

Post-OCR Paragraph Recognition by Graph Convolutional Networks

Renshen Wang, Yasuhisa Fujii, Ashok C. Popat
Google Research

{rewang, yasuhisaf, popat}@google.com

Abstract

We propose a new approach for paragraph recognition in document images by spatial graph convolutional networks (GCN) applied on OCR text boxes. Two steps, namely line splitting and line clustering, are performed to extract paragraphs from the lines in OCR results. Each step uses a β -skeleton graph constructed from bounding boxes, where the graph edges provide efficient support for graph convolution operations. With pure layout input features, the GCN model size is 3~4 orders of magnitude smaller compared to R-CNN based models, while achieving comparable or better accuracies on PubLayNet and other datasets. Furthermore, the GCN models show good generalization from synthetic training data to real-world images, and good adaptivity for variable document styles.

1. Introduction

Document image understanding is a task to recognize, structure, and understand the contents of document images, and is a key technology to digitally process and consume such images, which are ubiquitous and can be found in numerous applications. Document image understanding enables the conversion of such documents into a digital format with rich structure and semantic information and makes them available for subsequent information tasks.

A document can be represented by its semantic structure and physical structure [13]. The task to recover the semantic structure is called logical layout analysis [5] or semantic structure extraction [35] while the task to recover the physical structure is called geometric (physical, or structural) layout analysis [5]. These tasks are critical subproblems of document image understanding.

A paragraph is a semantic unit of writing consisting of one or more sentences that usually develops one main idea. Paragraphs are basic constituents of semantic structure and thus paragraph boundary estimation is an important building block of logical layout analysis. Moreover, paragraphs are often appropriate as processing units for various downstream tasks such as translation and information extraction




Figure 1. Examples of paragraphs in printed text. Paragraphs may have complex shapes when wrapped around figures or other types of document entities.

because they are self-contained and have rich semantic information. Therefore, developing a generic paragraph estimation algorithm is of great interest by itself.

Paragraphs are usually rendered in a geometric layout structure according to broadly accepted typographical rules. In this work, we exclude semantic paragraphs that can span over multiple text columns or pages, and only consider physical/geometrical paragraphs. There are usually clear visual cues to identify such paragraphs, but the task of estimating paragraphs is non-trivial as shown in Fig. 1.

Previous studies have attempted to develop a paragraph estimation method by defining handcrafted rules based on careful observations [23, 29, 4, 28] or by learning an object detection model to identify the regions of paragraphs from an image [35, 38]. For the former approaches, it is usually challenging to define a robust set of heuristics even for a limited domain, and hence machine-learning-based solutions are generally preferable. The latter approaches tend to have difficulty dealing with diverse aspect ratios and text shapes, and the wide range of degradations observed in real-world applications such as image skews and perspective distortions.

In this paper, we propose to apply graph convolutional networks (GCNs) in a post-processing step of an optical

character recognition (OCR) system for paragraph recognition. Recent advancements in graph neural (convolutional) networks [26, 34] have enabled deep learning on non-Euclidian data. GCNs can learn spatial relationships among entities combining information from multiple sources and provide a natural way to learn the non-linear mapping from OCR results to paragraphs.

More specifically, we design two classifiers based on GCNs — one for line splitting and one for line clustering. A word graph is constructed for the first stage and a line graph for the second stage. Both are constructed based on the β -skeleton algorithm [15] that produces a graph with good connectivity and sparsity.

To fully utilize the models' capability, it is desirable to have a diverse set of document styles in the training data. We create synthetic data sets from web pages where the page styles are randomly modified in the web scraping engine. By leveraging open web sites like Wikipedia [1] for source material to render in randomized styles, we have access to practically unlimited document data.

We evaluate the 2-step models on both the PubLayNet [38] and our own datasets. We show that GCN based models can be small and efficient by taking OCR produced bounding boxes as input, and are also capable of generating highly accurate results. Moreover, with synthesized training data from a browser-based rendering engine, these models can be a step towards a reverse rendering engine that recovers comprehensive layout structure from document images.

2. Related Work

2.1. Page Segmentation

A lot of previous work have studied the page segmentation task, including CRF based approaches [22, 30, 18], CNN based approaches [35, 17] and mixed algorithms [21].

While the pixel masks from a segmentation can tell us where the paragraphs are, they do not produce individual paragraphs. For example, when text is dense and paragraphs are only hinted by subtle indentations, the adjacency graph in [21] produces many false positive edges that form multi-paragraph text components.

As a result, the problem we are trying to solve is different. Our work takes OCR result (text lines and words) rather than the image as input, and the goal is to recognize the paragraphs among the lines so as to improve the overall structure of the OCR engine output.

2.2. Geometric and Rule-based Approaches

Early studies have proposed geometric methods [4, 3] and rule-based methods [23, 29, 28]. Both categories have algorithms to find column gaps by searching for white space [3] or text alignment [28].

Limitations of these approaches include susceptibility




Figure 2. Example of multiple short paragraphs densely packed and rotated into a non axis-aligned direction. The right side shows the region proposal boxes for object detection models.

to input noise and false positive column boundaries from monospace font families. Especially when handling scene text with perspective distortions from camera angles, rule based algorithms can be fragile and inconsistent.

2.3. Image Based Detection

The PubLayNet paper [38] provides a large dataset for multiple types of document entities, as well as two object detection models F-RCNN [24] and M-RCNN [11] trained to detect these entities. Both show good metrics in evaluations, but with some inherent limitations.

- Cost: Object detection models are typically large in size and expensive in computation. When used together with an OCR engine to retrieve text paragraphs, it seems wasteful to bypass the OCR results and attempt to detect paragraphs independently.
- Quality: Paragraph bounding boxes may have high aspect ratios and are sometimes tightly packed. In Fig. 2, several short paragraphs are printed with dense text and rotated by 45 degrees. The region proposals required to detect all the paragraphs are highly overlapped, so some detections will be dropped by non-maximum suppression (NMS). Rotational RCNN models [14] can mitigate this issue by inclined NMS, but further increase the computational cost while still facing a more difficult task with rotated or warped inputs.

2.4. Graph Neural Networks

Graph neural/convolutional networks have been used to extract document entities like tables [25] and curved lines [37, 20]. These work show that graph neural networks are flexible for handling various types of entities with complex shapes. One possible limitation from these approaches is on graph construction – the axis-aligned visibility graph in [25] can usually handle scanned documents but not scene text with image rotations and distortions, and the KNN graph in [37, 20] can form isolated components that restrict graph operations.

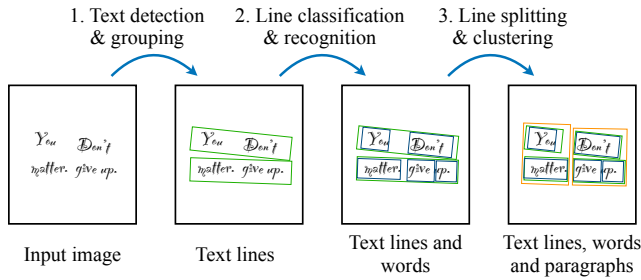


Figure 3. Overview workflow of a typical OCR engine (stage 1 and 2) and our post-OCR paragraph recognition (stage 3). Lines are shown in green boxes, words in blue, and paragraphs in orange.

3. Proposed Method

A typical general purpose OCR engine produces a set of text lines with recognized transcriptions [33]. To find paragraphs, we can consider a bottom-up approach to cluster text lines into paragraphs.

As shown in Fig. 3, the detected lines from stage 1 provide rudimentary layout information, but may not match the true text lines. The image in this example contains 2 text columns, each column containing a sentence which also forms a paragraph. The text line detector (stage 1) tries to find the longest curved fitted baselines, thus not able to split the lines by the 2-column layout. It is after stage 2 when the word boxes are available that we can perform a post-OCR layout analysis. We propose a 2-step process, namely line splitting and line clustering, to cluster the words and lines into paragraphs.

Both the line splitting and line clustering are non-trivial tasks for general-purpose paragraph estimation – the input images can be skewed or warped, and the layout styles can vary among different types of documents, e.g. newspapers, books, signs, web pages, handwritten letters, etc. Even though the concept of paragraph is mostly consistent across all document categories, the appearance of a paragraph can differ by many factors such as word spacing, line spacing, indentation, text flowing around figures, etc. Such variations make it difficult, if not impossible, to have a straightforward algorithm that identifies all the paragraphs.

We design the two steps based on graph convolutional neural networks (GCN) [34, 8] that takes input features from the coordinate values of OCR output boxes, together with a β -skeleton graph [15] constructed from these boxes. Neither the original image nor text transcriptions are included in the input, so the models are small, fast, and entirely focused on the layout structure.

- Step 1: Line splitting. Raw text lines from OCR line detectors may cross multiple columns, and thus need to be split into shorter lines. A GCN node classifier takes word boxes to predict splitting points in lines.
- Step 2: Line clustering. The refined lines produced by

step 1 are clustered into paragraphs. A GCN edge classifier takes line boxes to predict clustering operations on pairs of neighboring lines.

Output of each model is applied to the OCR text lines, with some additional error-correction heuristics (e.g. lines too far apart should not be clustered).

3.1. β -skeleton on Boxes

A graph is a key part of the GCN model input. We want a graph with high connectivity for effective message passing in graph convolutions, while also being sparse for computational efficiency.

Visibility graphs have been used in previous studies [7, 25], where edges are made by “lines-of-sight”. However, they are unsuitable for our models because of the edge density. Fig. 4(a) shows the visibility graph built on two rows of boxes, where any pairs of boxes on different rows are connected. This means word connections between text lines may get $O(n^2)$ number of edges. If we limit the lines-of-sight to be axis aligned like Fig. 4(b), then the graph becomes too sparse, even producing disconnected components in some cases. In comparison, k-nearest-neighbor graphs used in [37, 20] are more scalable, but can also produce dense and isolated components.




Figure 4. Comparison among different types of graphs constructed on an example set of boxes.




Figure 5. Building a box β -skeleton from point based β -skeleton. Left side: intersecting boxes are first connected with edges of length 0. Right side: Non-internal peripheral points (in green) are connected with β -skeleton edges which are then collapsed into box edges. Edge lengths are approximate. The middle line points are added so that no edges can go through the boxes.

By changing “lines-of-sight” into “balls-of-sight” in visibility graphs, we get a β -skeleton graph [15] with $\beta = 1$. In such a graph, two boxes are connected if they can both touch a circle that does not intersect with any other boxes. It provides a good balance between connectivity and sparsity. As shown in Fig. 4(c), a β -skeleton graph does not have excessive connections between rows of boxes. With $\beta = 1$, it is a subgraph of a Delaunay triangulation [2] with number of edges bounded by $O(n)$. Yet, it provides good connectivity within any local cluster of boxes, and the whole graph is guaranteed to be one connected component.

The original β -skeleton graph is defined on a point set. To apply it to rectangular boxes, we build a graph on peripheral points of all the box as in Fig. 5, and keep at most one edge between each pair of boxes.

3.2. Message Passing on Graphs

Our graph convolutional network resembles the message passing neural network (MPNN) [9] and GraphSage [10]. We use the term “message passing phase” from [9] to describe the graph level operations in our models. In this phase, repeated steps of “message passing” are performed based on a message function M and node update function U . At step t , a message $M(h_v^t, h_w^t)$ is passed along every edge e_{vw} in the graph where h_v^t and h_w^t are the hidden states of node v and w . Let $N(v)$ denote the neighbors of node v in the graph, the aggregated message by average pooling received by v is

$$m_v^{t+1} = \frac{\sum_{w \in N(v)} M(h_v^t, h_w^t)}{|N(v)|} \quad (1)$$

and the updated hidden state is

$$h_v^{t+1} = U(h_v^t, m_v^{t+1}) \quad (2)$$

Alternatively, we can use attention weighted pooling [32] to enhance message aggregation. Consequently, the model is also called a graph attention network (GAT), where calculation of m_v^{t+1} is replaced by

$$m_v^{t+1} = \frac{\sum_{w \in N(v)} \exp(e_{vw}^t) M(h_v^t, h_w^t)}{\sum_{w \in N(v)} \exp(e_{vw}^t)} \quad (3)$$

and e_{vw}^t is computed from a shared attention mechanism a , for which we use the dot product self-attention in [31]. So

$$e_{vw}^t = a(h_v^t, h_w^t) = K(h_w^t) \cdot Q(h_v^t) \quad (4)$$

where K is a shared key function and Q is a shared query function.

3.3. Splitting Lines

When multi-column text blocks are present in a document page, splitting lines across columns is a necessary first step [4, 28]. Note that the horizontal spacings between words is not a reliable signal for this task, as when the typography alignment of the text is “justified,” i.e. the text falls flush with both sides, these word spacings may be stretched to fill the full column width. In Fig. 8, the 2nd-to-last left line has word spacings larger than the column gap. This is common in documents with tightly packed text such as newspapers.

We use the GCN model shown in Fig. 6 to predict the splitting points, or tab-stops as in [28]. Each graph node is a word bounding box. Graph edges are the β -skeleton edges built as described in section 3.1. The model output contains two sets of node classification results – whether each word is a “line start” and whether it is a “line end”.

Fig. 8 shows a β -skeleton graph constructed from the word bounding boxes. Since words are aligned on either side of the two text columns, a set of words with their left edges all aligned are likely on the left boundary of a column, i.e. these words are line starts. Similarly, a set of words with right edges aligned are likely on the right boundary, i.e. they are line ends. The β -skeleton edges can connect aligned words in neighboring lines, and pass box alignment signals for the effective learning of the GCN model.




Figure 6. Overview of the line splitting model. In the output, line start nodes are marked green and line end nodes are marked orange.




Figure 7. Overview of the line clustering model. In the output, positive edges are marked pink.




Figure 8. Line splitting signal from word box alignment propagating through β -skeleton edges. The resulting predictions are equivalent to tab-stop detection.




Figure 9. Example of paragraph line clustering by indentations. Light blue edges indicate the β -skeleton constructed on line bounding boxes, and pink edges indicate that the connected lines are clustered into paragraphs.

3.4. Clustering Lines

After splitting all the lines into “true” lines, the remaining task is to cluster them into paragraphs. Again we use a GCN, but each graph node is a line bounding box, and the output is edge classification similar to link predictions in [19, 36]. We define a positive edge to connect two consecutive lines in the same paragraph. Note that it is possible to have non-consecutive lines in the same paragraph being connected by a β -skeleton edge. Such edges are defined as negative to make the task easier to learn.

Fig. 7 is an overview of the line clustering model. The input consists of line bounding boxes, and an additional “node-to-edge” step is added for the final edge output:

$$m'_{e=(v,w)} = \frac{M'(h_v, h_w) + M'(h_w, h_v)}{2} \quad (5)$$

The model predicts whether two lines belong to the same paragraph on each pair of lines connected with a β -skeleton edge. The predictions are made from multiple types of context like indentations (Fig. 9) and line spacings.

4. Experiments

We experiment with the 2-step GCN models and evaluate the end-to-end performance on both the open PubLayNet dataset, our synthetic web-scraped set, and a human annotated image set. The OCR engine is from Google Cloud Vision API DOCUMENT_TEXT_DETECTION v2021, and GCN setup details are in the appendix.

The GCN models are compared against other approaches. Besides the F-RCNN and M-RCNN from [38], we train an F-RCNN model with additional quadrilateral outputs for rotated boxes, denoted by “F-RCNN-Q” in following subsections. It uses a ResNet-101 [12] backbone at ~ 200 MB in size. In contrast, the GCN models are only ~ 100 KB each. A rule-based heuristic algorithm in our production system is also used as baseline.

4.1. Datasets

4.1.1 PubLayNet

PubLayNet [38] contains a large amount of document images with ground truth annotations: 340K in the training set and 12K in the development/validation (dev) set. The testing set ground truth has not been released at the time of this writing, so we use the dev set for evaluation.

4.1.2 Web Synthetic Page Layout

Data diversity is a crucial necessity for handling all types of inputs. By taking advantage of high quality and publicly available web documents, as well as a powerful rendering engine used in modern browsers, we can generate synthetic training data with a web scraper.

We use a browser-based web scraper to retrieve a list of Wikipedia [1] pages, where each result includes the image rendered in the browser as well as the HTML DOM (document object model) tree. The DOM tree contains the

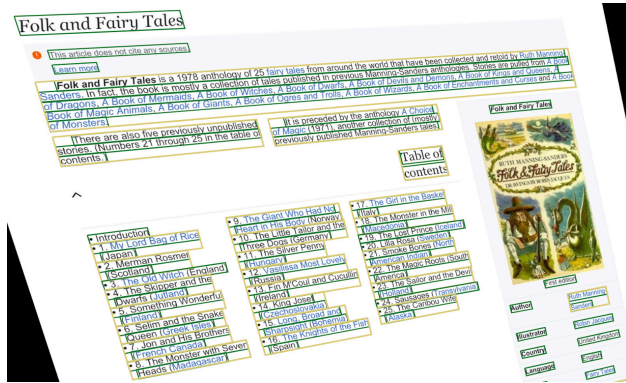


Figure 10. Training data example from web scraping with randomized style changes and data augmentation. Green boxes indicate line ground truth labels and yellow boxes indicate multi-line paragraph ground truth labels.

complete document structure and detailed locations of all the rendered elements, from which we can reconstruct the ground truth line bounding boxes. Each line bounding box is an axis-aligned rectangle covering a line of text. For paragraph ground truth, the HTML tag `<p>` conveniently indicates a paragraph node, and all the text lines under this node belong to this paragraph.

By running extension scripts in the browser, we can randomly change layout styles of web pages and diversify our data. For example, to generate double-column text for a certain division of a page, we can use “`div.style.columnCount = 2`.” And to emulate the effect of camera angles on scene text, we further augment the data by perform randomized rotations and perspective projections on each scraped page. These methods produce a great variety of layout styles that can be encountered in the real world. Fig. 10 shows an example page from the augmented web synthetic data.

4.1.3 Human Annotated Paragraph Dataset

We have a human annotated set with real-world images – 25K in English for training and a few hundred for testing in each available language. The images are collected from books, documents or objects with printed text, and sent to a team of raters who draw ground truth polygons for paragraphs. Example images are shown in Fig. 13, 14 and 15.

4.2. Evaluation Metrics

We measure the end-to-end performance of our OCR-GCN models by IoU based metrics such as the COCO $mAP@IoU[.50:.95]$ used in [38] so the results are comparable. The average precision (AP) for mAP is usually calculated on a precision-recall curve. But since our models produce binary predictions, we have only one output set of paragraph bounding boxes, i.e. only one point on the precision-recall curve. So $AP = precision \times recall$.

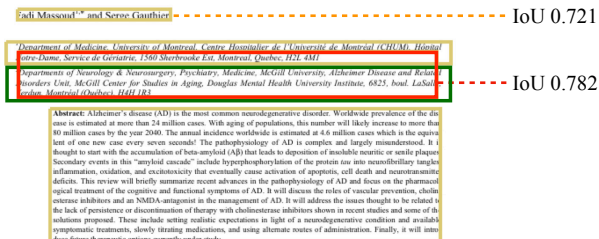


Figure 11. A paragraph detection example on an image from PubLayNet [38]. Yellow boxes are correct detections in terms of enclosed words, the red box is a wrong detection and the green box is ground truth. A single-line correct detection has lower IoU than a multi-line wrong one, which calls for $F1_{var}$ score with variable IoU thresholds.

For a better evaluation on paragraphs, we introduce an F1-score of variable IoU thresholds ($F1_{var}$ for short). As shown in Fig. 11, a single-line paragraph has a lower IoU even though it is correctly detected, while a 4-line detection (in red) has a higher IoU with a missed line. This is caused by boundary errors at character scale rather than at paragraph scale. This error is larger for post-OCR methods since the OCR engine is not trained to fit the paragraph training data. If we have line-level ground truth in each paragraph, and adjust IoU thresholds T_{iou} by

$$T_{iou} = \min\left(1 - \frac{1}{1 + \#lines}, 0.95\right) \quad (6)$$

the single-line paragraph will have IoU threshold 0.5, the 5-line one will have IoU threshold 0.833, and both cases in Fig. 11 can be more reasonably scored.

Both PubLayNet [38] and our web synthetic dataset have line level ground truth to support this $F1_{var}$ metric. For the human annotated set without line annotations, we fall back to a fixed IoU threshold of 0.5.

4.3. PubLayNet Evaluations

The PubLayNet dataset has five types of layout elements: text, title, list, figure and table. For our task, we take text and title bounding boxes as paragraph ground truth, and set all other types as “don’t-care” for both training and testing.

Table 1 shows that F-RCNN-Q matches the mAP scores in [38]. The GCN models are worse in this metric because there is only one point in the precision-recall curve, and the OCR engine is not trained to produce bounding boxes that match the ground truth. In the bottom row of Table 1, “OCR + Ground Truth” is computed by clustering OCR words into paragraphs based on ground truth boxes, which is the upper bound for all post-OCR methods. For mAP scores, even the upper bound is lower than the scores of image based models. However, if we measure by $F1_{var}$ scores defined above, OCR + GCNs can match image based models with a slight advantage. Fig. 12 shows some GCN produced examples.

Table 1. Paragraph mAP@IoU[.50:.95] score and $F1_{var}$ score comparisons. All models are tested on the PubLayNet development set. Numbers for mAP in the first 2 rows are from [38].

Model	Training Set	mAP	$F1_{var}$
F-RCNN [38]	PubLayNet training	0.910	-
M-RCNN [38]	PubLayNet training	0.916	-
F-RCNN-Q	PubLayNet training	0.914	0.945
Tesseract [27]	-	0.571	0.707
OCR + Heuristic	-	0.302	0.364
OCR + GCNs	Augmented web synthetic	0.748	0.867
OCR + GCNs	PubLayNet training	0.842	0.959
OCR + Ground Truth	-	0.892	0.997

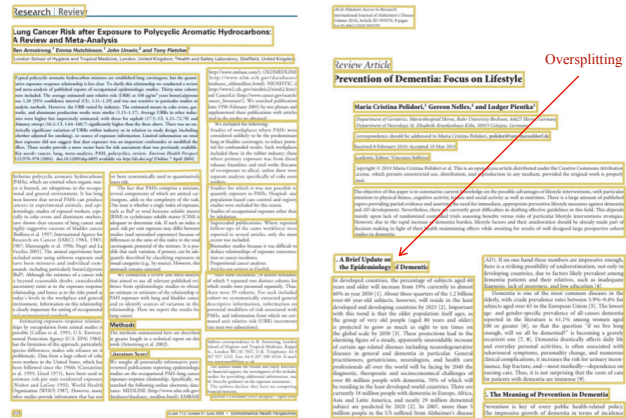


Figure 12. Representative PubLayNet examples of paragraphs from OCR and GCN line splitting + clustering. The left one gets all the correct predictions, whereas the right one has a line splitting error at section 1’s title.

The high $F1_{var}$ score on “OCR + Ground Truth” also shows that the OCR engine we use has a very high recall on text detection. The only reason it is lower than one is from ground truth variations – a small fraction of single-line paragraphs have IoU lower than 0.5.

4.4. Web Synthetic Evaluations

The synthetic dataset from web scraping gives a more difficult test for these models by its aggressive style variations. Data augmentation further increases the difficulty especially for image based detection models.

In Table 2, we can see the $F1_{var}$ score of the image based F-RCNN-Q model decreases sharply as the task difficulty increases. At “Augmented web synthetic” with images like Fig. 10, detection is essentially broken, not only from non-max suppression drops shown in Fig. 2, but also from worse box predictions.

In contrast, the GCN models are much less affected by layout style variations and data augmentations. The $F1_{var}$ score change is minimal between augmented and non-augmented datasets. So GCN models will have a greater advantage for scene text when input images are rotated.

Table 2. Paragraph $F1_{var}$ score comparison across different types of models and datasets. Data difficulty increases monotonically from PubLayNet to Augmented web synthetic.

Model	Data Source for Training & Test	$F1_{var}$
F-RCNN-Q	PubLayNet	0.945
	Web synthetic	0.722
	Augmented web synthetic	0.547
OCR + GCNs	PubLayNet	0.959
	Web synthetic	0.830
	Augmented web synthetic	0.827

Table 3. Paragraph F1-scores tested on the real-world test set with paragraph annotations. Fixed IoU threshold 0.5 is used since there is no line-level ground truth to support variable thresholds.

Model	Training Data	F1@IoU0.5
F-RCNN-Q	Augmented web synthetic	0.030
F-RCNN-Q	Annotated data (pre-trained on PubLayNet)	0.607
OCR + Heuristic	-	0.602
OCR + GCNs	Augmented web synthetic	0.614
OCR + GCNs	Annotated data	0.671
OCR + GCNs	Augmented synthetic + Annotated	0.671
OCR + Ground Truth	-	0.960

4.5. Real-word Dataset Evaluations

The human annotated dataset can potentially show the models’ performance in real-world applications. The annotated set is relatively small, so the F-RCNN-Q model needs to be pre-trained on PubLayNet, while the GCN models are small enough to be trained entirely on this set. Evaluation metric for this set is F1@IoU0.5.

Table 3 shows comparisons across different models and different training sets. Note that Faster R-CNN trained from synthetic web data does not work at all for real-world images, whereas the OCR+GCN models can generalize well.

Fig. 13 and Fig. 14 show some examples of OCR + GCNs produced paragraphs. The right image in Fig. 13 shows the effectiveness of the augmented web synthetic data, as there are no similar images in the annotated set. On the other hand, the right table in Fig. 14 is not recognized since our models only takes bounding box coordinates as input. Using GCNs for table detection like [25] is another interesting topic but out of the scope of this paper.

To verify the robustness of the GCN models for language and script diversity, we test them on a multi-language evaluation set. The GCN models are trained with additional synthetic data from Wikipedia pages in Chinese, Japanese and Korean. Table 4 once again shows the generalizability of GCN models. F-RCNN-Q is not trained in the three Asian languages for the lack of training data.

The GCN models are also flexible in handling text lines written in vertical directions, which are common in

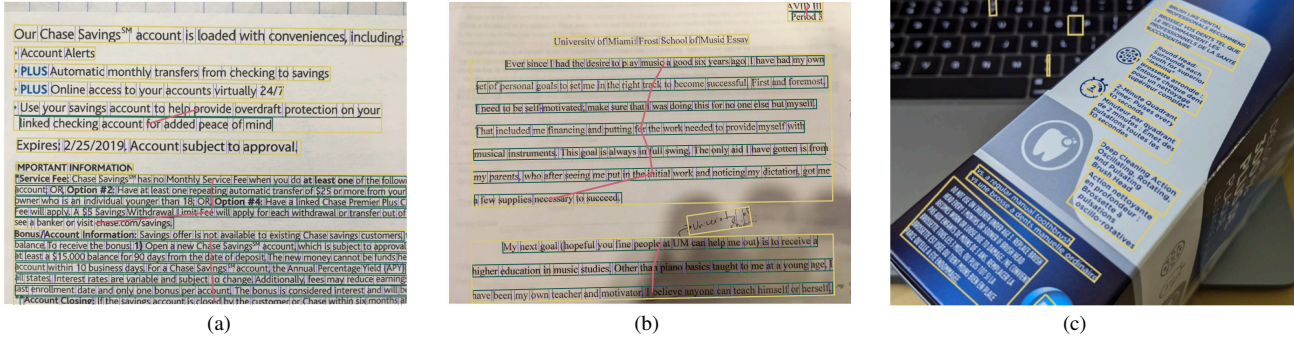


Figure 13. Representative examples of real-world images with OCR followed by GCN line splitting and line clustering. Paragraphs shown in yellow boxes, lines in green and words in blue. Pink line segments indicate positive line clustering predictions.

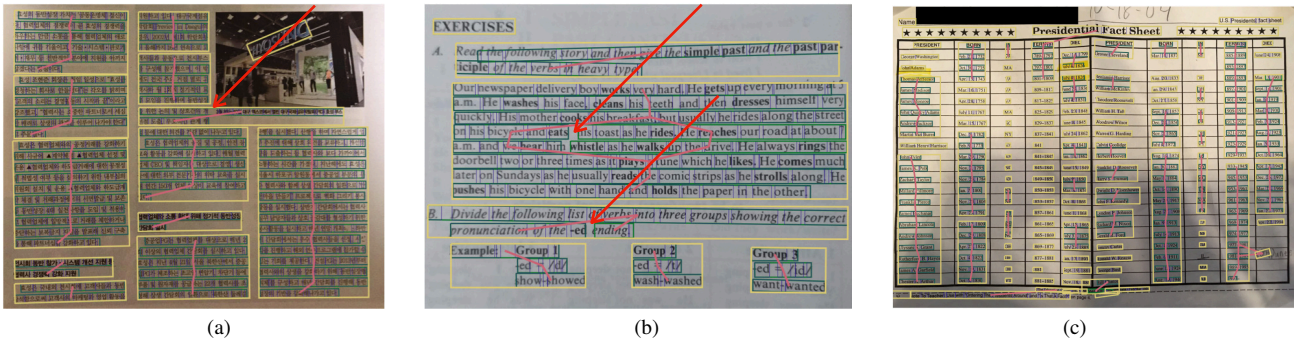


Figure 14. Paragraph errors in real-world images. (a) Line under splitting. (b) Line over splitting. (c) Over clustering table elements.

Japanese and Chinese, and also appear in Korean. Although we don't have much training data with vertical lines, the bounding box structures of lines and symbols in these languages remain the same when the lines are written vertically, as if they were written horizontally while the image is rotated clockwise by 90 degrees. Fig. 15 shows such an example. Since our models are trained to handle all rotation angles, such paragraphs can be correctly recognized.

5. Conclusions and Future Work

We demonstrate that GCN models can be powerful and efficient for the task of paragraph recognition. Provided with a good OCR engine, they can match image based

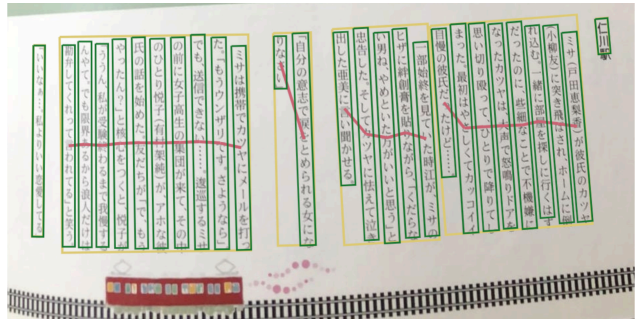


Figure 15. Example of paragraphs from text lines with vertical writing direction.

Table 4. F1@IoU0.5 scores tested on the multi-language evaluation set.

Language	OCR + Heuristic	F-RCNN -Q	OCR + GCNs	OCR + Ground Truth
English	0.429	0.513	0.544	0.890
French	0.438	0.557	0.553	0.885
German	0.427	0.538	0.566	0.873
Italian	0.455	0.545	0.556	0.862
Spanish	0.449	0.597	0.616	0.885
Chinese	0.370	-	0.485	0.790
Japanese	0.398	-	0.487	0.772
Korean	0.400	-	0.547	0.807

models with much lower requirement on training data and computation resources, and significantly beat them on non-axis-aligned inputs with complex layout styles. The graph convolutions in these models give them unique advantages in dealing with different levels of page elements and their structural relations.

Future work include extending the GCN models to find more types of entities and extract document structural information. Joining image based CNN backbones with GCN may work better for entities with non-text components like checkboxes and grid lines. In addition, reading order among entities will be helpful if we want to identify semantic paragraphs that span across multiple columns/pages.

References

- [1] Wikipedia, the free encyclopedia.
- [2] Mark de Berg, Otfried Cheong, Marc van Kreveld, and Mark Overmars. *Computational Geometry: Algorithms and Applications*. Springer-Verlag TELOS, Santa Clara, CA, USA, 3rd ed. edition, 2008.
- [3] T. M. Breuel. An algorithm for finding maximal white-space rectangles at arbitrary orientations for document layout analysis. In *Seventh International Conference on Document Analysis and Recognition, 2003. Proceedings.*, pages 66–70 vol.1, 2003.
- [4] Thomas M. Breuel. High performance document layout analysis. In *Proceedings of the Symposium on Document Image Understanding Technology*, pages 209–218, Greenbelt, MD, 2003.
- [5] R. Cattoni, T. Coianiz, S. Messelodi, and C. M. Modena. Geometric layout analysis techniques for document image understanding: A review. Technical Report 9703-09, IRST, Trento, Italy, 1998.
- [6] Christian Clausner, Apostolos Antonacopoulos, and Stefan Pletschacher. ICDAR2019 competition on recognition of documents with complex layouts - RDCL2019. In *2019 International Conference on Document Analysis and Recognition, ICDAR 2019, Sydney, Australia, September 20-25, 2019*, pages 1521–1526. IEEE, 2019.
- [7] Brian L. Davis, Bryan S. Morse, Scott Cohen, Brian L. Price, and Chris Tensmeyer. Deep visual template-free form parsing. In *2019 International Conference on Document Analysis and Recognition, ICDAR 2019, Sydney, Australia, September 20-25, 2019*, pages 134–141. IEEE, 2019.
- [8] David K Duvenaud, Dougal Maclaurin, Jorge Iparraguirre, Rafael Bombarell, Timothy Hirzel, Alan Aspuru-Guzik, and Ryan P Adams. Convolutional networks on graphs for learning molecular fingerprints. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28, pages 2224–2232. Curran Associates, Inc., 2015.
- [9] Justin Gilmer, Samuel S. Schoenholz, Patrick F. Riley, Oriol Vinyals, and George E. Dahl. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17*, page 1263–1272. JMLR.org, 2017.
- [10] William L. Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 1024–1034, 2017.
- [11] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask R-CNN. *IEEE Trans. Pattern Anal. Mach. Intell.*, 42(2):386–397, 2020.
- [12] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016.
- [13] W. Horak. Office document architecture and office document interchange formats: Current status of international standardization. *Computer*, 18:50–60, October 1985.
- [14] Y. Jiang, X. Zhu, X. Wang, S. Yang, W. Li, H. Wang, P. Fu, and Z. Luo. R2 cnn: Rotational region cnn for arbitrarily-oriented scene text detection. In *2018 24th International Conference on Pattern Recognition (ICPR)*, pages 3610–3615, 2018.
- [15] David G. Kirkpatrick and John D. Radke. A framework for computational morphology. *Machine Intelligence and Pattern Recognition*, 2:217–248, 1985.
- [16] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [17] Joonho Lee, Hideaki Hayashi, Wataru Ohyama, and Seiichi Uchida. Page segmentation using a convolutional neural network with trainable co-occurrence features. In *2019 International Conference on Document Analysis and Recognition, ICDAR 2019, Sydney, Australia, September 20-25, 2019*, pages 1023–1028. IEEE, 2019.
- [18] Yixin Li, Yajun Zou, and Jinwen Ma. Deeplayout: A semantic segmentation approach to page layout analysis. In De-Shuang Huang, M. Michael Gromiha, Kyungsook Han, and Abir Hussain, editors, *Intelligent Computing Methodologies*, pages 266–277, Cham, 2018. Springer International Publishing.
- [19] David Liben-Nowell and Jon Kleinberg. The link-prediction problem for social networks. *Journal of the American Society for Information Science and Technology*, 58(7):1019–1031, May 2007.
- [20] Chixiang Ma, Lei Sun, Zhuoyao Zhong, and Qiang Huo. Re-LaTeX: Exploiting visual relationships for arbitrary-shaped scene text detection with graph convolutional networks. *Pattern Recognition*, 111:107684, mar 2021.
- [21] Ana Lucia Lima Marreiros Maia, Frank Dennis Julca-Aguilar, and Nina Sumiko Tomita Hirata. A machine learning approach for graph-based page segmentation. In *31st SIBGRAPI Conference on Graphics, Patterns and Images, SIBGRAPI 2018, Paraná, Brazil, October 29 - Nov. 1, 2018*, pages 424–431. IEEE Computer Society, 2018.
- [22] S. Nicolas, J. Dardenne, T. Paquet, and L. Heutte. Document image segmentation using a 2d conditional random field model. In *Ninth International Conference on Document Analysis and Recognition (ICDAR 2007)*, volume 1, pages 407–411, 2007.
- [23] Debashish Niyogi and Sargur N. Srihari. A rule-based system for document understanding. In *Proceedings of the Fifth AAAI National Conference on Artificial Intelligence, AAAI'86*, page 789–793. AAAI Press, 1986.
- [24] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(6):1137–1149, 2017.
- [25] Pau Riba, Anjan Dutta, Lutz Goldmann, Alicia Fornés, Oriol Ramos Terrades, and Josep Lladós. Table detection in invoice documents by graph neural networks. In *2019 International Conference on Document Analysis and Recognition*

- dition, *ICDAR 2019, Sydney, Australia, September 20-25, 2019*, pages 122–127. IEEE, 2019.
- [26] F. Scarselli, M. Gori, A. C. Tsoi, M. Hagenbuchner, and G. Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2009.
- [27] Ray Smith. An overview of the Tesseract OCR engine. In *Proc. 9th IEEE Intl. Conf. on Document Analysis and Recognition (ICDAR)*, pages 629–633, 2007.
- [28] Raymond W. Smith. Hybrid page layout analysis via tab-stop detection. In *10th International Conference on Document Analysis and Recognition, ICDAR 2009, Barcelona, Spain, 26-29 July 2009*, pages 241–245. IEEE Computer Society, 2009.
- [29] Sargur N. Srihari and G. W. Zack. Document image analysis. In *Proceedings of the 8th International Conference on Pattern Recognition*, page 434–436. AAAI Press, 1986.
- [30] Xin Tao, Zhi Tang, Canhui Xu, and Yongtao Wang. Logical labeling of fixed layout pdf documents using multiple contexts. In *2014 11th IAPR International Workshop on Document Analysis Systems*, pages 360–364, 2014.
- [31] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *CoRR*, abs/1706.03762, 2017.
- [32] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018*, 2018.
- [33] Jake Walker, Yasuhisa Fujii, and Ashok Popat. A web-based ocr service for documents. In *2018 13th IAPR International Workshop on Document Analysis Systems*, pages 21–22, 2018.
- [34] Z. Wu, S. Pan, F. Chen, G. Long, C. Zhang, and P. S. Yu. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, pages 1–21, 2020.
- [35] Xiao Yang, Ersin Yumer, Paul Asente, Mike Kraley, Daniel Kifer, and C Lee Giles. Learning to extract semantic structure from documents using multimodal fully convolutional neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [36] Muhan Zhang and Yixin Chen. Link prediction based on graph neural networks. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada*, pages 5171–5181, 2018.
- [37] Shi-Xue Zhang, Xiaobin Zhu, Jie-Bo Hou, Chang Liu, Chun Yang, Hongfa Wang, and Yin xu. Deep relational reasoning graph network for arbitrary shape text detection. pages 9696–9705, 06 2020.
- [38] Xu Zhong, Jianbin Tang, and Antonio Jimeno-Yepes. Publaynet: Largest dataset ever for document layout analysis. In *2019 International Conference on Document Analysis and Recognition, ICDAR 2019, Sydney, Australia, September 20-25, 2019*, pages 1015–1022. IEEE, 2019.