

# TreeForm: End-to-end Annotation and Evaluation for Form Document Parsing

Ran Zmigrod, Zhiqiang Ma, Armineh Nourbakhsh, Sameena Shah  
J.P. Morgan AI Research  
{first\_name}.{last\_name}@jpmchase.com

## Abstract

Visually Rich Form Understanding (VRFU) poses a complex research problem due to the documents' highly structured nature and yet highly variable style and content. Current annotation schemes decompose form understanding and omit key hierarchical structure, making development and evaluation of end-to-end models difficult. In this paper, we propose a novel F1 metric to evaluate form parsers and describe a new content-agnostic, tree-based annotation scheme for VRFU: **TreeForm**. We provide methods to convert previous annotation schemes into TreeForm structures and evaluate TreeForm predictions using a modified version of the normalized tree-edit distance. We present initial baselines for our end-to-end performance metric and the TreeForm edit distance, averaged over the FUNSD and XFUND datasets, of 61.5 and 26.4 respectively. We hope that TreeForm encourages deeper research in annotating, modeling, and evaluating the complexities of form-like documents.

## 1 Introduction

Visually rich document understanding (VRDU) has been a growing field in multimodal AI research. VRDU takes document images as input and applies tasks such as document classification (Gu et al., 2021; Kim et al., 2022; Gu et al., 2022), information extraction (Borchmann et al., 2021; Wang et al., 2022), visual question-answering (VQA) (Mathew et al., 2021; Tito et al., 2021; Li et al., 2022), *inter alia*. Consequently, state-of-the-art VRDU models rely on image-to-text understanding, such as optical character recognition (OCR), as well as multimodal models that can exploit textual, visual, and spatial features of documents.

VRFU is a subset of VRDU that focuses on forms, i.e., documents that contain a collection of hierarchical key-value pairs (in various styles) regarding specific entities. Forms are highly structured documents, and so can be directly parsed into

WINSTON & STRAWN  
MIAMI TRIAL SITE, 3000 FIRST UNION FINANCIAL CENTER, 200 SOUTH BISCAYNE BOULEVARD, MIAMI, FLORIDA 33131  
FAX: (305) 400-6108

200 PARK AVENUE, NEW YORK, NY 10022-6103, 212-859-8770  
1400 L STREET, N.W., WASHINGTON, DC 20005-3502, 202-371-3700  
35 WEST WALKER, CHICAGO, IL 60601, 312-859-8500  
21 AVENUE VICTOR HUGO, 75110 PARIS, FRANCE, 33-1-43-54-82-82

Fax Number: 305 400-6107

FROM: Kevin Narko  
DATE: 10/13/99

CHARGEBACK: 4162/158

Please Deliver as Soon as Possible To:

RECIPIENT	COMPANY	FAX NO.	PHONE NO.
John Mulderig	Philip Morris	917-663-5796	917-663-3056
Gregory Little	Philip Morris	917-663-5979	

Figure 1: Excerpt of a FUNSD form. Headers are marked in **burgundy**, questions are marked **green**, and answers are marked in **blue**. Entity links provided by the FUNSD annotation schemes are marked in **orange**. Links in **pink** were created for TreeForm.

a structured object. Parsing receipts, which share a similar though smaller and simpler structure, have been the focus of much document parsing research (Majumder et al., 2020; Borchmann et al., 2021; Gao et al., 2021). In this work, we focus on the broader field of form parsing.

The FUNSD dataset (Jaume et al., 2019), and its multilingual counter-part XFUND (Xu et al., 2022), are standard VRFU datasets used in the literature. These datasets split form parsing into entity labeling and entity linking tasks. Unfortunately, this annotation scheme is unable to fully express the structure of a form. Moreover, to the best of our knowledge, no holistic approach exists for evaluating complete FUNSD-type predictions. Few other form parsing datasets exist, such as FUNSD+ (Zagami and Helm, 2022) and the National Archives Forms (NAF) dataset (Davis et al., 2019). We do not address these directly but many of our contributions can be extended to these annotation schemes.

In this paper, we tackle the problem of complete form parsing through a tree-based approach that

successfully captures the hierarchical structure of forms which is often missed in modern annotation schemes. Our contributions are summarized below:

1. We provide novel end-to-end metrics for evaluating FUNSD-type annotations.
2. We present a novel tree-based representation of forms, **TreeForm**, that is conveniently stored as a single JSON object. We also utilize the greedy-aligned tree-edit distance (GANTED) proposed by [Davis et al. \(2022\)](#) to evaluate TreeForm predictions.
3. We detail a method to transform *any* FUNSD-type annotations into TreeForm. Our method captures additional hierarchical and tabular information, which FUNSD annotations do not contain. The additional information gained from TreeForm is visualized in Figure 1.<sup>1</sup>
4. We evaluate TreeForm and our novel metrics using state-of-the-art models ([Kim et al., 2022](#); [Xu et al., 2022](#)) on the FUNSD and XFUND datasets. Our baselines achieve a node-alignment accuracy of 0.22, end-to-end F1 score of 61.5, and TreeForm GANted score of 14.5 across all languages.

## 2 Limitations of Form Parsing Datasets

Most VRFU research employs FUNSD-type annotations ([Jaume et al., 2019](#)). This annotation scheme contains two components: First, a list of semantic entities defined by a group of tokens and a semantic label (header, question, or answer).<sup>2</sup> Second, a list of directed links between entities to denote form structure. An example of an FUNSD annotation is given in App. C. While the FUNSD (and XFUND) project was a fundamental step in VRFU, its approach contains underlying issues ([Vu and Nguyen, 2020](#)). Other non-FUNSD schematic VRFU datasets exist ([Davis et al., 2019](#); [Zagami and Helm, 2022](#)). However, they are not widely used and are not without similar issues. Table 1 provides an overview of form parsing datasets.

**Task Decomposition.** The FUNSD annotation scheme proposes that VRFU should be decomposed into semantic entity labeling and entity link-

<sup>1</sup>Note that for ease of visualization, Figure 1 does not contain boxes for every single annotation. In practice, every piece of information annotated in FUNSD (or the original annotation scheme) will be captured in TreeForm.

<sup>2</sup>The original annotations contain a fourth entity, *other*, for text that does not conform to one of the three aforementioned labels. We follow recent work and ignore these labels ([Xu et al., 2021b,a, 2022](#)).

Dataset	Lang.	Scheme	Train	Test
FUNSD	EN	FUNSD	149	50
XFUND	Multi	FUNSD	1043	350
FUNSD+	EN	FUNSD+	1023	116
NAF	EN	NAF	741	63

Table 1: Form Understanding Datasets. XFUND is equally split over seven languages (149 training documents, 50 test documents): DE, ES, FR, IT, JA, PT, ZH.

ing. This makes a proper end-to-end evaluation of form understanding models difficult as combining the two task metrics does not represent an accurate joint evaluation. Indeed, works that report individual F1 scores do not attempt to combine these ([Carbonell et al., 2020](#); [Li et al., 2021b](#); [Gemelli et al., 2022a](#); [Hong et al., 2022](#)) while much research omits the entity linking task ([Li et al., 2021a](#); [Appalaraju et al., 2021](#); [Xu et al., 2021b,a](#); [Chen et al., 2022](#); [Luo et al., 2022](#)).

**Inconsistent Annotations.** The FUNSD annotations contain several inconsistencies in both their entity labeling and entity linking components. Inconsistencies arise in hierarchical forms, handwritten input, and missing entity links. While many of these inconsistencies were improved in ([Vu and Nguyen, 2020](#)), some issues still persist

**Table Recognition.** Tables are a natural component of forms, however, they break the common question-answer structure. The FUNSD annotation scheme handles tables as columns, where a column header is a question and each value in the column is an answer linked to the question. Row headers are sometimes also annotated and linked, but this is a less common pattern in the dataset. Also, as pointed out by [Davis et al. \(2022\)](#), row and column alignment is not present in the annotations, making table parsing or reconstruction not possible in FUNSD.

## 3 An End-to-end Evaluation of FUNSD

In this section, we propose a new F1 metric to evaluate FUNSD-type predictions on their combined entity labeling and entity linking performance. Our metric, inspired by the labeled attachment score (LAS) ([Kübler et al., 2009](#)), is based on constructing a tree of the FUNSD annotations. Each node in the tree represents a semantic FUNSD entity, and each edge in the tree is a link between entities, labeled with the entity label of the child node. Entities with no incoming links have an incoming edge from a dummy root node, labeled with the

entity’s label. An example of this structure is given in Figure 2b. Let  $T$  be the set of predicted edges and  $G$  be the set of ground truth edges. We can then respectively define the precision and recall as:

$$\text{prec}(T, G) \stackrel{\text{def}}{=} \frac{|T \cap G|}{|T|} \quad \text{rec}(T, G) \stackrel{\text{def}}{=} \frac{|T \cap G|}{|G|}$$

These can be combined as normal to yield an F1 metric that examines both entity labels and links. Note that this metric equally punishes incorrect edges regardless of where they lie in the nested tree. For example, mis-predicting a top level node (i.e., one connected to the dummy root node) is punished equally to mis-predicting the edge connecting an answer to a question.

The above F1 metric does not consider node-alignment (i.e., word grouping) in its evaluation. Indeed, it assumes that the nodes of the predicted and ground truth trees are aligned. Node alignment can be done through a greedy algorithm. For each predicted entity, we compute the normalized Levenshtein distance (Mathew et al., 2021) to each true entity.<sup>3</sup> We then greedily select alignments that meet a certain threshold (edit distance less than 0.4), until no possible alignments remain. We define the node-alignment accuracy (NAA) metric as the mean normalized Levenshtein distance between each pair of aligned nodes.<sup>4</sup>

## 4 A Tree-based Annotation Scheme

In this section we present a novel tree-based form annotation scheme, **TreeForm**. Unlike our F1 and NAA metrics, TreeForm does not assume the FUNSD task decomposition. Instead, TreeForm provides a fully structured object that enables complete representation of a form.

A TreeForm node is represented by a JSON object (i.e., a dictionary) whose key and value respectively contain the node’s (textual) value and a dictionary where each key is either another node or an edge label that leads to a group of child nodes. Therefore, a TreeForm tree is represented using a single JSON object which efficiently contains all form information necessary for any down-stream VRFU task. Davis et al. (2022) also parses complete forms into JSON objects, however, their approach requires multiple objects per form and so

<sup>3</sup>We choose the Levenshtein distance to follow related work and to enable measuring the edit distance in a more meaningful way as we can assign context-specific penalties to additions, deletions, and alterations.

<sup>4</sup>For unaligned nodes, we use the normalized Levenshtein distance between a node and the empty string, i.e., 1.

annotations are less readable. TreeForm considers forms to be comprised of three key structures: headers, question-answer pairs, and tables; each of these structures is represented in a tree structure which is detailed in App. B. An example of a TreeForm annotation is given in Figure 2c.

### 4.1 Evaluating TreeForm

We assess TreeForm predictions using a variant of the normalized tree-edit distance (nTED) (Hwang et al., 2021b), a distance metric that informs us how much work is needed to convert the predicted tree into the ground-truth tree (and so a lower nTED is better). Specifically, we follow the greedy-aligned nTED (GANTED) metric proposed by Davis et al. (2022). This is a more holistic metric than those proposed in Section 3 as it does not assume any task decomposition.

### 4.2 Converting FUNSD into TreeForm

We propose a set of transformations and heuristics to convert FUNSD-type annotations into TreeForm structures. We first apply simple conversions for straight-forward cases such as single-answer questions and hierarchical headers which fit directly into the TreeForm annotation scheme. Next, we elaborate on the more complex heuristics we have chosen to create the most accurate and complete TreeForm structure. Note an important caveat to our transformations, TreeForm cannot attain information that was not annotated in FUNSD (or the original annotation scheme).

**Discarding Incomplete Information.** TreeForm does not include free-form text, unanswered questions, or unprompted answers. Any such annotations in the dataset are discarded here.

**Choosing the Form Title.** If headers exist in the form, we aim to assign a form title to each TreeForm structure. Since FUNSD does not necessarily nest headers correctly (as seen in Figure 1), we consider the form title to be the non-nested header that is spatially highest on the page. This seems like a sensible heuristic as we typically expect titles to be at the top of documents.<sup>5</sup>

**Constructing Tables.** Table structure is not directly recoverable from FUNSD annotations (as

<sup>5</sup>This heuristic is not perfect, as for example, the form title in Figure 1 could be “Fax Cover Sheet” rather than “WINSTON & STRAWN”; this also raises an inconsistency issue as it is unclear whether the latter header is indeed a header entity.



Figure 2: Different annotation schemes for excerpt of FUNSD form given in Figure 1.

discussed in Section 2). We utilize bounding box details to align columns and rows such that we can approximate the full table. Similar heuristics were also used in (Davis et al., 2022). Some questions with multiple answers in FUNSD may indicate a multi-line answer rather than a table column (or row). Therefore, we require multiple answers to start roughly in the same horizontal (or vertical) position to be considered a table entry.

## 5 Experiments

To evaluate TreeForm as well as our end-to-end FUNSD metrics, we fine-tune LayoutXML (Xu et al., 2022) and Donut (Kim et al., 2022) on each of the FUNSD and XFUND datasets.<sup>6</sup> We chose these two models as they represent the two cur-

<sup>6</sup>We fine-tune one variant of each model on FUNSD annotations and another on TreeForm. Details are given in App. D.

rent approaches to form parsing: Pipelined<sup>7</sup> (LayoutXML) and end-to-end (Donut). We describe pre- and post-processing steps required for applying the FUNSD and TreeForm evaluations to each model in App. D. Importantly, we performed post-processing steps to the output of LayoutXML to enable evaluation using our metrics. The results are given in Table 2.

For all languages, our end-to-end F1 metric is lower than one or both of the standard labeling and linking F1 scores. This indicates that while past work may have high labeling (or linking) F1 scores, they may not have full form understanding capabilities. This is further exacerbated as our F1 metric requires the NAA score to provide a holistic overview. As expected, LayoutXML outperforms Donut with respect to NAA as it has access to OCR tokens. However, Donut still outperforms Lay-

<sup>7</sup>The pipelined approach first predicts the entity labels and uses said predictions to perform the entity linking task.



Metric	Model	FUNSD			XFUND				
		EN	DE	ES	FR	IT	JA	PT	ZH
Labeling F1 ↑	LayoutXLM	77.6	79.5	74.1	78.9	77.3	79.1	79.4	88.9
Linking F1 ↑		48.8	57.3	63.8	61.4	53.5	64.4	53.0	67.8
NAA ↓	LayoutXLM	<b>0.21</b>	<b>0.24</b>	<b>0.28</b>	<b>0.23</b>	<b>0.23</b>	<b>0.24</b>	<b>0.23</b>	<b>0.12</b>
	Donut	0.69	0.58	0.69	0.61	0.64	0.75	0.75	0.74
Tree F1 ↑	LayoutXLM	62.6	57.0	<b>60.6</b>	<b>65.7</b>	59.6	<b>59.3</b>	<b>59.9</b>	<b>66.9</b>
	Donut	<b>64.8</b>	<b>61.5</b>	53.0	62.4	<b>62.9</b>	51.7	46.8	55.6
GAnTED ↓	LayoutXLM	18.0	27.7	20.7	22.3	20.6	7.2	25.1	<b>7.3</b>
	Donut	<b>15.1</b>	<b>21.0</b>	<b>15.8</b>	<b>13.6</b>	<b>17.5</b>	<b>6.7</b>	<b>17.3</b>	8.9

Table 2: Standard and end-to-end evaluation metrics for FUNSD and TreeForm annotations.

outXLM on F1 for some languages, suggesting it is better at understanding hierarchical structure.

We consider the median GAnTED score for evaluating TreeForm predictions. Donut outperformed LayoutXLM in all languages except for Chinese. This is expected as Donut was designed with document parsing in mind and so could be directly fine-tuned on TreeForm data whereas LayoutXLM was pipelined using the FUNSD annotation scheme (see App. D). As GAnTED has not yet been widely used in the literature, it is difficult to assess the effectiveness of these systems. Nevertheless, Dessurt (Davis et al., 2022) reported a GAnTED score of 23.4 on the FUNSD dataset which aligns with the scores in Table 2.<sup>8</sup> Additionally, the GAnTED scores seem to be similar across languages which attests to the generalizability of TreeForm.

## 6 Conclusion

In this paper, we described the challenges of current form understanding and the limitations of the current datasets and their annotation schemes. We introduced a novel F1 metric that can be applied to current annotation schemes and further proposed a new tree-based annotation scheme, TreeForm, that enables complete form parsing. We applied state-of-the-art models (LayoutXLM and Donut) to provide the first TreeForm baselines. We envision future work to create a new TreeForm dataset, that contains correct and consistent annotations.

## Disclaimer

This paper was prepared for informational purposes by the Artificial Intelligence Research group of JP Morgan Chase & Co and its affiliates (“JP Morgan”), and is not a product of the Research Depart-

<sup>8</sup>In the experimental set-up of Davis et al. (2022), a flatter tree representation (with some other key differences) is used and so the results are not directly comparable.

ment of JP Morgan. JP Morgan makes no representation and warranty whatsoever and disclaims all liability, for the completeness, accuracy or reliability of the information contained herein. This document is not intended as investment research or investment advice, or a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction, and shall not constitute a solicitation under any jurisdiction or to any person, if such solicitation under such jurisdiction or to such person would be unlawful.

## References

- Milan Aggarwal, Hires Gupta, Mausoom Sarkar, and Balaji Krishnamurthy. 2020a. [Form2Seq : A framework for higher-order form structure extraction](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3830–3840, Online. Association for Computational Linguistics.
- Milan Aggarwal, Mausoom Sarkar, Hires Gupta, and Balaji Krishnamurthy. 2020b. [Multi-modal association based grouping for form structure extraction](#). In *The IEEE Winter Conference on Applications of Computer Vision*, pages 2075–2084.
- Srikar Appalaraju, Bhavan Jasani, Bhargava Urala Kota, Yusheng Xie, and R. Manmatha. 2021. [Docformer: End-to-end transformer for document understanding](#). In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 973–983. IEEE.
- Lukasz Borchmann, Michal Pietruszka, Tomasz Stanislawek, Dawid Jurkiewicz, Michal Turski, Karolina Szyndler, and Filip Gralinski. 2021. [DUE: End-to-end document understanding benchmark](#). In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021, December 2021, virtual*.

- Manuel Carbonell, Pau Riba, Mauricio Villegas, Alicia Fornés, and Josep Lladós. 2020. [Named entity recognition and relation extraction with graph neural networks in semi structured documents](#). In *25th International Conference on Pattern Recognition, ICPR 2020, Virtual Event / Milan, Italy, January 10-15, 2021*, pages 9622–9627. IEEE.
- Jingye Chen, Tengchao Lv, Lei Cui, Cha Zhang, and Furu Wei. 2022. [XDoc: Unified pre-training for cross-format document understanding](#). *CoRR*, abs/2210.02849.
- Brian L. Davis, Bryan S. Morse, Scott Cohen, Brian L. Price, and Chris Tensmeyer. 2019. [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.
- Brian L. Davis, Bryan S. Morse, Brian L. Price, Chris Tensmeyer, and Curtis Wigington. 2021. [Visual FUDGE: form understanding via dynamic graph editing](#). In *16th International Conference on Document Analysis and Recognition, ICDAR 2021, Lausanne, Switzerland, September 5-10, 2021, Proceedings, Part I*, volume 12821 of *Lecture Notes in Computer Science*, pages 416–431. Springer.
- Brian L. Davis, Bryan S. Morse, Bryan Price, Chris Tensmeyer, Curtis Wigington, and Vlad I. Morariu. 2022. [End-to-end document recognition and understanding with dessurt](#). *CoRR*, abs/2203.16618.
- Mingfei Gao, Zeyuan Chen, Nikhil Naik, Kazuma Hashimoto, Caiming Xiong, and Ran Xu. 2021. [Field extraction from forms with unlabeled data](#). *CoRR*, abs/2110.04282.
- Andrea Gemelli, Sanket Biswas, Enrico Civitelli, Josep Lladós, and Simone Marinai. 2022a. [Doc2Graph: A task agnostic document understanding framework based on graph neural networks](#). *CoRR*, abs/2208.11168.
- Andrea Gemelli, Emanuele Vivoli, and Simone Marinai. 2022b. [Graph neural networks and representation embedding for table extraction in PDF documents](#). In *26th International Conference on Pattern Recognition, ICPR 2022, Montreal, QC, Canada, August 21-25, 2022*, pages 1719–1726. IEEE.
- Jiuxiang Gu, Jason Kuen, Vlad I. Morariu, Handong Zhao, Nikolaos Barmpalios, Rajiv Jain, Ani Nenkova, and Tong Sun. 2022. [Unified pretraining framework for document understanding](#). *CoRR*, abs/2204.10939.
- Jiuxiang Gu, Jason Kuen, Vlad I. Morariu, Handong Zhao, Rajiv Jain, Nikolaos Barmpalios, Ani Nenkova, and Tong Sun. 2021. [UniDoc: Unified pretraining framework for document understanding](#). In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 39–50.
- Jonathan Herzig, Pawel Krzysztof Nowak, Thomas Müller, Francesco Piccinno, and Julian Martin Eisen-schlos. 2020. [Tapas: Weakly supervised table parsing via pre-training](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 4320–4333. Association for Computational Linguistics.
- Teakgyu Hong, Donghyun Kim, Mingi Ji, Wonseok Hwang, Daehyun Nam, and Sungrae Park. 2022. [BROS: A pre-trained language model focusing on text and layout for better key information extraction from documents](#). In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 10767–10775. AAAI Press.
- Zheng Huang, Kai Chen, Jianhua He, Xiang Bai, Dimosthenis Karatzas, Shijian Lu, and C. V. Jawahar. 2019. [ICDAR2019 competition on scanned receipt OCR and information extraction](#). In *2019 International Conference on Document Analysis and Recognition, ICDAR 2019, Sydney, Australia, September 20-25, 2019*, pages 1516–1520. IEEE.
- Wonseok Hwang, Seonghyeon Kim, Minjoon Seo, Jinyeong Yim, Seunghyun Park, Sungrae Park, Junyeop Lee, Bado Lee, and Hwalsuk Lee. 2019. [Post-OCR parsing: Building simple and robust parser via BIO tagging](#). In *Workshop on Document Intelligence at NeurIPS 2019*.
- Wonseok Hwang, Hyunji Lee, Jinyeong Yim, Geewook Kim, and Minjoon Seo. 2021a. [Cost-effective end-to-end information extraction for semi-structured document images](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 3375–3383. Association for Computational Linguistics.
- Wonseok Hwang, Jinyeong Yim, Seunghyun Park, Sohee Yang, and Minjoon Seo. 2021b. [Spatial dependency parsing for semi-structured document information extraction](#). In *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021, Online Event, August 1-6, 2021*, volume ACL/IJCNLP 2021 of *Findings of ACL*, pages 330–343. Association for Computational Linguistics.
- Guillaume Jaume, Hazim Kemal Ekenel, and Jean-Philippe Thiran. 2019. [FUNSD: A dataset for form understanding in noisy scanned documents](#). In *2nd International Workshop on Open Services and Tools for Document Analysis, OST@ICDAR 2019, Sydney, Australia, September 22-25, 2019*, pages 1–6. IEEE.
- 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 *Computer Vision - ECCV*

- 2022 - 17th European Conference, Tel Aviv, Israel, October 23-27, 2022, *Proceedings, Part XXVIII*, volume 13688 of *Lecture Notes in Computer Science*, pages 498–517. Springer.
- Sandra Kübler, Ryan T. McDonald, and Joakim Nivre. 2009. *Dependency Parsing*. Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers.
- Kenton Lee, Mandar Joshi, Iulia Raluca Turc, Hexiang Hu, Fangyu Liu, Julian Martin Eisenschlos, Urvasi Khandelwal, Peter Shaw, Ming-Wei Chang, and Kristina Toutanova. 2023. [Pix2Struct: Screenshot parsing as pretraining for visual language understanding](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 18893–18912. PMLR.
- Chenliang Li, Bin Bi, Ming Yan, Wei Wang, Songfang Huang, Fei Huang, and Luo Si. 2021a. [Structurallm: Structural pre-training for form understanding](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 6309–6318. Association for Computational Linguistics.
- Xiangpeng Li, Bo Wu, Jingkuan Song, Lianli Gao, Pengpeng Zeng, and Chuang Gan. 2022. [Text-instance graph: Exploring the relational semantics for text-based visual question answering](#). *Pattern Recognit.*, 124:108455.
- Yulin Li, Yuxi Qian, Yuechen Yu, Xiameng Qin, Chengquan Zhang, Yan Liu, Kun Yao, Junyu Han, Jingtuo Liu, and Errui Ding. 2021b. [Structext: Structured text understanding with multi-modal transformers](#). In *MM '21: ACM Multimedia Conference, Virtual Event, China, October 20 - 24, 2021*, pages 1912–1920. ACM.
- Chuwei Luo, Guozhi Tang, Qi Zheng, Cong Yao, Lianwen Jin, Chenliang Li, Yang Xue, and Luo Si. 2022. [Bi-VLDoc: Bidirectional vision-language modeling for visually-rich document understanding](#). *CoRR*, abs/2206.13155.
- 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, ACL 2020, Online, July 5-10, 2020*, pages 6495–6504. Association for Computational Linguistics.
- Minesh Mathew, Dimosthenis Karatzas, and C. V. Jawahar. 2021. [DocVQA: A dataset for VQA on document images](#). In *IEEE Winter Conference on Applications of Computer Vision, WACV 2021, Waikoloa, HI, USA, January 3-8, 2021*, pages 2199–2208. IEEE.
- Puneet Mathur, Rajiv Jain, Ashutosh Mehra, Jiuxiang Gu, Franck Deroncourt, Quan Tran, Verena Kaynig-Fittkau, Ani Nenkova, Dinesh Manocha, Vlad I. Morariu, et al. 2023. [LayerDoc: Layer-wise extraction of spatial hierarchical structure in visually-rich documents](#). In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3610–3620.
- 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*.
- Shah Rukh Qasim, Hassan Mahmood, and Faisal Shafait. 2019. [Rethinking table recognition using graph neural networks](#). In *2019 International Conference on Document Analysis and Recognition, ICDAR 2019, Sydney, Australia, September 20-25, 2019*, pages 142–147. IEEE.
- Pau Riba, Anjan Dutta, Lutz Goldmann, Alicia Fornés, Oriol Ramos Terrades, and Josep Lladós. 2019. [Table detection in invoice documents by graph neural networks](#). In *2019 International Conference on Document Analysis and Recognition, ICDAR 2019, Sydney, Australia, September 20-25, 2019*, pages 122–127. IEEE.
- Mausoom Sarkar, Milan Aggarwal, Arneh Jain, Hires Gupta, and Balaji Krishnamurthy. 2020. [Document structure extraction using prior based high resolution hierarchical semantic segmentation](#). In *European Conference on Computer Vision*, pages 649–666. Springer.
- Sebastian Schreiber, Stefan Agne, Ivo Wolf, Andreas Dengel, and Sheraz Ahmed. 2017. [Deepdesrt: Deep learning for detection and structure recognition of tables in document images](#). In *14th IAPR International Conference on Document Analysis and Recognition, ICDAR 2017, Kyoto, Japan, November 9-15, 2017*, pages 1162–1167. IEEE.
- Štěpán Šimsa, Milan Šulc, Michal Uříčář, Yash Patel, Ahmed Hamdi, Matěj Kocián, Matyáš Skalický, Jiří Matas, Antoine Doucet, Mickaël Coustaty, and Dimosthenis Karatzas. 2023. [DocILE benchmark for document information localization and extraction](#).
- Rubèn Tito, Minesh Mathew, C. V. Jawahar, Ernest Valveny, and Dimosthenis Karatzas. 2021. [ICDAR 2021 competition on document visual question answering](#). In *16th International Conference on Document Analysis and Recognition, ICDAR 2021, Lausanne, Switzerland, September 5-10, 2021, Proceedings, Part IV*, volume 12824 of *Lecture Notes in Computer Science*, pages 635–649. Springer.
- Hieu M. Vu and Diep . Nguyen. 2020. [Revising FUNSD dataset for key-value detection in document images](#). *CoRR*, abs/2010.05322.

- Wenjin Wang, Zhengjie Huang, Bin Luo, Qianglong Chen, Qiming Peng, Yinxu Pan, Weichong Yin, Shikun Feng, Yu Sun, Dianhai Yu, and Yin Zhang. 2022. [mmLayout: Multi-grained multimodal transformer for document understanding](#). In *MM '22: The 30th ACM International Conference on Multimedia, Lisboa, Portugal, October 10 - 14, 2022*, pages 4877–4886. ACM.
- 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. 2021a. [LayoutLMv2: Multi-modal pre-training for visually-rich document understanding](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pages 2579–2591. Association for Computational Linguistics.
- Yiheng Xu, Tengchao Lv, Lei Cui, Guoxin Wang, Yijuan Lu, Dinei Florêncio, Cha Zhang, and Furu Wei. 2021b. [LayoutXLM: Multimodal pre-training for multilingual visually-rich document understanding](#). *CoRR*, abs/2104.08836.
- Yiheng Xu, Tengchao Lv, Lei Cui, Guoxin Wang, Yijuan Lu, Dinei Florencio, Cha Zhang, and Furu Wei. 2022. [XFUND: A benchmark dataset for multilingual visually rich form understanding](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3214–3224, Dublin, Ireland. Association for Computational Linguistics.
- Davide Zagami and Christopher Helm. 2022. [FUNSD+:](#) [A larger and revised funsd dataset](#).
- Richard Zanibbi, Dorothea Blostein, and James R. Cordy. 2004. A survey of table recognition: Models, observations, transformations, and inferences. *Document Analysis and Recognition*, 7:1–16.
- Xu Zhong, Elaheh ShafieiBavani, and Antonio Jimeno-Yepes. 2020. [Image-based table recognition: Data, model, and evaluation](#). In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XXI*, volume 12366 of *Lecture Notes in Computer Science*, pages 564–580. Springer.
- Xu Zhong, Jianbin Tang, and Antonio Jimeno-Yepes. 2019. [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.



## A Related Work

We discuss work regarding form information extraction (i.e., entity labeling and entity linking) in the main paper. Here, we briefly describe other relevant VRFU tasks as well as other domains for document parsing.

**Form Structure Extraction.** Complete form parsing is related to form structure extraction, a task which aims to learn the structure and type of field values in forms such as a text field or a checkbox (Aggarwal et al., 2020b,a; Sarkar et al., 2020; Gao et al., 2021; Mathur et al., 2023). The tasks differ in that structure extraction does not aim to extract the answers of a form (indeed the input forms do not need to be filled in), nor does it aim to necessarily learn hierarchical structure. As such, its datasets do not need to be filled in forms, and can just be form templates. We note that structure extraction is a possible approach to form parsing.

**Table Detection.** Tables are a natural occurrence in semi-structured documents and so have merited much study in the literature (Riba et al., 2019; Qasim et al., 2019; Zhong et al., 2019). There has been a plethora of work on table understanding that dates back several decades; (Zanibbi et al., 2004) provides a through analysis of initial methods for table detection. More recent work has leveraged image, text, and spatial features to train neural networks to achieve F1 scores of over 80 for cell and table header detection (Schreiber et al., 2017; Herzig et al., 2020; Zhong et al., 2020; Gemelli et al., 2022b).

**Receipt and Invoice Parsing.** Document parsing is the task of assigning structure to a document image. While work exists for a variety of document types such as namecards (Hwang et al., 2019, 2021a) and even forms (Davis et al., 2022), receipts and invoices have been the most prominent domain for document parsing (Huang et al., 2019; Majumder et al., 2020; Borchmann et al., 2021; Gao et al., 2021; Šimsa et al., 2023). This is largely due to shared characteristics among receipts (e.g., item names and prices) and their structured layout. Indeed, the end-to-end model used in this work, Donut (Kim et al., 2022), was chosen due to its effectiveness in invoice parsing.

## B Detailed TreeForm Annotation Scheme

In this section, we describe the TreeForm construction in more detail. Specifically we discuss two

versions of TreeForm: Concise and non-concise. We view forms as being comprised of four components: headers, question-answer pairs, tables, and free-form text.<sup>9</sup> In the non-concise TreeForm, we use tree leaves to capture all textual information of a form. Non-leaf nodes are then used to describe the type of textual information (e.g., question, answer) as well as structural information (e.g., question-answer pairs, table entries). A similar structure is suggested in the supplementary materials of (Davis et al., 2022). TreeForm differs in that leaves of the tree always correspond to the textual content of the form.<sup>10</sup> We present the non-concise version of TreeForm to show the underlying structure of the form; an example of this annotation scheme is given in Figure 2c. We hope that future annotations of TreeForm will not need this more verbose version and will immediately use the structure described in App. B.1.

**Headers.** Headers are represented by an intermediate header node with a connected leaf node containing the header value (if a value is given). A header node contains the sub-tree of all information that falls under that header (and associated section). Therefore, the header node associated with the form title (if one exists) is the root of the tree.

**Question-answer Pairs.** Question-answer pairs are represented by a chained question node and answer node. The question node has two children: a leaf node containing the question text, and an answer node that has its own leaf node containing the answer text. In this work, we have opted to omit answer-less questions and question-less answers, however, TreeForm can be extended to include these.

**Tables.** Tables can be viewed as a collection of entries where each entry has a potential header, and the same set of question-answer pairs. We thus represent tables through nested trees. Each entry to the table forms its own node, and contains a leaf node with the entry header or table header if either exist. We then connect each question-answer pair of the entry as described above. TreeForm considers ta-

<sup>9</sup>Question-answer pairs cover any single question response, including checkboxes or similar structures.

<sup>10</sup>We may also want to include bounding box information in TreeForm annotations, this can be done by attaching a child to each text value that contains the node’s bounding box. Note that (Davis et al., 2022) did not offer bounding box information in their proposed tree scheme.

bles to be read in row-major order, i.e., we consider each row to be an entry to the table with a potential entry header (row name), and each column to be a specific question that an entry answers.

**Free-form Text.** Much like recent work that ignores the *other* entity of FUNSD, we believe that free-form text does not capture the important structural information of a form, and so we do not include such components in TreeForm. Nevertheless, they can be included as leaf nodes connected to non-leaf nodes (e.g., headers) that they relate to.

### B.1 Concise TreeForm

The above structure contains additional information to enable structured prediction and evaluation. However, it is possible to push textual information up the tree to construct a much denser structure. We can replace intermediate header and answer nodes by their textual information to reduce both the depth and breadth of the tree. This can also be done for table entry names (if they exist). Similarly, we can condense question-answer pairs into a question node whose value is the question text with a child answer node whose value is the answer text. This concise TreeForm structure removes any unnecessary labeling; for a well structured and well annotated form, this would likely mean all node values will contain the content of the form. We provide an example of the concise TreeForm structure in Figure 2d.

## C Representation of Different Annotation Schemes

Figure 2 provides examples of JSON annotations for FUNSD, aggregated FUNSD (used for NAA and tree F1 metric), non-concise TreeForm, and concise TreeForm annotations.

## D Experimental Set-up

In this section, we discuss the necessary pre- and post-processing steps we took to fine-tune evaluate LayoutXLM and Donut for TreeForm.

### D.1 LayoutXLM

LayoutXLM (Xu et al., 2022) is a multilingual layout language model for document understanding. It is an extension of the popular LayoutLM models (Xu et al., 2021b,a) that were pre-trained for English documents only. The model takes as input the document image as well tokenized OCR output such that it knows the ground-truth text values in

the form. The fine-tuned model first predicts the entities which are then used to predict the entity links such that it provides FUNSD-like predictions.<sup>11</sup> We can then evaluate the model’s NAA and F1 performance as described in Section 3. We can further apply the same transformations to the predictions described in Section 4.2 to get the TreeForm predictions. Since bounding boxes are already known for each token, they can be used to apply the TreeForm transformations previously described.

### D.2 Donut

Donut (Kim et al., 2022) is an end-to-end (document) image-to-JSON model. It is pre-trained for a variety of VRDU tasks such as document classification, VQA, and document parsing. Specifically, Donut was shown to do well in complete document parsing for receipts and invoices, and so applying it to forms is a natural progression. We note that Donut was developed concurrently with a similar end-to-end model, Dessurt (Davis et al., 2022). While Dessurt builds on work more closely related to VRFU (Davis et al., 2019, 2021), we chose to run experiments using Donut as the model is more accessible for custom training purposes. Furthermore, Donut was pre-trained with additional synthetic data in Chinese, English, Korean, and Japanese while Dessurt was pre-trained purely on English data. Therefore, Donut is more suited to be fine-tuned for multiple languages than Dessurt. More recently, Pix2Struct (Lee et al., 2023) has been introduced and generally achieves better performance than Donut. However, much like Dessurt, it was only trained on English data and so we use Donut in this work.

In order to fine-tune the model<sup>12</sup> to predict FUNSD-like annotations, we create JSON objects that represent the trees constructed in Section 3. This does mean that if a node has multiple parents, its text is repeated several times in the JSON representation. To fine-tune Donut for TreeForm annotations, we use the non-concise TreeForm representation as the model expects all textual components to be leaves in the JSON tree. For both annotation

<sup>11</sup>We fine-tuned the model using the recommended commands and configurations given by (Xu et al., 2022) at <https://github.com/microsoft/unilm/tree/master/layoutxlm>. Due to resource availability, we use the base version of the model and only trained LayoutXLM with a single GPU. Consequently, our models did not reproduce the results of (Xu et al., 2022).

<sup>12</sup>Donut is available at <https://github.com/clovaai/donut>.

schemes, we follow the training configuration of (Kim et al., 2022) for the Consolidated Receipt Dataset (CORD) dataset (Park et al., 2019), which predicts a full parse tree for receipts. We make two minor changes to the training-set up. Firstly, we use a maximum sequence length of 1024 rather than 768 due to forms naturally containing more information than receipts. Secondly, we fine-tune on the transformed FUNSD dataset for 50 epochs rather than the 30 epochs used for CORD as our training set of 149 documents (per dataset) is much smaller than the 800 of CORD.

Donut is a generative model, and as such its input is not always formatted in line with the tree structure of the ground truth. Consequently, we applied a few greedy post-processing transformations that either created valid annotations, or discarded non-viable structures.<sup>13</sup> Furthermore, we observed that the model seemed to sometimes suffer from duplicating text and tree components.<sup>14</sup> As such, we applied a further greedy transformation that removed similar looking leaves and their paths if their normalized Levenshtein distance was greater than 0.6. When discarding a path, we always kept the path containing the longest text. We used a similar heuristic to remove repeated long entities (greater than 20 characters) for the FUNSD-type prediction. We only applied the heuristic for long entities as we expect more repetition as previously described.

---

<sup>13</sup>E.g., unanswered questions, empty trees, *inter alia*.

<sup>14</sup>We suspected this may be due to the increased maximum sequence length, but also saw similar issues when using 768 as in (Kim et al., 2022).