understanding is the representation ability of the Visual Encoder and Vision-to-Text module for visual text and structure information. Thus, during the Unified Structure Learning, we freeze the LLM and tune the Visual Encoder and H-Reducer. The MAM is also optimized to help the LLM better distinguish visual features and texts parsed from the image. During the stage of Multi-task Fine-tuning, the model

mainly learns how to follow the user's instructions to give answers based on visually-situated text and structure understanding capabilities acquired in the first stage. Therefore, the Visual Encoder is frozen and other modules are tuned. More implementation details can be found in Appendix B.1.

4 Experiments

4.1 Main Results

We evaluate the Visual Document Understanding performance on 10 text-rich image benchmarks, covering documents (DocVQA (Mathew et al., 2021), InfoVQA (Mathew et al., 2022), DeepForm (Svetlichnaya, 2020), KLC (Stanislawek et al., 2021)), tables (WTQ (Pasupat and Liang, 2015), TabFact (Chen et al., 2020)), charts (ChartQA (Masry et al., 2022)), natural images (TextVQA (Singh et al., 2019), TextCaps (Sidorov et al., 2020)), and webpage screenshots (VisualMRC (Tanaka et al., 2021)). We compare DocOwl 1.5 with state-of-the-art OCR-free models, including both Multimodal Large Language Models adapted for recognizing texts and much smaller models trained only for document understanding. The detailed comparison of model settings can be found in Appendix B.2.

As shown in Table 1, previous MLLMs with more than 7B parameters underperform domain-specific models with less than 1B parameters, showing that the document understanding is still a shortcoming for existing MLLMs. Our DocOwl 1.5 outperforms both domain-specific models and MLLMs with similar sizes on all 10 benchmarks. This validates that DocOwl 1.5 is much stronger on visual document understanding across 5 domains, covering visual question answering, information retrieval, natural language inference, and image captioning tasks. Besides, with much fewer unnatural data (3M vs 9M) and parameters (8.1B vs 17.3B), DocOwl 1.5 outperforms CogAgent (Hong et al., 2023) on InfoVQA and ChartQA, and achieves comparable performance on DocVQA. This suggests that our unified structure learning with DocStruct4M is more efficient in learning printed text recognition and how to analyze documents. However, our model still underperforms CogAgent on TextVQA, which requires the ability of scene text recognition and general knowledge about natural objects. The primary reason is that scene texts are more diverse in shapes than printed texts and CogAgent is trained on 98M samples of scene text recog-

Table 1: Comparison with OCR-free methods on various types of text-rich image understanding tasks. The superscript ‘*’ refers to models separately fine-tuned on each downstream task, rather than generalists. The *underline* means the best performance among models with <10B parameters.

Model	Size	Doc VQA	Info VQA	Deep Form	KLC	WTQ	Tab Fact	Chart QA	Text VQA	Text Caps	Visual MRC
Dessurt*	<1B	63.2	-	-	-	-	-	-	-	-	-
Donut*	<1B	67.5	11.6	61.6	30.0	18.8	54.6	41.8	43.5	74.4	93.91
Pix2Struct _{base} *	<1B	72.1	38.2	-	-	-	-	56.0	-	88.0	-
Pix2Struct _{large} *	1.3B	76.6	40.0	-	-	-	-	58.6	-	95.5	-
DocPeida	7.0B	47.1	15.2	-	-	-	-	46.9	60.2	-	-
DocOwl	7.1B	62.2	38.2	42.6	30.3	26.9	60.2	57.4	52.6	111.9	188.8
QwenVL	9.6B	65.1	35.4	-	-	-	-	65.7	63.8	-	-
UReader	7.1B	65.4	42.2	49.5	32.8	29.4	67.6	59.3	57.6	118.4	221.7
Monkey	9.8B	66.5	36.1	40.6	32.8	25.3	-	-	67.6	93.2	-
Vary	7.0B	76.3	-	-	-	-	-	66.1	-	-	-
CogAgent	17.3B	81.6	44.5	-	-	-	-	68.4	76.1	-	-
DocOwl-1.5	8.1B	81.6	50.4	68.8	37.9	39.8	80.4	70.5	<u>68.8</u>	132.0	239.5

dition from LAION-2B (Schuhmann et al., 2022) and COYO-700M (Byeon et al., 2022), much more than the natural images (1M) in DocStruct4M. In this work, we mainly focus on improving the unified structure comprehension of visual documents and leave further scaling up data on natural scenes as future work.

4.2 Ablation Study

As shown in Table 2, we further perform a comprehensive ablation study to validate the effectiveness of our H-Reducer and Unified Structure Learning.

Firstly, initializing from a stronger general MLLMs brings better performance on text-rich images (r2 vs r1), showing general vision-and-language knowledge benefits visual document understanding. Tuning the visual encoder during multi-task fine-tuning significantly improves the document understanding performance (r3 vs r2). This suggests that the visual representation of document images may be the main shortcoming of MLLMs and inspires us to design Unified Structure Learning to enhance the representation ability of the visual encoder for visual texts and structure. **Effectiveness of H-Reducer.** When using the Shape-adaptive Cropping Module, the image resolution supported by the MLLM is the product of the cropping number and basic resolution of each crop. With the Abstractor as the vision-to-text module, reducing the cropping number causes an obvious performance decrease (r4 vs r3) on documents. However, with a smaller cropping number, the H-Reducer achieves better performance than the Abstractor (r6 vs r3), showing that $448^2 \times 9 \approx 2^{21}$ is an acceptable resolution for existing benchmarks

and the H-Reducer is stronger on maintaining rich text information during vision-and-language feature alignment. Compared with the CAbstractor (Cha et al., 2024), H-Reducer achieves similar performance (r6 vs r5) while using 10% parameters of C-Abstractor. This shows merging tokens with 1 learnable convolution layer is more parameter-efficient than multiple 1x1 convolution layers with mean pooling in document understanding. Besides, we further compare different settings of the merging shape in the convolution layer. With the same number of merged tokens, the model with the 1x4 merging shape achieves better performance than the one with the 2x2 merging shape on document and table datasets but slightly worse performance on chart understanding (r7 vs r6). This is consistent with the common sense that documents and tables mainly organize texts in the left-to-right order while the semantic structures of charts are much more flexible. A square merging shape is more suited to encode visual features in the form of bars, lines, or pies while the 1x4 merging shape is more appropriate for general document understanding. Further extending the 1x4 merging shape horizontally and vertically decreases the length of visual features but at the cost of performance degradation (r8-r10). Considering the overall performance on all text-rich images, we finally choose the 1x4 as the merging shape in H-Reducer.

Besides proving the effectiveness of H-Reducer through downstream text-rich image understanding performance, we further directly compare the text localization performance after the Unified Structure Learning. Experiments

Table 2: Ablation study of model setting. ‘Crop’ refers to the maximum number of cropped images. ‘CropPos’ means using learnable embeddings (‘Emb’) or textual tokens (‘Text’) to represent the position of cropped images. ‘Parsing’ and ‘MTL’ refer to structure-aware parsing tasks and the Multi-grained Text Location task, respectively. ‘Owl(224)’ and ‘Owl2(448)’ refer to mPLUG-Owl with 224 resolution and mPLUG-Owl2 with 448 resolution.

	Init	Model Architecture			Structure Learning	Multi-task Tuning		DocVQA	TabFact	ChartQA
		V2T	Crop	CropPos		ViT	LLM			
r1	Owl(224)	Abstractor	20	Emb	×	×	×	65.4	67.6	59.3
r2	Owl2(448)	Abstractor	20	Emb	×	×	×	66.3	69.8	60.6
r3	Owl2(448)	Abstractor	20	Emb	×	✓	×	71.4	70.3	64.2
r4	Owl2(448)	Abstractor	9	Emb	×	✓	×	68.0	70.0	64.2
r5	Owl2(448)	CAbstractor	9	Emb	×	✓	×	72.6	72.3	65.2
r6	Owl2(448)	H-Reducer(1x4)	9	Emb	×	✓	×	72.8	72.9	65.0
r7	Owl2(448)	H-Reducer(2x2)	9	Emb	×	✓	×	71.8	72.1	65.2
r8	Owl2(448)	H-Reducer(2x4)	9	Emb	×	✓	×	71.4	71.1	66.0
r9	Owl2(448)	H-Reducer(1x8)	9	Emb	×	✓	×	69.9	71.2	64.4
r10	Owl2(448)	H-Reducer(2x8)	9	Emb	×	✓	×	69.2	70.2	65.6
r11	Owl2(448)	H-Reducer(1x4)	9	Emb	Parsing	×	×	77.7	76.5	67.5
r12	Owl2(448)	H-Reducer(1x4)	9	Emb	Parsing	×	✓	78.9	78.1	68.1
r13	Owl2(448)	H-Reducer(1x4)	9	Text	Parsing	×	✓	79.8	77.7	69.1
r14	Owl2(448)	H-Reducer(1x4)	9	Text	Parsing+MTL	×	✓	81.6	80.4	70.5

on a carefully-built Multi-grained Text Localization evaluation set DocLocal4K validate that H-Reducer(1x4) is better than H-Reducer(2x2) and Abstractor in preserving spatial features. More details about DocLocal4K and localization performance can be found in Appendix B.4.

Effectiveness of Unified Structure Learning. After determining the vision-to-text module, we perform two-stage training with Unified Structure Learning. With only the structure-aware parsing tasks, there is significant improvement across different domains (r11 vs r6). This validates that fine-tuning the visual encoder and H-Reducer with structure-aware parsing tasks greatly helps MLLMs understand text-rich images. Further tuning the parameters of LLM brings slight improvement (r12 vs r11), suggesting that general language knowledge is not the main obstacle to visual document understanding. By replacing the learnable crop position embeddings with special textual tokens, the model achieves better performance (r13 vs r12), showing that the LLM can well understand the relative positions of multiple cropped images with just simple textual indicators. Finally, by introducing Multi-grained Text Localization tasks, DocOwl 1.5 achieves the best performance, validating that correlating visual texts with concrete positions helps comprehend documents more accurately.

Effectiveness of the Two-stage Training. Besides two-stage training, we also try one-stage joint training of the structure learning and downstream tasks and gradually increase samples from DocStruct4M. Our experiments validate that the two-stage training could better enhance basic text recognition and

structure parsing abilities and is more beneficial and efficient for downstream document understanding. More details can be found in Appendix B.5.

4.3 Qualitative Results

Besides quantitative results, we further present some qualitative results of visual document understanding on different domains of images. As shown in Fig. 4(a) and (b), both models answer the question with texts in the image. DocOwl 1.5 can better understand the structure of two documents and give correct answers. In Fig. 4(c), due to the learning of parsing chart with Markdown codes, DocOwl 1.5 can better understand the chart and successfully correlate the x/y axis. Fig. 4(d) shows that although inconsistent with the ground truth, DocOwl 1.5 gives another correct answer with the help of stronger structure understanding on tables. More qualitative results can be found in Appendix B.6.

5 Conclusion

To enhance the Visual Document Understanding performance of MLLMs, we first propose Unified Structure Learning across 5 domains of text-rich images, including both structure-aware parsing tasks and multi-grained text localization tasks. To better maintain structure and spatial information during vision-and-language feature alignment, we design a simple and effective vision-to-text module, named H-Reducer. It mainly utilizes a convolution layer to aggregate horizontally neighboring visual features. To support the Unified Structure Learning, we build DocStruct4M by collecting publicly available images and carefully constructing structure-

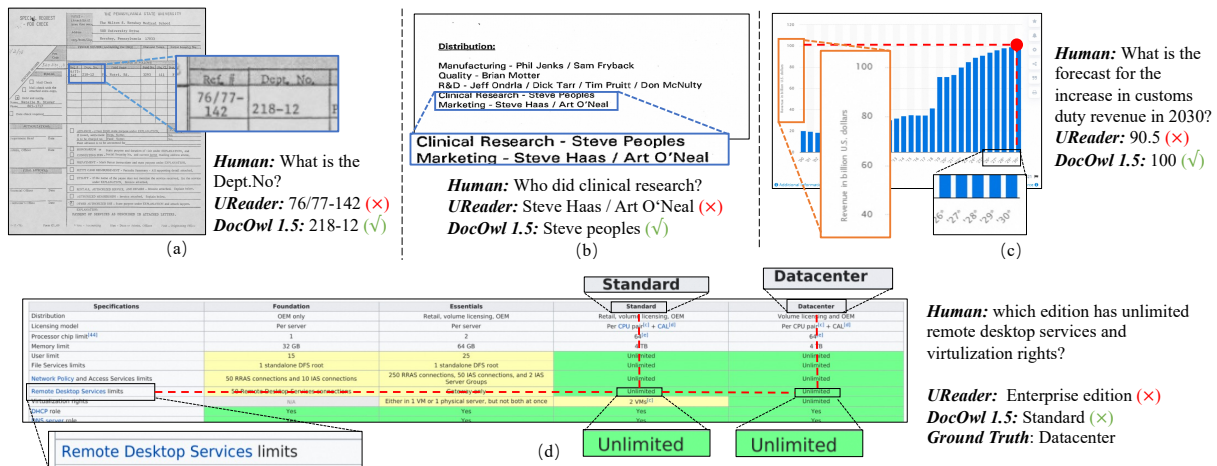


Figure 4: Qualitative results of DocOwl 1.5 and UReader on different domains of images.

aware text sequences and multi-grained pairs of texts and bounding boxes. With Unified Structure Learning, our model DocOwl 1.5 achieves state-of-the-art OCR-free performance on 10 visual document understanding benchmarks.

6 Limitation

In this work, we mainly focus on enhancing the OCR-free document understanding performance of the Multimodal Large Language Model through unified structure learning and pay less attention to the efficiency of vision encoding. Our H-Reducer could preserve spatial information and produce fewer visual tokens than MLP. However, a high-resolution image can still occupy a maximum of 2,560 tokens during the language modeling of the LLM decoder. This is not efficient enough when understanding a document with multiple pages (e.g., 20 pages will occupy a maximum of 51,200 tokens). How to encode high-resolution images with fewer tokens and maintain the document understanding performance are left to future work.

7 Ethics Statement

Our DocOwl 1.5 is initialized from a general Multimodal Large Language Model, which is trained on massive vision-and-language data from the web. Thus, it may also suffer from issues of LLMs such as toxic language and bias (Bender et al., 2021). However, the Unified Structure Learning proposed in this work focuses on parsing texts in publicly available images and introduces few biases relevant to ethical issues.

References

- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob Menick, Sebastian Borgeaud, Andy Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karen Simonyan. 2022. Flamingo: a visual language model for few-shot learning. *ArXiv*, abs/2204.14198.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Sheng-guang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. 2023a. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023b. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*.
- Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 610–623.
- Lukasz Borchmann, Michal Pietruszka, Tomasz Stanislawek, Dawid Jurkiewicz, Michal Turski, Karolina

- Szyndler, and Filip Gralinski. 2021. DUE: end-to-end document understanding benchmark. In *NeurIPS Datasets and Benchmarks*.
- Minwoo Byeon, Beomhee Park, Haecheon Kim, Sungjun Lee, Woonhyuk Baek, and Saehoon Kim. 2022. Coyo-700m: Image-text pair dataset. <https://github.com/kakaobrain/coyo-dataset>.
- Junbum Cha, Wooyoung Kang, Jonghwan Mun, and Byungseok Roh. 2024. Honeybee: Locality-enhanced projector for multimodal llm. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Jun Chen, Deyao Zhu, Xiaoqian Shen, Xiang Li, Zechun Liu, Pengchuan Zhang, Raghuraman Krishnamoorthi, Vikas Chandra, Yunyang Xiong, and Mohamed Elhoseiny. 2023. Minigt-v2: large language model as a unified interface for vision-language multi-task learning. *CoRR*, abs/2310.09478.
- Wenhu Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyu Zhou, and William Yang Wang. 2020. Tabfact : A large-scale dataset for table-based fact verification. In *International Conference on Learning Representations (ICLR)*, Addis Ababa, Ethiopia.
- Xingyu Chen, Zihan Zhao, Lu Chen, JiaBao Ji, Danyang Zhang, Ao Luo, Yuxuan Xiong, and Kai Yu. 2021. Websrc: A dataset for web-based structural reading comprehension. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4173–4185.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven C. H. Hoi. 2023. Instructblip: Towards general-purpose vision-language models with instruction tuning. *CoRR*, abs/2305.06500.
- Xiang Deng, Huan Sun, Alyssa Lees, You Wu, and Cong Yu. 2022. TURL: table understanding through representation learning. *SIGMOD Rec.*, 51(1):33–40.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*. OpenReview.net.
- Hao Feng, Qi Liu, Hao Liu, Wengang Zhou, Houqiang Li, and Can Huang. 2023. Docpedia: Unleashing the power of large multimodal model in the frequency domain for versatile document understanding. *CoRR*, abs/2311.11810.
- Yucheng Han, Chi Zhang, Xin Chen, Xu Yang, Zhibin Wang, Gang Yu, Bin Fu, and Hanwang Zhang. 2023. Chartllama: A multimodal LLM for chart understanding and generation. *CoRR*, abs/2311.16483.
- Adam W. Harley, Alex Ufkes, and Konstantinos G. Derpanis. 2015. Evaluation of deep convolutional nets for document image classification and retrieval. In *ICDAR*, pages 991–995. IEEE Computer Society.
- Wenyi Hong, Weihan Wang, Qingsong Lv, Jiazheng Xu, Wenmeng Yu, Junhui Ji, Yan Wang, Zihan Wang, Yuxuan Zhang, Juanzi Li, Bin Xu, Yuxiao Dong, Ming Ding, and Jie Tang. 2023. Cogagent: A visual language model for GUI agents. *CoRR*, abs/2312.08914.
- Anwen Hu, Shizhe Chen, and Qin Jin. 2021. Question-controlled text-aware image captioning. In *ACM Multimedia*, pages 3097–3105. ACM.
- Anwen Hu, Yaya Shi, Haiyang Xu, Jiabo Ye, Qinghao Ye, Ming Yan, Chenliang Li, Qi Qian, Ji Zhang, and Fei Huang. 2023. mplug-paperowl: Scientific diagram analysis with the multimodal large language model. *arXiv preprint arXiv:2311.18248*.
- Yupan Huang, Tengchao Lv, Lei Cui, Yutong Lu, and Furu Wei. 2022. Layoutlmv3: Pre-training for document AI with unified text and image masking. In *ACM Multimedia*, pages 4083–4091. ACM.
- Kushal Kafle, Brian L. Price, Scott Cohen, and Christopher Kanan. 2018. DVQA: understanding data visualizations via question answering. In *CVPR*, pages 5648–5656. Computer Vision Foundation / IEEE Computer Society.
- Samira Ebrahimi Kahou, Vincent Michalski, Adam Atkinson, Ákos Kádár, Adam Trischler, and Yoshua Bengio. 2018. Figureqa: An annotated figure dataset for visual reasoning. In *ICLR (Workshop)*. OpenReview.net.
- Shankar Kantharaj, Rixie Tiffany Ko Leong, Xiang Lin, Ahmed Masry, Megh Thakkar, Enamul Hoque, and Shafiq R. Joty. 2022. Chart-to-text: A large-scale benchmark for chart summarization. In *ACL (1)*, pages 4005–4023. Association for Computational Linguistics.
- Geewook Kim, Teakgyu Hong, Moonbin Yim, Jeongyeon Nam, Jinyoung Park, Jinyeong Yim, Wonseok Hwang, Sangdoon Yun, Dongyoon Han, and Seunghyun Park. 2022. Ocr-free document understanding transformer. In *ECCV (28)*, volume 13688 of *Lecture Notes in Computer Science*, pages 498–517. Springer.
- Kenton Lee, Mandar Joshi, Iulia Raluca Turc, Hexiang Hu, Fangyu Liu, Julian Martin Eisenschlos, Urvashi Khandelwal, Peter Shaw, Ming-Wei Chang, and Kristina Toutanova. 2023. Pix2struct: Screenshot parsing as pretraining for visual language understanding. In *ICML*, volume 202 of *Proceedings of Machine Learning Research*, pages 18893–18912. PMLR.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven C. H. Hoi. 2023a. BLIP-2: bootstrapping language-image pre-training with frozen image encoders and large language models. *CoRR*, abs/2301.12597.

- Zhang Li, Biao Yang, Qiang Liu, Zhiyin Ma, Shuo Zhang, Jingxu Yang, Yabo Sun, Yuliang Liu, and Xiang Bai. 2023b. Monkey: Image resolution and text label are important things for large multi-modal models. *CoRR*, abs/2311.06607.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2023a. Improved baselines with visual instruction tuning. *CoRR*, abs/2310.03744.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023b. Visual instruction tuning. *CoRR*, abs/2304.08485.
- Ahmed Masry, Do Xuan Long, Jia Qing Tan, Shafiq R. Joty, and Enamul Hoque. 2022. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. In *ACL (Findings)*, pages 2263–2279. Association for Computational Linguistics.
- Minesh Mathew, Viraj Bagal, Rubèn Tito, Dimosthenis Karatzas, Ernest Valveny, and C. V. Jawahar. 2022. Infographicvqa. In *WACV*, pages 2582–2591. IEEE.
- Minesh Mathew, Dimosthenis Karatzas, and C. V. Jawahar. 2021. Docvqa: A dataset for VQA on document images. In *WACV*, pages 2199–2208. IEEE.
- Fanqing Meng, Wenqi Shao, Quanfeng Lu, Peng Gao, Kaipeng Zhang, Yu Qiao, and Ping Luo. 2024. Chartassistant: A universal chart multimodal language model via chart-to-table pre-training and multitask instruction tuning. *CoRR*, abs/2401.02384.
- Nitesh Methani, Pritha Ganguly, Mitesh M. Khapra, and Pratyush Kumar. 2020. Plotqa: Reasoning over scientific plots. In *WACV*, pages 1516–1525. IEEE.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Panupong Pasupat and Percy Liang. 2015. Compositional semantic parsing on semi-structured tables. In *ACL (1)*, pages 1470–1480. The Association for Computer Linguistics.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning transferable visual models from natural language supervision. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. 2022. LAION-5B: an open large-scale dataset for training next generation image-text models. In *NeurIPS*.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *ACL (1)*, pages 2556–2565. Association for Computational Linguistics.
- Oleksii Sidorov, Ronghang Hu, Marcus Rohrbach, and Amanpreet Singh. 2020. Textcaps: A dataset for image captioning with reading comprehension. In *ECCV (2)*, volume 12347 of *Lecture Notes in Computer Science*, pages 742–758. Springer.
- Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. 2019. Towards VQA models that can read. In *CVPR*, pages 8317–8326. Computer Vision Foundation / IEEE.
- Tomasz Stanislawek, Filip Gralinski, Anna Wróblewska, Dawid Lipinski, Agnieszka Kaliska, Paulina Rosalska, Bartosz Topolski, and Przemyslaw Biecek. 2021. Kleister: Key information extraction datasets involving long documents with complex layouts. In *ICDAR (1)*, volume 12821 of *Lecture Notes in Computer Science*, pages 564–579. Springer.
- S Svetlichnaya. 2020. Deepform: Understand structured documents at scale.
- Ryota Tanaka, Kyosuke Nishida, and Sen Yoshida. 2021. Visualmrc: Machine reading comprehension on document images. In *AAAI*, pages 13878–13888. AAAI Press.
- Benny J. Tang, Angie Boggust, and Arvind Satyanarayan. 2023a. Vistext: A benchmark for semantically rich chart captioning. In *ACL (1)*, pages 7268–7298. Association for Computational Linguistics.
- Zineng Tang, Ziyi Yang, Guoxin Wang, Yuwei Fang, Yang Liu, Chenguang Zhu, Michael Zeng, Cha Zhang, and Mohit Bansal. 2023b. Unifying vision, text, and layout for universal document processing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19254–19264.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Michał Turski, Tomasz Stanislawek, Karol Kaczmarek, Paweł Dyda, and Filip Gralinski. 2023. Ccpdf: Building a high quality corpus for visually rich documents from web crawl data. In *ICDAR (3)*, volume 14189 of *Lecture Notes in Computer Science*, pages 348–365. Springer.
- Vicuna. 2023. Vicuna: An open chatbot impressing gpt-4. <https://github.com/lm-sys/FastChat>.

- Weihan Wang, Qingsong Lv, Wenmeng Yu, Wenyi Hong, Ji Qi, Yan Wang, Junhui Ji, Zhuoyi Yang, Lei Zhao, Xixuan Song, Jiazheng Xu, Bin Xu, Juanzi Li, Yuxiao Dong, Ming Ding, and Jie Tang. 2023a. Cogvlm: Visual expert for pretrained language models. *CoRR*, abs/2311.03079.
- Wenjin Wang, Yunhao Li, Yixin Ou, and Yin Zhang. 2023b. Layout and task aware instruction prompt for zero-shot document image question answering. *CoRR*, abs/2306.00526.
- Haoran Wei, Lingyu Kong, Jinyue Chen, Liang Zhao, Zheng Ge, Jinrong Yang, Jianjian Sun, Chunrui Han, and Xiangyu Zhang. 2023. Vary: Scaling up the vision vocabulary for large vision-language models. *CoRR*, abs/2312.06109.
- Yang Xu, Yiheng Xu, Tengchao Lv, Lei Cui, Furu Wei, Guoxin Wang, Yijuan Lu, Dinei A. F. Florêncio, Cha Zhang, Wanxiang Che, Min Zhang, and Lidong Zhou. 2021. Layoutlmv2: Multi-modal pre-training for visually-rich document understanding. In *ACL/IJCNLP (1)*, pages 2579–2591. Association for Computational Linguistics.
- Zhengyuan Yang, Yijuan Lu, Jianfeng Wang, Xi Yin, Dinei Florêncio, Lijuan Wang, Cha Zhang, Lei Zhang, and Jiebo Luo. 2021. TAP: text-aware pre-training for text-vqa and text-caption. In *CVPR*, pages 8751–8761. Computer Vision Foundation / IEEE.
- Jiabo Ye, Anwen Hu, Haiyang Xu, Qinghao Ye, Ming Yan, Yuhao Dan, Chenlin Zhao, Guohai Xu, Chenliang Li, Junfeng Tian, Qian Qi, Ji Zhang, and Fei Huang. 2023a. mplug-docowl: Modularized multimodal large language model for document understanding. *CoRR*, abs/2307.02499.
- Jiabo Ye, Anwen Hu, Haiyang Xu, Qinghao Ye, Ming Yan, Guohai Xu, Chenliang Li, Junfeng Tian, Qi Qian, Ji Zhang, Qin Jin, Liang He, Xin Lin, and Fei Huang. 2023b. Ureader: Universal ocr-free visually-situated language understanding with multi-modal large language model. In *EMNLP (Findings)*, pages 2841–2858. Association for Computational Linguistics.
- Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, Chenliang Li, Yuanhong Xu, Hehong Chen, Junfeng Tian, Qian Qi, Ji Zhang, and Fei Huang. 2023c. mplug-owl: Modularization empowers large language models with multimodality. *CoRR*, abs/2304.14178.
- Qinghao Ye, Haiyang Xu, Jiabo Ye, Ming Yan, Anwen Hu, Haowei Liu, Qi Qian, Ji Zhang, Fei Huang, and Jingren Zhou. 2023d. mplug-owl2: Revolutionizing multi-modal large language model with modality collaboration. *CoRR*, abs/2311.04257.
- Duzhen Zhang, Yahan Yu, Chenxing Li, Jiahua Dong, Dan Su, Chenhui Chu, and Dong Yu. 2024. Mmllms: Recent advances in multimodal large language models. *arXiv preprint arXiv:2401.13601*.
- Liang Zhang, Anwen Hu, Jing Zhang, Shuo Hu, and Qin Jin. 2023a. MPMQA: multimodal question answering on product manuals. *CoRR*, abs/2304.09660.
- Tianshu Zhang, Xiang Yue, Yifei Li, and Huan Sun. 2023b. Tablellama: Towards open large generalist models for tables. *CoRR*, abs/2311.09206.
- Xu Zhong, Elaheh ShafieiBavani, and Antonio Jimeno-Yepes. 2020. Image-based table recognition: Data, model, and evaluation. In *ECCV (21)*, volume 12366 of *Lecture Notes in Computer Science*, pages 564–580. Springer.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2023. Minigpt-4: Enhancing vision-language understanding with advanced large language models.

A DocStruct4M

A.1 Document/Webpage Parsing

We choose CCpdf (Turski et al., 2023), RVL-CDIP (Harley et al., 2015), VisualMRC (Tanaka et al., 2021) and datasets encapsulated in DUE-Benchmark (Borchmann et al., 2021) (DocVQA (Mathew et al., 2021), InfoVQA (Mathew et al., 2022), DeepForm (Svetlichnaya, 2020), KLC (Stanislawek et al., 2021), WTQ (Pasupat and Liang, 2015), TabFact (Chen et al., 2020)) to support the Document/Webpage Parsing task.

CCpdf (Turski et al., 2023) is a multi-lingual PDF dataset built upon webpages from Common Crawl¹, covering diverse domains of documents, such as industry, academic, and medical. In this work, we mainly focus on English Document Understanding and drop PDFs detected as other languages. RVL-CDIP contains 16 categories of industry documents, such as ‘letter’, ‘email’, and ‘scientific reports’. We further remove some categories with flipping and blurring texts, such as ‘handwritten’ and ‘form’. DUE-Benchmark is a collection of available and reformulated datasets over various document domains and layouts featuring tables, graphs, lists, and infographics. VisualMRC is a webpage screenshot dataset across 35 websites. OCR annotations in VisualMRC are aligned with local regions, thus, we follow them to utilize crops of a screenshot as input for this parsing task. For CCpdf and DUE-Benchmark, a PDF-parsing tool pdfplumber² can be directly used to generate structure-aware text sequence with a PDF page as the input. For RVL-CDIP and VisualMRC, there are no PDF files, just annotations of bounding boxes of texts. As an alternative, akin to the LATIN-Prompt (Wang et al., 2023b), we insert the line feeds and spaces by calculating and comparing the horizontal and vertical distances of bounding boxes. To avoid too many space characters resulting in sparse texts, we further limit the maximum number of consecutive spaces to 4. This strategy allows us to construct structure-aware text sequences in the same style as pdfplumber.

A.2 Table Parsing

We choose TURL (Deng et al., 2022) and PubTabNet (Zhong et al., 2020) to do the structure-aware table parsing task, where tables are col-

lected from Wikipedia pages and scientific articles, respectively. Without cells across rows and columns, tables in TURL can be directly represented with Markdown codes. Due to lacking table images in TURL, we transfer tables into HTML codes and render table images with variations in background color and font size. PubTabNet contains pairs of table images and HTML codes. We convert HTML codes into Markdown style and add ‘<ROWSPAN=x>’ or ‘<COLSPAN=y>’ before the value when attributes ‘rowspan=x’ or ‘colspan=y’ are set in the ‘<td>’ label.

A.3 Chart Parsing

We adopt PlotQA (Methani et al., 2020), FigureQA (Kahou et al., 2018), DVQA (Kafle et al., 2018), and ChartQA (Masry et al., 2022) to support the structure-aware chart parsing task. These datasets cover charts on both synthetic (Kahou et al., 2018; Kafle et al., 2018) data and data from real-world sources (Methani et al., 2020; Masry et al., 2022). Chart types include vertical bar, horizontal bar, line, dot line, and pie chart. Source data of the chart is provided in the JSON (Methani et al., 2020; Kahou et al., 2018; Methani et al., 2020) or CSV format (Masry et al., 2022), both can be conveniently converted to Markdown codes. However, some raw values are not suitable as standard answers for parsing because there are too many significant digits to be represented on the chart. Therefore, to reduce the difficulty of estimating values and make the model focus more on structural understanding, we keep 4 significant digits for all values.

A.4 Natural Image Parsing

We utilize OCR-CC (Yang et al., 2021) to support the Natural Image Parsing task. OCR-CC is a subset of Conceptual Caption (Sharma et al., 2018), which contains images with scene texts detected by the Microsoft Azure OCR system.

A.5 Multi-grained Text Localization

We set four granularities of texts for both Text Grounding and Text Recognition tasks: word, phrase, line, and block. The ‘word’ is the smallest granularity of the bounding box, referring to only 1 word. To ensure that the word is visible and the answer is unique, words that are too small (normalized area < 0.001) and words that appear multiple times in the same image are excluded from candidates. The ‘line’ consists of texts that are judged to

¹<https://commoncrawl.org>

²<https://github.com/jsvine/pdfplumber>



Task	Dataset	Num
Document Parsing	CCpdf	938,058
	RVL-CDIP	159,418
	DUE	56,809
	VisualMRC	7,016
Table Parsing	TURL	200,000
	PubTabNet	200,000
Chart Parsing	ChartQA	18,317
	FigureQA	99,714
	PlotQA	157,070
	DQA	200,000
Natural Image Parsing	OCR-CC	1,000,000
Multi-Grained Text Localization	Word	114,344
	Phrase	113,130
	Line	378,432
	Block	394,094
Total		4,036,402

Figure 5: Detailed statistics of DocStruct4M.

Table 3: Different settings of OCR-free Visual Document Understanding models. ‘Open’ refers to whether all OCR learning data is open-source.

Model	Init	Resolution	OCR Learning				Open
			Text	Bbox	Size	Domain	
Donut (Kim et al., 2022)	-	2560x1920	✓	×	13M	Synthetic, Doc	✓
Pix2Struct (Lee et al., 2023)	-	2 ¹⁹ (shape variable)	✓	×	80M	Web	×
QwenVL (Bai et al., 2023b)	-	448x448	✓	×	24.8M	Synthetic, Doc, Web	×
Monkey (Li et al., 2023b)	QwenVL (Bai et al., 2023b)	896x896	×	×	-	-	-
URReader (Ye et al., 2023b)	Owl (Ye et al., 2023c)	224x224(x20 crops)	✓	×	0.1M	Doc, Table, Chart, Web, Natural	✓
DocPedia (Feng et al., 2023)	-	2560x2560	✓	✓	0.9M	Doc	×
CogAgent (Hong et al., 2023)	CogVLM (Wang et al., 2023a)	1120x1120	✓	✓	107M	Synthetic, Nature, Doc, Web	×
DocOwl 1.5	Owl2 (Ye et al., 2023d)	448x448(x9 crops)	✓	✓	4M	Doc, Table, Chart, Web, Natural	✓

Table 4: The detailed statistics of DocReason25K. The ‘Avg Length’ refers to the average token length of the answer.

Dataset	Image	Sample	Avg Length
DocVQA	1,491	5,119	79.2
InfoVQA	1,614	5,421	95.4
WTQ	850	5,994	77.7
VisualMRC	1,927	5,263	103.4
ChartQA	1,252	1,827	106.9
TextVQA	1,612	2,253	88.0
ALL	8,746	25,877	89.9

be horizontally parallel by vertical distance, and the ‘phrase’ is comprised of multiple adjacent words within the same line. The ‘block’ is a combination of multiple successive lines, ranging from 2 to half of the total lines. The text sequences of word-level and phrase-level question answering are much shorter than the other two. Therefore, in order to learn localization more efficiently, each word-level or phrase-level sample consists of up to 5 question-answer pairs for the same image. As for the representation of bounding boxes, we transfer each continuous value in the normalized bounding

box into a discrete position token, ranging from 0 to 999.

The bounding box annotation is necessary for constructing samples for Multi-grained Text Localization tasks. Therefore, we take DocVQA, InfoVQA, WTQ, TabFact, DeepForm, KLC, ChartQA, VisualMRC, and TextVQA (Singh et al., 2019) for this task, across domains of the document, table, chart, webpage, and natural image.

A.6 Detailed Statistic

DocStruct4M comprises 4M training samples to support the learning of Struct-aware Document Parsing, Table Parsing, Chart Parsing, Natural Image Parsing, and Multi-grained Text Localization. Fig. 5 shows the detailed statistics of DocStruct4M.

A.7 Data License

To build DocStruct4M, We collect 17 publicly available datasets, whose licenses range from MIT (Turski et al., 2023; Svetlichnaya, 2020; Chen et al., 2020; Methani et al., 2020; Yang et al., 2021), CC BY 4.0 (Singh et al., 2019), CC NC 4.0 (Kafle et al., 2018), Apache License 2.0 (Deng et al., 2022)

Table 5: Comparison with OCR-free methods on various types of text-rich image understanding tasks. The *underline* means the best performance among models with <10B parameters.

Model	Size	Doc VQA	Info VQA	Deep Form	KLC	WTQ	Tab Fact	Chart QA	Text VQA	Text Caps	Visual MRC
Previous SOTA	-	81.6	44.5	61.6	32.8	29.4	67.6	68.4	76.1	118.4	221.7
DocOwl-1.5	8.1B	81.6	50.4	68.8	37.9	39.8	80.4	70.5	<u>68.8</u>	132.0	239.5
DocOwl-1.5-Chat	8.1B	82.2	50.7	68.8	38.7	40.6	80.2	70.2	68.6	131.6	246.4

to GPL3.0 (Masry et al., 2022). There is no personally identifying information or offensive content in these datasets. We further produce struct-aware text sequence as the annotation and will release the DocStruct4M with the CC NC 4.0 for research purposes.

B Experimentns

B.1 Implementation Details

DocOwl 1.5 is initialized from mPLUG-Owl2 (Ye et al., 2023d), which utilizes the ViT/L-14 (Dosovitskiy et al., 2021) as the Visual Encoder and a 7B Large Language Model with the Modality Adaptive Module as the language decoder. According to the aspect ratio and resolution, each image is cropped into up to 9 sub-images with a fixed resolution of 448x448. Each sub-image is encoded to 1,024 features by the ViT/L-14 and then reduced to 256 features by the H-Reducer. The model is trained with 12,000 iterations on DocStruct4M, with the learning rate and batch size set as $1e-4$ and 1,024. It costs about 128 A100 days. During the Multi-task finetuning, the model is trained for 6,500 iterations with the batch size set as 256 and the learning rate set as $2e-5$. This further costs about 24 A100 days.

B.2 Baselines

Table 3 shows the comparison of DocOwl 1.5 and state-of-the-art OCR-free methods on model initialization, image resolution, and training data of OCR learning.

B.3 DocOwl 1.5-Chat

Existing benchmarks mainly evaluate the document understanding ability by answering the question with simple phrases and neglect detailed explanations. In this work, to better leverage the strong language reasoning ability of Large Language Models on Visual Document Understanding, we build a small instruction-tuning set with detailed explanations on text-rich image understanding, namely DocReason25K. Based on

raw questions from DocVQA (Mathew et al., 2021), InfoVQA (Mathew et al., 2022), WTQ (Pasupat and Liang, 2015), VisualMRC (Tanaka et al., 2021), ChartQA (Masry et al., 2022) and TextVQA (Singh et al., 2019), we collect detailed explanations with ChatGPT³. Text contents are dominant information on documents, tables or webpage screenshots. Therefore, for DocVQA, InfoVQA, WTQ, and VisualMRC, we take the structure-aware text sequence of the image as the input to gpt-3.5-turbo-0301 and prompt it to answer the question with simple answers and detailed explanations. As for ChartQA and TextVQA, we take the image as the input and utilize the gpt-4-vision-preview to answer the question with detailed explanations. In order to filter out samples where ChartGPT answers incorrectly, we further prompt gpt-3.5-turbo-0301 to judge whether the answer given by ChartGPT is consistent with the concise human-annotated ground-truth answer. Compared with raw questions in benchmark datasets, questions in DocReason25K are added with a prompt ‘Answer the question with detailed explanation’. Detailed statistics of DocReason25K are presented in Table 4. DocOwl 1.5-Chat is trained by combining downstream datasets with DocReason25K and performing multi-task tuning after Unified Structure Learning.

By removing the prompt of detailed explanation, DocOwl 1.5-Chat can also be evaluated on concise-answer benchmarks. As shown in Table 5, DocOwl 1.5-Chat achieves comparable or slightly better performance than DocOwl 1.5, showing that a small amount of detailed explanatory data may better help the model understand the semantics of text-rich images. Qualitative results of DocOwl 1.5-Chat on detailed explanation can be found in Appendix B.6.

³<https://openai.com/chatgpt>

Table 6: The detailed statistic of DocLocal4K.

Task	Text Granularity				Doc	Image Domain				
	Word	Phrase	Line	Block		Table	Chart	Web	Natural	
Text Recognition	622	499	522	482	1,004	491	229	267	134	
Text Grounding	595	542	503	485	1,011	524	240	242	108	

Table 7: Multi-grained text localization performance of models with different vision-to-text modules.

Module	Iter	Text Grounding					Text Recognition				
		Word	Phrase	Line	Block	ALL	Word	Phrase	Line	Block	ALL
Abstractor	1,800	10.92	25.83	34.59	87.01	37.69	30.68	28.58	40.12	32.73	33.03
H-Reducer(2x2)	1,800	14.19	34.87	43.94	89.07	43.94	37.20	38.33	48.68	41.99	41.55
H-Reducer(1x4)	1,800	17.82	39.30	53.28	90.52	48.28	39.60	41.84	55.37	49.84	46.66
H-Reducer(1x4)	12,000	70.42	76.38	85.88	91.34	80.38	70.10	67.86	73.88	70.70	70.63

B.4 Text Localization Evaluation

Besides proving the effectiveness of H-Reducer through downstream text-rich image understanding performance in Table 2, we further directly compare the text localization performance after the Unified Structure Learning to validate its superiority in preserving spatial features. We build a text localization evaluation set DocLocal4K with 4,250 samples balanced on 4 granularities and covering both text recognition and text grounding tasks. The detailed statistics of DocLocal4K are shown in Table 6. Considering that document images are much more diverse and complex than other images, there are more samples in this domain than others. The IOU@0.5 is used to evaluate the text grounding performance. As for text recognition, the word, phrase, line, and block granularity is evaluated with BLEU1, BLEU2, BLEU3, and BLEU4 (Papineni et al., 2002), respectively. As shown in Table 7, when trained with the same iterations, the H-Reducer achieves much better performance on both Text Recognition and Text Grounding tasks, showing that H-Reducer with the 1x4 merging shape helps the LLM better understand concrete positions in images.

B.5 Effectiveness of the Two-stage Training

Table 8 shows the performance comparison on DocVQA of two-stage training and one-stage joint training of the structure learning and downstream tasks. The epoch is gradually reduced because we didn’t observe performance improvements with more iterations. For joint training, the model improves significantly on DocVQA as the samples of Unified Structure Learning increase when it is be-

low 1M. However, as the Unified Structure Learning samples are further increased, the improvement of the model becomes subtle and its performance is not as good as the one using two-stage training. This shows that the two-stage training could better enhance basic text recognition and structure parsing abilities and is more beneficial and efficient for downstream document understanding.

B.6 Qualitative Results

Question Answering with Detailed Explanations.

Fig. 6 and Fig. 7 present qualitative results of detailed explanations. Through a small amount of reasoning training, DocOwl 1.5-Chat can well inherit the reasoning ability of LLM and provide detailed explanations about the answer. However, as presented in Fig. 7(c), like most general Multimodal large Language Models (Ye et al., 2023c,d; Bai et al., 2023b), DocOwl 1.5-Chat may also suffer from the hallucination problem in Visual Document Understanding. In this work, we mainly focus on enhancing the unified structure understanding ability of MLLMs and leave how to resolve the hallucination problem in OCR-free document understanding as future work.

Structure-aware Parsing. As shown in Fig. 8, DocOwl 1.5 could parse a document image by using line feeds and spaces to represent the structure of text contents. Besides parsing the whole document, as shown in Fig. 9, it could also parse texts from the middle of the image according to human instruction. Fig. 10 presents qualitative results of structure-aware table parsing through extended Markdown syntax on tables with cells spanning multiple columns or not. Furthermore, Fig. 11 shows some cases of parsing different types

Table 8: The comparison of two-stage training and one-stage joint training with increasing samples from DocStruct4M. For a fair comparison, the LLM is frozen for both two-stage and one-stage training. The batch size of one-stage training is always set as 256, the same as the Multi-task Tuning in two-stage training.

	One-Stage					Two-Stage
DocStruct4M samples	0.0M	0.5M	1.0M	2.0M	4.0M	4.0M
Benchmark samples	0.6M	0.6M	0.6M	0.6M	0.6M	0.6M
Epoch/iteration	7/18k	6/25k	6/37k	4/40k	3/54k	3/12k + 3/6.5k
Cost (A100 days)	60.0	83.3	123.3	133.3	180.0	144.8
DocVQA	72.8	75.5	78.6	78.8	78.9	79.9

of charts into Markdown codes, including vertical bar, horizontal bar, pie, and line charts. When all data points are presented in the chart, DocOwl 1.5 can accurately align statistic objects with corresponding numbers. It makes some mistakes in Fig. 11(d) because estimating the concrete numbers is quite challenging when no data points are provided. Finally, as shown in Fig. 12, DocOwl 1.5 can both describe the content of natural images and read scene texts.

Multi-grained Text Localization. Fig. 13 and Fig. 14 show qualitative results of text grounding and text recognition at granularities of word, phrase, line and block. The image domains range from documents, webpages, charts, and tables to natural images.

Additional Direct Account Information Required
Complete and Return by **Aug 28**

A. MacWilliams
 SIS # Direct Account Name Information Required
 158120 L. MacWilliams Discount Drug EFT or Non-EFT?
 30 (SM) Cases (15 Cases each (4 styles) To arrive 8/21/95 EFT \$
 100 Displays (#49890) To arrive 8/18/95 (Acco-Ship)
 30 (SM) Cases (15 Cases each (4 styles) To arrive 8/29/95 EFT \$
 100 Displays (#49890) To arrive 8/29/95 (Acco-Ship)

B. Cook
 SIS # Direct Account Name Information Required
 482000 P & C PO #
 22 (SM) Cases (8 Co. each style (4 styles) 2 To arrive 8/18/95 \$
 120 Displays (#49890) To arrive 8/21/95

L. MacWilliams
 SIS # Direct Account Name Information Required
 158120 L. MacWilliams Discount Drug PO #
 24 (6M) Cases (6 each style (4 styles) To arrive 8/18/95 \$
 2 Displays per store will be shipped directly to stores to arrive 8/21/95

Attention: J. Campbell

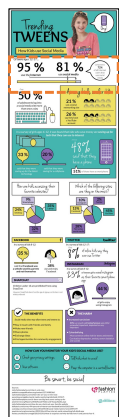
What is the SIS # of L. MacWilliams? Answer the question with detailed explanation.



DocOwl 1.5-Chat:

SIS # 158120.
 The SIS # of L. MacWilliams is 158120, as stated in the text under the section "L. MacWilliams SIS # Direct Account Name Information Required PO #".

(a)



What percentage of teenagers from the age group 12-17 didn't use the Internet? Answer the question with detailed explanation.



DocOwl 1.5-Chat:

According to the text, 95% of teenagers ages 12-17 use the Internet. Therefore, the percentage of teenagers from the age group 12-17 who didn't use the Internet would be $100\% - 95\% = 5\%$.

(b)

Date	Time	Opponent#	Rank#	Site	TV	Result	Attendance
September 1	2:30 PM	#9 (FCS) Northern Iowa*	#12	Camp Randall Stadium • Madison, WI	BTN	W 26-21	79,568
September 8	3:00 PM	at Oregon State*	#13	Reser Stadium • Corvallis, OR	FX	L 7-10	42,189
September 15	7:00 PM	Utah State*	#22	Camp Randall Stadium • Madison, WI	BTN ^[2]	W 16-14	79,332
September 22	11:00 AM	UTEP*	#24	Camp Randall Stadium • Madison, WI	ESPN2	W 37-26	79,806
September 29	7:00 PM	at #20 Nebraska	#23	Memorial Stadium • Lincoln, NE	ABC	L 27-30	85,962
October 6	2:30 PM	Illinois		Camp Randall Stadium • Madison, WI	ABC/ESPN2	W 31-14	80,096
October 13	11:00 AM	at Purdue		Ross-Ade Stadium • West Lafayette, IN	BTN	W 38-14	46,007
October 20	11:00 AM	Minnesota		Camp Randall Stadium • Madison, WI (Paul Bunyan's Axe)	ESPNU	W 38-13	80,587
October 27	2:30 PM	Michigan State†	#25	Camp Randall Stadium • Madison, WI	ABC/ESPN2	L 13-16 OT	80,538
November 10	11:00 AM	at Indiana		Memorial Stadium • Bloomington, IN	ESPN2	W 62-14	43,240
November 17	2:30 PM	Ohio State		Camp Randall Stadium • Madison, WI	ABC/ESPN2	L 14-21 OT	80,112
November 24	2:30 PM	at Penn State		Beaver Stadium • University Park, PA	ESPN2	L 21-24 OT	93,505
December 1	7:00 PM	vs. #14 Nebraska		Lucas Oil Stadium • Indianapolis, IN (Big Ten Championship Game)	FOX	W 70-31	41,260
January 1, 2013	4:10 PM	vs. #8 Stanford	#23	Rose Bowl • Pasadena, CA (Rose Bowl)	ESPN	L 14-20	93,259

*Non-conference game. †Homecoming. ²Rankings from Coaches Poll released prior to game. All times are in Central Time.

what is the last date and time on the table? Answer the question with detailed explanation.

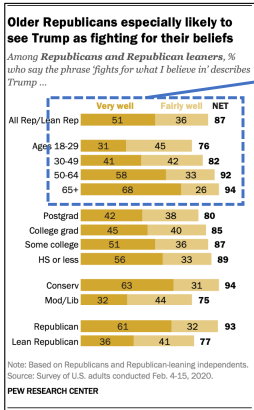


DocOwl 1.5-Chat:

The last date and time on the chart is January 1, 2013 at 4:10 PM. The chart lists all the games played by a certain team, along with the date, time, opponent, rank, site, and result. The last game listed on the chart is against Stanford, which took place on January 1, 2013 at 4:10 PM. Therefore, the last date and time on the chart is January 1, 2013 at 4:10 PM.

(c)

Figure 6: Qualitative results of question answering with detailed explanations. Some regions are enlarged for better visualization.



	Very well	Fairly well	NET
p/Lean Rep	51	36	87
Ages 18-29	31	45	76
30-49	41	42	82
50-64	58	33	92
65+	68	26	94



What's the value of the Very well bar in the 65+ age group? Answer the question with detailed explanation.



DocOwl 1.5-Chat:

68%.
 In the image, which appears to be a chart from the Pew Research Center, the bar representing "Very well" for the age group 65 and older is marked at 68%. This indicates that 68% of Republicans and Republican leaners in this age group believe that Trump is fighting for their beliefs "Very well."

(a)



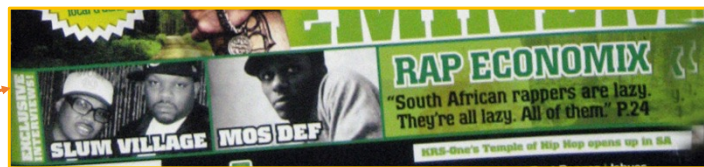
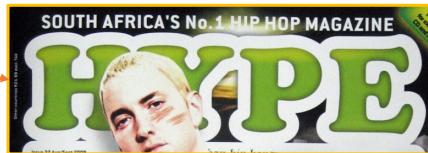
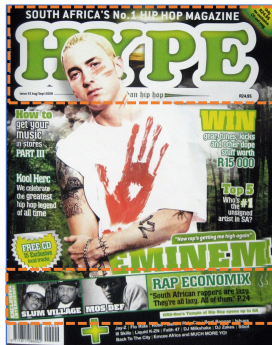
what is written on the blue shirt on the left? Answer the question with detailed explanation.



DocOwl 1.5-Chat:

Fly Emirates.
 The blue shirt on the left has the word "Fly Emirates" written across the chest. The logo is in white lettering against a blue background, which is the color scheme typically used for the airline's branding.

(b)



what nationality of rappers are lazy? Answer the question with detailed explanation.



DocOwl 1.5-Chat:

South African rappers are lazy. The text in the image states that "South African rappers are lazy." This is a reference to a quote from an article titled "Rap Economix" in the magazine Hype, which is South Africa's No. 1 hip hop magazine. The article discusses the perception that South African rappers are lazy and that they rely on the success of Kool & the Gang's "Stayin' Alive" instead of creating their own music. The article also mentions that this perception is changing, with new artists emerging and breaking away from the stereotype.

(c)

Figure 7: Qualitative results of question answering with detailed explanations. Hallucination in answers are marked in red.

Tropical Birding Trip Report NAMIBIA NOV DEC 2021	
Tinkling Cisticola	<i>Cisticola rufilatus</i>
Seen near Kayova.	
Chirping Cisticola	<i>Cisticola pipiens</i>
Seen at Kayova.	
Piping Cisticola	<i>Cisticola fulvicapilla</i>
Seen near Kayova.	
Zitting Cisticola	<i>Cisticola juncidis</i>
Seen at Etosha & Windhoek.	
Desert Cisticola	<i>Cisticola aridulus</i>
Seen at Etosha.	
Reed Warblers and Allies (Acrocephalidae)	
Icterine Warbler	<i>Hippolais icterina</i>
Seen at Etosha & Windhoek.	
Sedge Warbler	<i>Acrocephalus schoenobaenus</i>
Seen at Kayova.	
African Reed Warbler	<i>Acrocephalus baeticus</i>
Seen at Walvis Bay & Windhoek.	
Lesser Swamp Warbler	<i>Acrocephalus gracillirostris</i>
Seen at Kayova & Windhoek.	
Swallows (Hirundinidae)	
Plain Martin	<i>Riparia paludicola</i>
Seen at Kayova.	
Bank Swallow	<i>Riparia riparia</i>
Seen at Caprivi.	
www.tropicalbirding.com +1-409-515-9110 info@tropicalbirding.com	





DocOwl 1.5:



Extract words from the graphic.

```
<doc> Tropical Birding Trip Report NAMIBIA NOV DEC 2021
Tinkling Cisticola Cisticola rufilatus
Seen near Kayova.
Chirping Cisticola Cisticola pipiens
Seen at Kayova.
Piping Cisticola Cisticola fulvicapilla
Seen near Kayova.
Zitting Cisticola Cisticola juncidis
Seen at Etosha & Windhoek.
Desert Cisticola Cisticola aridulus
Seen at Etosha.
Reed Warblers and Allies (Acrocephalidae)
Icterine Warbler Hippolais icterina
Seen at Etosha & Windhoek.
Sedge Warbler Acrocephalus schoenobaenus
Seen at Kayova.
African Reed Warbler Acrocephalus baeticus
Seen at Walvis Bay & Windhoek.
Lesser Swamp Warbler Acrocephalus gracillirostris
Seen at Kayova & Windhoek.
Swallows (Hirundinidae)
Plain Martin Riparia paludicola
Seen at Kayova.
Bank Swallow Riparia riparia
Seen at Caprivi.
www.tropicalbirding.com +1-409-515-9110 info@tropicalbirding.com </doc>
```

Figure 8: A qualitative result of structure-aware document parsing.

5.4. Quality Criteria for a Blended Learning Platform

5.4.1. Introduction

Our objective is to create a list of absolutely necessary features for the eLearning platform (LMS) used in a blended learning course. These features are focused on the technical and pedagogical background as well as on the learner's situation. In addition, recommendations for useful features and tools are defined that used either inside the LCMS or as additional tools.

The target of the blended learning platform is the provision of high quality learning that will lead to the development of the 21st century skills. These are characterized, according to Kong et al (2014) by three emphases:

First emphasis is on skills development in both formal and informal learning contexts (Cox, 2013; Huang, Kinshuk, & Spector, 2013). The learners will be engaged in a seamless learning environment to coherently apply various generic skills for in-school teacher-led learning process initiated in digital classrooms and after-school learner-initiated learning process in social learning platforms/tools according to individual needs (Milrad, Wong, Sharples, Hwang, Looi, & Ogata, 2013; Otero, Milrad, Rogers, Santos, Verissimo, & Tones, 2011; Wong & Looi, 2011).

Second emphasis is on skills development through both individualized and collaborative learning approaches. On their own or with peers, learners take responsibilities to apply various generic skills to plan goals, implement tasks, monitor progress and evaluate outcomes in their learning process (Kicken, Brand-Gruwel, Merriënboer, & Slot, 2009; Norris & Soloway, 2009). The feedback for learners in a minimal but sufficient amount identifies individual needs and directions for future improvement (Caballero, van Riesen, Alvarez, Nussbaum, De Jong, 2014; Sims, 2003; Van Merriënboer, & Sluismans, 2009).

Third emphasis is on skills development supported by evidence of improvement and awareness of progress. The learning process in the e-learning environment can be designed in a range of activities in authentic learning contexts. Rich evidence of improvement and productive failure is collected by learners performance during the learning process; indications on applying 21st century skills for processing real-life information, reflecting on problem-solving ways, articulating tacit knowledge and




Figure 5-24: Learning with Tech trends³¹

³¹ Adapted from Olenka Villavicencio (2013): Learning with Tech trends (Retrieved from http://olevilla.blogspot.gr/2013_07_01_archive.html) and Planning for Personalization (retrieved from <http://education.vermont.gov/plb-working-group/main>).
A guidance to Blended Learning 79



DocOwl 1.5:

```
Continue reading the text: <doc> 5.4. Quality Criteria for a Blended Learning Platform
5.4.1. Introduction
Our objective is to create a list of absolutely necessary features for the eLearning platform (LMS) used
in a blended Learning course. These features are focused on the technical and pedagogical background
as well as on the learner's situation. In addition, recommends for useful features and tools are defined
that used either inside the LCMS or as additional tools.
The target of the blended learning platform is the provision of high quality learning that will lead to the
development of the 21st century skills. These are characterized, according to Kong et al.(2014) by three
emphases:
First emphasis is on skills development in both formal and informal learning contexts (Cox, 2013; Huang,
Kinshuk, & Spector, 2013). The learners will be engaged in a seamless learning environment to
coherently apply various generic skills for in-school teacher-led learning process initiated in digital
classrooms and after-school learner-initiated learning process in social learning platforms/tools
according to individual needs (Milrad, Wong, Sharples, Hwang, Looi, & Ogata, 2013; Otero, Milrad,
Rogers, Santos, Verissimo, & Tones, 2011; Wong & Looi, 2011).
Second emphasis is on skills development through both individualized and collaborative learning
approaches. On their own or with peers, learners take responsibilities to apply various generic skills to
plan goals, implement tasks, monitor progress and evaluate outcomes in their learning process (Kicken,
Brand-Gruwel, Merriënboer, & Slot, 2009; Norris & Soloway, 2009). The feedback for learners in a
minimal but sufficient amount identifies individual needs and directions for future improvement
(Caballero, van Riesen, Alvarez, Nussbaum, De Jong, 2014; Sims, 2003; Van Merriënboer, &
Sluismans,
2009).
Figure 5-24: Learning with Tech trends31
Third emphasis is on skills development supported by evidence of improvement and awareness of
progress. The learning process in the e-learning environment can be designed in a range of activities in
authentic learning contexts. Rich evidence of improvement and productive failure is collected by
learners performance during the learning process; indications on applying 21st century skills for
processing real-life information, reflecting on problem-solving ways, articulating tacit knowledge and
31 Adapted from Olenka Villavicencio (2013): Learning with Tech trends
(Retrieved from http://olevilla.blogspot.gr/2013_07_01_archive.html) and Planning for
Personalization (retrieved from http://education.vermont.gov/plb-working-group/main).
A guidance to Blended Learning 79 </doc>
```

Figure 9: A qualitative result of structure-aware document parsing from the middle of the image. The red dotted box is only used to mark the location of the answers for better visualization and is not included in the input image.

LAN exposure	N (%) [*]		Crude results (n = 61,129)		Adjusted results ^{**} (n = 47,817)	
	Cases	Referents	Odds ratios	95% CI	Odds ratios	95% CI
Low	27,121 (63)	10,970 (60)		Ref.		Ref.
Medium	5,974 (14)	2,623 (14)	1.13	(1.07-1.19)	1.06	(0.97-1.16)
High	9,659 (23)	4,782 (26)	1.22	(1.18-1.28)	1.12	(1.04-1.20) [†]



DocOwl 1.5:



Convert the figure into Markdown language.

```
<md> | <COLSPAN=2> N (%)** | <COLSPAN=2> Crude results (n = 61,129) | <COLSPAN=2> Adjusted results** (n = 47,817) |
| LAN exposure | Cases | Referents | Odds ratios | 95% CI | Odds ratios | 95% CI |
| --- | --- | --- | --- | --- | --- | --- |
| Low | 27,121 (63) | 10,970 (60) | | Ref. | | Ref. |
| Medium | 5,974 (14) | 2,623 (14) | 1.13 | (1.07-1.19) | 1.06 | (0.97-1.16) |
| High | 9,659 (23) | 4,782 (26) | 1.22 | (1.18-1.28) | 1.12 | (1.04-1.20)† | </md>
```

(a)

region	population centers	notes
Bačka	Novi Sad , Subotica	part of Vojvodina ; partially in Hungary (Bácska)
Gornji Breg	Gornji Breg	sub-region of Bačka
Potisje	Bečej , Senta	sub-region of both Bačka and Banat
Šajkaška	Temerin , Zabalj	sub-region of Bačka
Telečka	Vrbas , Kula	sub-region of Bačka
Banat	Zrenjanin , Pančevo	part of Vojvodina ; partially in Romania , Hungary (Bánság) and Central Serbia
Gornje Livade		sub-region of Banat
Ilandžanski Rit	Ilandža , Lokve	sub-region of Banat
Pančevački Rit	Borča , Krnjača	sub-region of Banat; part of Central Serbia
Pomorišje	Novi Kneževac , Srpski Krstur	today considered sub-region of Banat; partially in Hungary and Romania
Repšte	Knjačanin	sub-region of Banat
Veliki Rit	north-west of Vršac	sub-region of Banat
Srem	Belgrade , Sremska Mitrovica	part of Vojvodina; partially in Central Serbia and Croatia (Srijem)
Crni Lug	Progar , Boljevci	sub-region of Podlužje
Kupinovski Kut	Kupinovo	sub-region of Podlužje
Podlužje	Jakovo , Hrtkovci	sub-region of Srem; partially in Central Serbia



Parse the picture content behind with markdown code: <md> | region | population centers | notes |
| --- | --- | --- |
| Bačka | Novi Sad , Subotica | part of Vojvodina ; partially in Hungary (Bácska) |
| Gornji Breg | Gornji Breg | sub-region of Bačka |
| Potisje | Bečej , Senta | sub-region of both Bačka and Banat |
| Šajkaška | Temerin , Zabalj | sub-region of Bačka |
| Telečka | Vrbas , Kula | sub-region of Bačka |
| Banat | Zrenjanin , Pančevo | part of Vojvodina ; partially in Romania , Hungary (Bánság) and Central Serbia |
| Gornje Livade | | sub-region of Banat |
| Ilandžanski Rit | Ilandža , Lokve

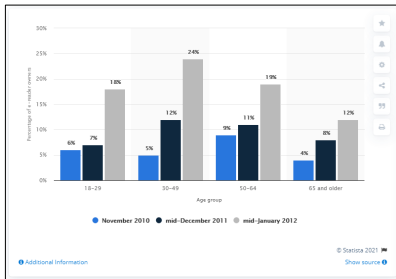


DocOwl 1.5:

```
| sub-region of Banat |
| Pančevački Rit | Borča , Knjača | sub-region of Banat; part of Central Serbia |
| Pomorišje | Novi Kneževac , Srpski Krstur | today considered sub-region of Banat; partially in Hungary and Romania |
| Repšte | Knjačanin | sub-region of Banat |
| Veliki Rit | north-west of Vršac | sub-region of Banat |
| Srem | Belgrade , Sremska Mitrovica | part of Vojvodina; partially in Central Serbia and Croatia (Srijem) |
| Crni Lug | Progar , Boljevci | sub-region of Podlužje |
| Kupinovski Kut | Kupinovo | sub-region of Podlužje |
| Podlužje | Jakovo , Hrtkovci | sub-region of Srem; partially in Central Serbia | </md>
```

(b)

Figure 10: Qualitative results of structure-aware table parsing on the table with cells spanning multiple columns (a) and structure-aware table parsing from the middle of the image (b). The red dotted box is only used to mark the location of the answers for better visualization and is not included in the input image.



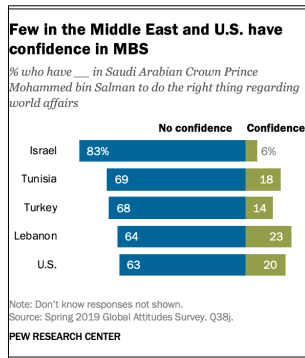
Translate the chart into Markdown format.



DocOwl 1.5:

```
<md> | Age group | November 2010 | mid-December 2011 | mid-January 2012 |
| --- | --- | --- | --- |
| 18-29 | 6% | 7% | 18% |
| 30-49 | 5% | 12% | 24% |
| 50-64 | 9% | 11% | 19% |
| 65 and older | 4% | 8% | 12% | </md>
```

(a)



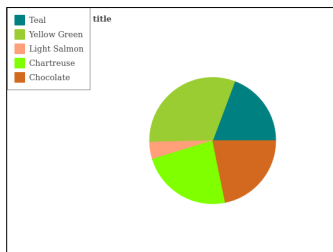
Render the figure in Markdown format.



DocOwl 1.5:

```
<md> | Entity | No confidence | Confidence |
| --- | --- | --- |
| U.S | 63 | 20 |
| Lebanon | 64 | 23 |
| Turkey | 68 | 14 |
| Tunisia | 69 | 18 |
| Israel | 83 | 6 | </md>
```

(b)



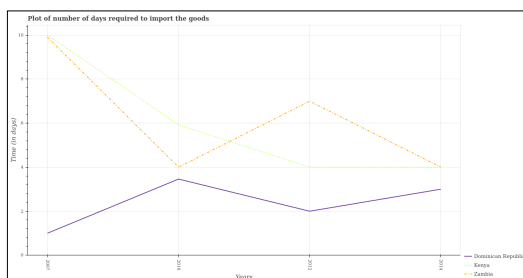
Convert the image into Markdown language.



DocOwl 1.5:

```
<md> | Teal | Yellow Green | Light Salmon | Chartreuse | Chocolate |
| --- | --- | --- | --- | --- |
| 19.5% | 31.1% | 4.17% | 24.2% | 21.9% | </md>
```

(c)



Parse the figure content behind with markdown code:

```
<md> | Years | Dominican Republic | Kenya | Zambia |
| --- | --- | --- | --- |
```



DocOwl 1.5:

```
| 2007 | 1 | 10 | 9.9 |
| 2010 | 3.47 | 4 | 6 |
| 2012 | 2 | 4 | 7 |
| 2014 | 3 | 4 | 4 | </md>
```

(d)

Figure 11: Qualitative results of structure-aware chart parsing on the charts of the vertical bar (a), horizontal bar (b), pie (c), and line (d). Incorrect words in the answer are marked in red.



(a alamy stock photo E7BDR1 www.alamy.com)



Detail the information and textual components depicted in the image.



DocOwl 1.5:

the main street in the town <ocr> a a a alamy NIMA RESTAURANT & BAR Special TOURS & DEL PSTAIRS PERIA DORA GENERAL SHOP SONAR ENTERPRISE a a a a a alamy photo stock E7BDR1 a www.alamy.com </ocr>

(a)



(PORTERHIRE)



Detail the information and textual components depicted in the image.

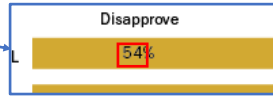
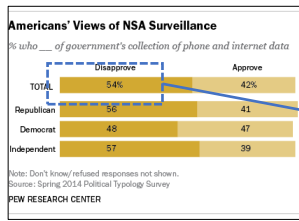


DocOwl 1.5:

a digger works to clear trees from the path of the railway line . <ocr> PORTERHIRE </ocr>

(b)

Figure 12: Qualitative results of natural image parsing. For better visualization, some regions are enlarged and labeled with corresponding scene texts. Incorrect words in the answer are marked in red.



Predict the bounding box of the text <ocr> 54 </ocr>



DocOwl 1.5:

<bbbox>329,341,371,381</bbbox>

(a)

Draw	Artist	Song	Points	Place
1	The Memories	"If It Means Losing You"	57	8th
2	Ann Breen	"Oh, Darling"	80	4th
3	Fran Meen	"Say That You Love Me"	66	6th
4	Dreams	"Sin Sin ("That's That")"	73	5th
5	Connor Stevens	"Count On Me"	88	3rd
6	Linda Martin and Friends	"All The People In The World"	105	2nd
7	Maggie Toal	"Feed Him With Love"	61	7th
8	Liam Reilly	"Somewhere In Europe"	130	1st



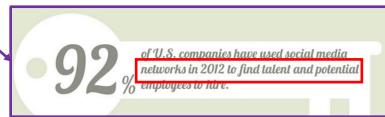
Give the bounding box of the text <ocr> Losing You" </ocr>



DocOwl 1.5:

<bbbox>564,214,696,272</bbbox>

(b)



Predict the bounding box of the text <ocr> networks in 2012 to find talent and potential </ocr>



DocOwl 1.5:

<bbbox>569,172,949,182</bbbox>

(c)

Statement of Council's and the Accounting Officer's responsibilities

Under Section 35 of the Nationality Act 1981, the Secretary of State for Culture, Media and Sport is responsible for the accounts for the Secretary of State for Culture, Media and Sport. The accounts are prepared on a consistent basis and are audited by the Comptroller and Auditor General.

In preparing the accounts, the Accounting Officer is required to comply with the requirements of the Government Financial Reporting Manual and in particular to:

Nicholas Capaldi
Accounting Officer
7 July 2017
Endorsed on behalf of Council

Phil George
7 July 2017
Endorsed on behalf of Council



Predict the bounding box of the text <ocr> particular to:
Nicholas Capaldi
• observe the Accounts Direction issued by the Accounting Officer Secretary of State for Culture, Media and Sport, including the relevant accounting and disclosure requirements, and apply suitable accounting policies on a consistent basis; Endorsed on behalf of Council:
• make judgements and estimates on a reasonable basis;
• state whether applicable accounting standards </ocr>



DocOwl 1.5:

<bbbox>71,437,761,658</bbbox>

In preparing the accounts, the Accounting Officer is required to comply with the requirements of the Government Financial Reporting Manual and in particular to:

Nicholas Capaldi
Accounting Officer
7 July 2017
Endorsed on behalf of Council

Phil George
7 July 2017
Endorsed on behalf of Council

(d)

Figure 13: Qualitative results of Multi-grained Text Grounding. Some regions are enlarged for better visualization. Bounding boxes predicted by DocOwl 1.5 are drawn in images as solid red boxes.

Detect the text in the bounding box
 <bbox>859,2,962,36</bbox>

DocOwl 1.5:

<ocr> Emirates </ocr>

(a)

Detect the text in the bounding box
 <bbox>128,430,256,455</bbox>

DocOwl 1.5:

<ocr> Produced capital </ocr>

(b)

Detect the text in the bounding box
 <bbox>76,674,833,682</bbox>

DocOwl 1.5:

<ocr> SPOT 1:30:00AM- 2:35:00APMRPer week (M),Tu,We 02/14/20 1:45AM (Mo) 01:00 DATELINE 1:35A MRBPK6033EH \$150.00 </ocr>

(c)

SPOT	Contract Line Remarks	CDR	Per week (M),Tu,We	02/10/20 2:16AM (Mo)	0L:00	DATELINE 1:35A	MRBPK6033EH	\$/S:00
SPOT	1:35:00AM- 2:35:00AM	Per week (M),Tu,We	02/11/20 1:53AM (Tu)	0L:00	DATELINE 1:35A	MRBPK6033EH	\$150.00	
SPOT	1:35:00AM- 2:35:00AM	Per week (M),Tu,We	02/12/20 1:45AM (We)	0L:00	DATELINE 1:35A	MRBPK6033EH	\$150.00	
SPOT	1:35:00AM- 2:35:00AM	Per week (M),Tu,We	02/12/20 2:04AM (We)	0L:00	DATELINE 1:35A	MRBPK6033EH	\$150.00	

(d)

Language Games [edit]

Wittgenstein believed that every word we speak is at part of a language game. For Wittgenstein language games seem similar to an inside joke. You would only get the joke if you were in on the joke. This is similar to language, you will only understand the language being used if you are familiar with the language. That is why Wittgenstein believed that religious language is meaningless, but only to the religious believers. They are all part of a group that regularly use that language, which has deep meaning to them. Non-believers would not find that religious language is meaningful, because we are not involved in that game. Wittgenstein refers to words as 'tools' because we use them to build our houses and as 'toys' because we play games with them.

Ludwig Wittgenstein [edit]

Born 1889, died 1951. An Austrian born Cambridge scholar, Wittgenstein was involved in his early years with the Logical Positivists of the Vienna Circle. Wittgenstein left Philosophy for a while but was persuaded to return in the 1920s. Wittgenstein's return prompted a different way of looking at religious language, he no longer sought to prove "whether" language refers, but looked at the way in which language refers to things. He decided that the best way to determine a word's meaning was to look at how it is used, and not to come up with a theoretical definition. Wittgenstein said this was about: "Getting back to the rough ground".

Detect the text in the bounding box <bbox>14,532,981,850</bbox>

DocOwl 1.5:

<ocr> Ludwig Wittgenstein [edit] \n Born 1889, died 1951. An Austrian born Cambridge scholar, Wittgenstein was involved in his early years with the Logical Positivists of the Vienna Circle. \n Wittgenstein left Philosophy for a while but was persuaded to return in the 1920s. Wittgenstein's return prompted a different way of looking at religious \n language, he no longer sought to prove "whether" language refers, but looked at the way in which language refers to things. \n He decided that the best way to determine a word's meaning was to look at how it is used, and not to come up with a theoretical definition. Wittgenstein \n said this was about: "Getting back to the rough ground". </ocr>

Figure 14: Qualitative results of Multi-grained Text Recognition. Some regions are enlarged for better visualization. Input bounding boxes are drawn in images as solid blue boxes. Incorrect words in answers are marked in red.