DocGraphLM: Documental Graph Language Model for Information Extraction

Dongsheng Wang
JPMorgan AI Research
London, UK
dongsheng.wang@jpmchase.com

Zhiqiang Ma
JPMorgan AI Research
New York, New York, USA
zhiqiang.ma@jpmchase.com

Armineh Nourbakhsh
JPMorgan AI Research
New York, New York, USA
armineh.nourbakhsh@jpmchase.com

Kang Gu
Dartmouth College
Hanover, New Hampshire, USA
Kang.Gu.GR@dartmouth.edu

Sameena Shah
JPMorgan AI Research
New York, New York, USA
sameena.shah@jpmchase.com

ABSTRACT

Advances in Visually Rich Document Understanding (VrDU) have enabled information extraction and question answering over documents with complex layouts. Two tropes of architectures have emerged—transformer-based models inspired by LLMs, and Graph Neural Networks. In this paper, we introduce DocGraphLM, a novel framework that combines pre-trained language models with graph semantics. To achieve this, we propose 1) a joint encoder architecture to represent documents, and 2) a novel link prediction approach to reconstruct document graphs. DocGraphLM predicts both directions and distances between nodes using a convergent joint loss function that prioritizes neighborhood restoration and downweighs distant node detection. Our experiments on three SotA datasets show consistent improvement on IE and QA tasks with the adoption of graph features. Moreover, we report that adopting the graph features accelerates convergence in the learning process druing training, despite being solely constructed through link prediction.

CCS CONCEPTS

• **Information systems** → *Document structure*; **Language models**; **Information extraction**.

KEYWORDS

language model, graph neural network, information extraction, visual document understanding

1 INTRODUCTION

Information extraction from visually-rich documents (VrDs), such as business forms, receipts, and invoices in the format of PDF or image has gained recent traction. Tasks such as field identification and extraction and entity linkage are crucial to digitizing VrDs and building information retrieval systems on the data. Tasks that require complex reasoning such as Visual Question Answering over documents require modeling the spatial, visual, and semantic signals in VrDs. Therefore, VrD Understanding is concerned with modeling the multi-modal content in image documents. Previous research has explored the use of encoding text, layout, and image features in a layout language model or multi-modal setting to improve downstream tasks. For example, LayoutLM and its variants [7, 23, 24] use image and layout information to enhance the representation of text, thereby improving performance on various tasks. However, models using Transformer mechanisms pose a challenge to representing spatially distant semantics, such as table cells far from their headers or contents across line breaks. In light of these limitations, a few studies [25, 27] have proposed using graph neural networks (GNNs) to model relationships and structures between text tokens or segments in documents. Although these models alone still underperform layout language models, they demonstrate the potential of incorporating additional structured information to improve document representation.

Motivated by this, we introduce a novel framework called DocGraphLM that integrates document graph semantics and the semantics derived from pre-trained language models to improve document representation. As depicted in Figure 1, the input to our model is embeddings of tokens, positions, and bounding boxes, which form the foundation of the document representation. To reconstruct the document graph, we propose a novel link prediction approach that predicts directions and distances between nodes by using a joint loss function, which balances the classification and regression loss. Additionally, the loss encourages close neighborhood restoration while downgrading detections on farther nodes. This is achieved by normalizing the distance through logarithmic transformation, treating nodes separated by a specific order-of-magnitude distance as semantically equidistant.

Our experiments on multiple datasets including FUNSD, CORD, and DocVQA, show the superiority of the model in a consistent manner. Furthermore, the incorporation of graph features is found
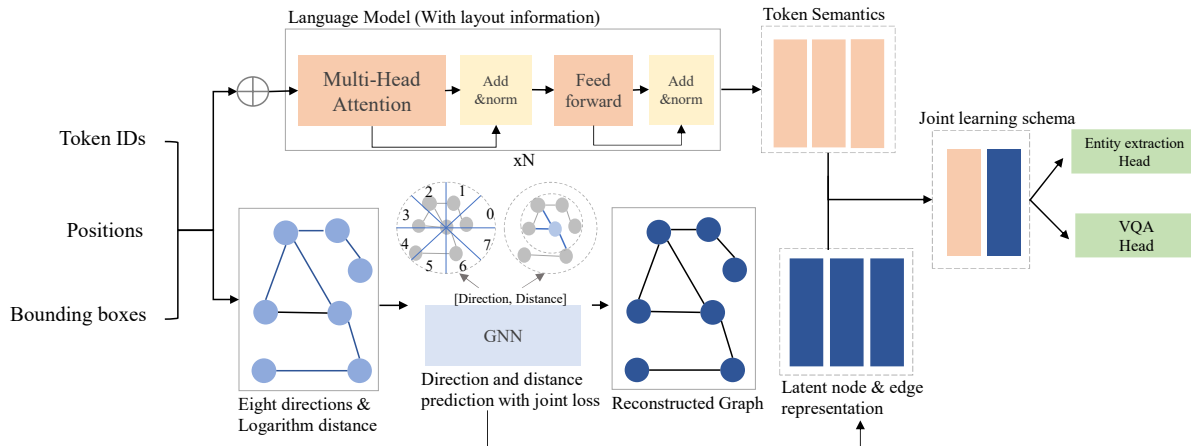
Figure 1: The model architecture of DocGraphLM.

to accelerate the learning process. We highlight the main contributions of our work as follows:

- we propose a novel architecture that integrates a graph neural network with pre-trained language model to enhance document representation;
- we introduce a link prediction approach to document graph reconstruction, and a joint loss function that emphasizes restoration on nearby neighbor nodes;
- lastly, the proposed graph neural features result in a consistent improvement in performance and faster convergence.

2 RELATED WORK

Transformer-based architectures have been successfully applied to layout understanding tasks, surpassing previous state-of-the-art (SotA) results [3, 13, 14, 16, 21, 22]. Studies such as LayoutLM [23] and LayoutLMv2 [24] fuse text embeddings with visual features using a region proposal network, allowing the models to be trained on objectives such as Masked Visual Language Model (MVLM) and spatial aware attention, resulting in improved performance on complex tasks such as VQA and form understanding. TILT [9] augments the attention by adding bias to capture relative 2-D positions, which has shown excellent performance on DocVQA leaderboard. StructuralLM [12] makes the most of the interactions of cells where each cell shares the same bounding boxes.

The use of GNNs [20] to represent documents allows information to propagate more flexibly. In GNN-based VrDU models, documents are often represented as graphs of tokens and/or sentences, and edges represent spatial relationships among them, e.g. capturing K-Nearest Neighbours. GNN-based models can be used for various document-grounded tasks such as text classification [25, 27] or key information extraction [2, 26]. However, their performance still lags behind that of layout language models. This is because graph representation alone is insufficient to capture the rich semantics of a document. In cases where GNN-based models substantially outperform layout language models, they are often larger and focused on specific tasks [10]. In this paper, we propose a framework that combines the rich semantics of layout language models with the robust structural signal captured by GNN models. We demonstrate how the addition of graph semantics can enhance the performance of layout language models on IE and QA tasks, and improve model convergence.

3 DOCGRAPHLM: DOCUMENT GRAPH LANGUAGE MODEL

3.1 Representing document as graph

In GNN, a graph consists of nodes and edges. In the context of representing document as graph, the nodes represent text segments (i.e. groups of adjacent words) and the relationships between them are represented as edges. Text segments from image documents can be obtained through Optical Character Recognition tools, which often capture the tokens as bounding boxes of various sizes.

To generate the edges between nodes, we adopt a novel heuristic named Direction Line-of-sight (D-LoS), instead of the commonly used K-nearest-neighbours (KNN) [19] or β-skeleton approach [11]. The KNN approach may result in dense, irrelevant rows or columns being treated as neighbours, ignoring the fact that some key-value pairs in a form can be farther apart nodes. To address this, we adopt the D-LoS approach, where we divide the 360-degree horizon surrounding a source node into eight discrete 45-degree sectors, and we determine the nearest node with respect to the source node within each sector. These eight sectors define eight directions with respect to the source node. This definition is inspired by the pre-training task reported in StrucTexT [14] which applies this approach to construct its graph representation.

Node representation. A node has two features — text semantics and node size. The text semantics can be obtained through token embeddings (e.g. from language models), while the node size is expressed by its dimensions on x and y coordinates, mathematically $M = \text{emb}([width, height])$ were $width = x_2 - x_1$ and $height = y_2 - y_1$, given that (x_1, y_1) and (x_2, y_2) are the coordinates of top left corner and bottom right corner of the segment bounding box. Intuitively, the node size is a significant indicator because it helps differentiate font size and potentially the semantic role of the segment, e.g., title, caption, and body. Thus, we denote a node input

as $E_u = \text{emb}(T_u) \oplus M_u$, where $u = \{1, 2, ..., N\}$ indicates the uth node in a document and T_u stands for the texts inside the node u.

We learn the node representation by reconstructing the document graph using GNN, expressed as $h_u^G = \text{GNN}(E_u)$. Details on learning h_u^G are described in Section 3.2.

Edge representation. To express the relationships between two nodes, we use their polar features, including relative distance and direction (one of eight possibilities). We compute the shortest Euclidean distance, d, between the two bounding boxes. To reduce the impact of distant nodes that may be less semantically relevant to the source node, we apply a distance smoothing technique with log transformation denoted as $e_{\text{dis}} = \log(d + 1)$. The relative direction $e_{\text{dir}} \in \{0, ..., 7\}$ for a pair of nodes is obtained from D-LoS. We define a linkage, denoted as $e_p = [e_{\text{dis}}, e_{\text{dir}}]$, to reconstruct the document graph in section 3.2.

3.2 Reconstructing graph by link prediction

We predict two key attributes of the linkages e_p to reconstruct the graph and frame the process as a multi-task learning problem.

The input to the GNN is the encoded node representations, and the representation is passed through the message passing mechanism on GNN, specifically:

$$h_u^{G,l+1} := \text{aggregate}(h_v^{G,l}, \forall v \in \mathcal{N}(u)), \qquad (1)$$

where l is the layer of neighbors, $\mathcal{N}(u)$ denotes the set of neighbors of node u, and $\text{aggregate}(\cdot)$ is an aggregation function that updates the node representation.

We jointly train the GNN on two tasks — predicting the distance and direction between nodes — to learn the node representation. For distance prediction, we define a regression head $\hat{y}_{u,v}^e$, which generates a scalar value through the dot-product of two node vectors, and uses a linear activation, as presented in Equation 2.

$$\hat{y}_{u,v}^e = Linear((h_u^G)^\top \times h_v^G) \qquad (2)$$

For direction prediction, we define a classification head $\hat{y}_{u,v}^d$ that assigns one of eight directions to each edge based on the element-wise product between two nodes, expressed as follows:

$$\hat{y}_{u,v}^d = \sigma((h_u^G \odot h_v^G) \times W) \qquad (3)$$

where $h_u^G \odot h_v^G$ is an element-wise product between two nodes and W is the learnable weight for the product vector. σ is a non-linear activation function.

We use MSE loss for distance regression and cross-entropy for the direction classification, respectively. Then, the joint loss is:

$$\begin{aligned}
loss = \sum_{(u,v) \in \text{batch}} [& (\lambda \cdot \text{loss}^{\text{MSE}}(\hat{y}_{u,v}^e, y_{u,v}^e) \\
& + (1 - \lambda) \cdot \text{loss}^{\text{CE}}(\hat{y}_{u,v}^d, y_{u,v}^d)] \cdot (1 - r_{u,v})
\end{aligned} \qquad (4)$$

where λ is a tunable hyper-parameter that balances the weights of the two losses, and $r_{u,v}$ is the normalization of the distance e_{dis}, constrained to the interval $[0, 1]$, so that the value of $1 - r_{u,v}$ downweights distant segments and favors nearby segments.

3.3 Joint representation

The joint node representation, h_u^C, is a combination of the language model representation h_u^L and the GNN representation h_u^G through

Table 1: Statistics of visual document datasets. The differences between DocVQA and DocVQA† is introduced in Section 4.1.

Dataset	No. labels	No. train	No. val	No. test
FUNSD	4	149	-	50
CORD	30	800	100	100
DocVQA	-	39,000	5,000	5,000
DocVQA†	-	32,553	4,400	5,000

an aggregation function f (e.g., concatenation, mean, or sum) represented as $h_u^C = f(h_u^L, h_u^G)$. In this work, we operationalize the aggregation function f with concatenation at the token level. The introduced node representations can be utilized as input for other models to facilitate downstream tasks, e.g., IE_Head(h_u^C) for entity extraction and QA_Head(h_u^C) for visual question answering task.

4 EXPERIMENTS

4.1 Datasets and baselines

We evaluate our models on two information extraction tasks across three commonly used datasets: FUNSD [8], CORD [18], and DocVQA [17]. FUNSD and CORD focus on entity-level extraction, while DocVQA concentrates on identifying answer spans in image documents in a question-answering task. Dataset statistics are shown in Table 1. Please refer to the citations for more details.

It is noted that the OCR files provided in DocVQA[1] contain a small number of imperfect OCR outputs, e.g., text misalignment and missing texts, which leads to failures in identifying the answers. We can only use 32,553 samples for training and 4,400 samples for validation. We denote the modified dataset as $DocVQA^\dagger$. In the interest of ensuring fair comparison in our experiments, we have maintained the use of the OCR outputs from the dataset.

As our baselines, we employ the SotA models that make use of different features, including RoBERTa [15], BROS [6], DocFormer-base [1], StructuralLM [12], LayoutLM [23], LayoutLMv3 [7] and Doc2Graph [4]. RoBERTa is transformer model without any layout or image features, BROS and StructuralLM adopt layout information solely, DocFormer and LayoutLMv3 utilizes both layout and image features, and Doc2Graph soly relies on document graph features.

4.2 Experimental setup

For FUNSD and CORD, we adopt the following training hyper-parameters: epoch = 20, learning rate = 5e-5, and batch size = 6, and trained our model on a single NVIDIA T4 Tensor Core GPU. For DocVQA, we apply the following training hyper-parameters: epoch = 5, learning rate = 5e-5, and batch size = 4.

We adopt GraphSage [5] as our GNN model, as it has been proven effective in document graph features[4]. For graph reconstruction, we set a constant value λ=0.5 throughout the experiment.

4.3 Results

The performance of DocGraphLM and other models on the FUNSD dataset are presented in Table 2. Our model reaches the best F1

---

[1]https://www.docvqa.org/

Table 2: Model performance comparison on FUNSD.

Model	F1	Precision	Recall
RoBERTa-base	65.37	61.17	70.20
Doc2Graph$^\diamond$[4]	82.25	-	-
StructuralLM_large$^\diamond$[12]	85.14	83.52	86.81
LayoutLM-base$^\diamond$[23]	78.66	75.97	81.55
LayoutLMv3-base[7]	88.16	86.70	87.7
BROS$^\diamond$[6]	83.05	81.16	85.02
DocFormer-base$^\diamond$[1]	83.34	80.76	86.09
DocGraphLM (RoBERTa-base)	67.03 (↑1.66)	62.92	70.0
DocGraphLM (LLMv3-base)	**88.77(↑0.61)**	**87.44**	**90.15**

Table 3: Model performance comparison CORD.

Model	F1	Precision	Recall
RoBERTa-base	48.99	42.77	57.34
LayoutLM-base$^\diamond$	94.80	95.03	94.58
LayoutLMv3-base	95.59	95.31	95.88
BROS$^\diamond$	95.36	95.58	95.14
DocFormer-base$^\diamond$	96.33	96.52	96.14
DocGraphLM (RoBERTa-base)	51.25 (↑2.26)	45.45	58.76
DocGraphLM (LayoutLMv3-base)	**96.93 (↑1.62)**	**96.86**	**97.01**

score at 88.77, achieved when it is paired with the LayoutLMv3-base model. On the other hand, RoBERTa-base (which does not leverage layout features) has the lowest F1 score of 65.37, but combining it with DocGraphLM results in a 1.66 point improvement. Please note scores with $\diamond$ are reported in the corresponding citations. The same notation applies to other tables.

For the CORD dataset, the performance comparisons are shown in Table 3, and the best performance is achieved by DocGraphLM (LayoutLMv3-base) with an F1 score of 96.93, followed closely by BROS. Similarly, even though RoBERTa-base alone achieves a much lower score, DocGraphLM (RoBERTa-base) increases the F1 score by 2.26 points.

Table 4 shows the model performance on the DocVQA test dataset. The performance scores are obtained by submitting our model output to the DocVQA leaderboard[2], as ground-truth answers are not provided to the public. Besides the overall score, the model's performances on sub-category tasks are also reported. DocGraphLM (with LayoutLMv3-base) outperforms others in almost every aspect except pure text semantics, which shows the model's ability to model multi-modal semnatics effectively. The table presents strong evidence towards the efficiency of DocGraphLM in improving document representations, when layout language models are augmented with our approach.

The superior performance across various datasets indicates that using the graph representation proposed in DocGraphLM leads to consistent improvements. A p-value less than 0.05 was received when comparing the models' performance across these datasets, indicating a statistically significant improvement from our model.

Table 4: Model performance comparison on DocVQA testing dataset. Scores are from DocVQA leaderboard.

Model	Score	Form	Table	Text
RoBERTa_base	60.40	71.75	54.23	61.35
LayoutLMv3_base	67.80	77.84	67.58	**70.55**
DocGraphLM (LayoutLMv3-base)	**69.84 (↑2.04)**	**79.73**	**68.48**	63.23

4.4 Impact on convergence

We also observed that the training convergence speed is often faster when supplementing the graph features than vanilla LayoutLM (V1 and V3 base models). For example, Figure 2 illustrates that the F1 score improves in a faster convergence rate within the first four epochs, when testing on the CORD dataset. This could be due to the graph features allowing the transformer to focus more on the nearby neighbours, which eventually results in a more effective information propagation process.
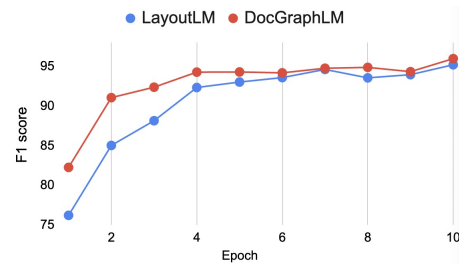
Figure 2: Model convergence speed comparison on CORD. The curves are generated from averaging over ten trials.

5 CONCLUSION AND FUTURE WORK

This paper presents a novel DocGraphLM framework incorporating graph semantics with pre-trained language models to improve document representation for VrDs. The proposed linkage prediction method reconstructs the distance and direction between nodes, increasingly down-weighting more distant linkages. Our experiments on multiple downstream tasks on various datasets show enhanced performance over LM-only baseline. Additionally, introducing the graph features accelerates the learning process. As a future direction, we plan to incorporate different pre-training techniques for different document segments. We will also examine the effect of different linkage representations for graph reconstruction.

REFERENCES

[1] Srikar Appalaraju, Bhavan Jasani, Bhargava Urala Kota, Yusheng Xie, and R Manmatha. 2021. Docformer: End-to-end transformer for document understanding. In *Proceedings of the IEEE/CVF international conference on computer vision*. 993–1003.

[2] Brian L. Davis, Bryan S. Morse, Brian L. Price, Chris Tensmeyer, and Curtis Wigington. 2021. Visual FUDGE: Form Understanding via Dynamic Graph Editing. *CoRR* abs/2105.08194 (2021). arXiv:2105.08194 https://arxiv.org/abs/2105.08194

[3] Lukasz Garncarek, Rafal Powalski, Tomasz Stanislawek, Bartosz Topolski, Piotr Halama, and Filip Gralinski. 2020. LAMBERT: Layout-Aware language Modeling using BERT for information extraction. *CoRR* abs/2002.08087 (2020). arXiv:2002.08087 https://arxiv.org/abs/2002.08087

[4] Andrea Gemelli, Sanket Biswas, Enrico Civitelli, Josep Lladós, and Simone Marinai. 2022. Doc2Graph: a Task Agnostic Document Understanding Framework based on Graph Neural Networks. *arXiv preprint arXiv:2208.11168* (2022).

[5] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. *Advances in neural information processing systems* 30 (2017).

[6] Teakgyu Hong, DongHyun Kim, Mingi Ji, Wonseok Hwang, Daehyun Nam, and Sungrae Park. 2020. BROS: a pre-trained language model for understanding texts in document. (2020).

[7] 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 *Proceedings of the 30th ACM International Conference on Multimedia*. 4083–4091.

[8] Guillaume Jaume, Hazim Kemal Ekenel, and Jean-Philippe Thiran. 2019. Funsd: A dataset for form understanding in noisy scanned documents. In *2019 International Conference on Document Analysis and Recognition Workshops (ICDARW)*, Vol. 2. IEEE, 1–6.

[9] Rafa l Powalski, Lukasz Borchmann, and Dawid Jurkiewicz. 2021. Going Full-TILT Boogie on Document Understanding with Text-Image-Layout Transformer. *arXiv preprint arXiv:2102.09550* (2021).

[10] Chen-Yu Lee, Chun-Liang Li, Timothy Dozat, Vincent Perot, Guolong Su, Nan Hua, Joshua Ainslie, Renshen Wang, Yasuhisa Fujii, and Tomas Pfister. 2022. Formnet: Structural encoding beyond sequential modeling in form document information extraction. *arXiv preprint arXiv:2203.08411* (2022).

[11] Chen-Yu Lee, Chun-Liang Li, Chu Wang, Renshen Wang, Yasuhisa Fujii, Siyang Qin, Ashok Popat, and Tomas Pfister. 2021. ROPE: Reading Order Equivariant Positional Encoding for Graph-based Document Information Extraction. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*. Association for Computational Linguistics, Online, 314–321. https://doi.org/10.18653/v1/2021.acl-short.41

[12] Chenliang Li, Bin Bi, Ming Yan, Wei Wang, Songfang Huang, Fei Huang, and Luo Si. 2021. Structurallm: Structural pre-training for form understanding. *arXiv preprint arXiv:2105.11210* (2021).

[13] Peizhao Li, Jiuxiang Gu, Jason Kuen, Vlad I Morariu, Handong Zhao, Rajiv Jain, Varun Manjunatha, and Hongfu Liu. 2021. Selfdoc: Self-supervised document representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5652–5660.

[14] Yulin Li, Yuxi Qian, Yuechen Yu, Xiameng Qin, Chengquan Zhang, Yan Liu, Kun Yao, Junyu Han, Jingtuo Liu, and Errui Ding. 2021. Structext: Structured text understanding with multi-modal transformers. In *Proceedings of the 29th ACM International Conference on Multimedia*. 1912–1920.

[15] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *ArXiv* abs/1907.11692 (2019).

[16] Bodhisattwa Prasad Majumder, Navneet Potti, Sandeep Tata, James Bradley Wendt, Qi Zhao, and Marc Najork. 2020. Representation learning for information extraction from form-like documents. In *proceedings of the 58th annual meeting of the Association for Computational Linguistics*. 6495–6504.

[17] Minesh Mathew, Dimosthenis Karatzas, and CV Jawahar. 2021. Docvqa: A dataset for vqa on document images. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*. 2200–2209.

[18] Seunghyun Park, Seung Shin, Bado Lee, Junyeop Lee, Jaeheung Surh, Minjoon Seo, and Hwalsuk Lee. 2019. CORD: a consolidated receipt dataset for post-OCR parsing. In *Workshop on Document Intelligence at NeurIPS 2019*.

[19] Yujie Qian, Enrico Santus, Zhijing Jin, Jiang Guo, and Regina Barzilay. 2019. GraphIE: A Graph-Based Framework for Information Extraction. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 751–761. https://doi.org/10.18653/v1/N19-1082

[20] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. 2008. The graph neural network model. *IEEE transactions on neural networks* 20, 1 (2008), 61–80.

[21] Zilong Wang, Yiheng Xu, Lei Cui, Jingbo Shang, and Furu Wei. 2021. Layoutreader: Pre-training of text and layout for reading order detection. *arXiv preprint arXiv:2108.11591* (2021).

[22] Zilong Wang, Mingjie Zhan, Xuebo Liu, and Ding Liang. 2020. Docstruct: A multimodal method to extract hierarchy structure in document for general form understanding. *arXiv preprint arXiv:2010.11685* (2020).

[23] Yiheng Xu, Minghao Li, Lei Cui, Shaohan Huang, Furu Wei, and Ming Zhou. 2020. Layoutlm: Pre-training of text and layout for document image understanding. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 1192–1200.

[24] Yang Xu, Yiheng Xu, Tengchao Lv, Lei Cui, Furu Wei, Guoxin Wang, Yijuan Lu, Dinei Florencio, Cha Zhang, Wanxiang Che, et al. 2020. Layoutlmv2: Multimodal pre-training for visually-rich document understanding. *arXiv preprint arXiv:2012.14740* (2020).

[25] Liang Yao, Chengsheng Mao, and Yuan Luo. 2019. Graph convolutional networks for text classification. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 7370–7377.

[26] Wenwen Yu, Ning Lu, Xianbiao Qi, Ping Gong, and Rong Xiao. 2021. PICK: processing key information extraction from documents using improved graph learning-convolutional networks. In *2020 25th International Conference on Pattern Recognition (ICPR)*. IEEE, 4363–4370.

[27] Yufeng Zhang, Xueli Yu, Zeyu Cui, Shu Wu, Zhongzhen Wen, and Liang Wang. 2020. Every document owns its structure: Inductive text classification via graph neural networks. *arXiv preprint arXiv:2004.13826* (2020).