

# Detect-Order-Construct: A Tree Construction based Approach for Hierarchical Document Structure Analysis

Jiawei Wang<sup>a,b,1,\*</sup>, Kai Hu<sup>a,b,1,\*</sup>, Zhuoyao Zhong<sup>b,1,\*</sup>, Lei Sun<sup>b,1</sup>, Qiang Huo<sup>b</sup>

<sup>a</sup>*Department of EEIS, University of Science and Technology of China, Hefei, 230026, China*

<sup>b</sup>*Microsoft Research Asia, Beijing, 100080, China*

---

## Abstract

Document structure analysis (aka document layout analysis) is crucial for understanding the physical layout and logical structure of documents, with applications in information retrieval, document summarization, knowledge extraction, etc. In this paper, we concentrate on Hierarchical Document Structure Analysis (HDSA) to explore hierarchical relationships within structured documents created using authoring software employing hierarchical schemas, such as LaTeX, Microsoft Word, and HTML. To comprehensively analyze hierarchical document structures, we propose a tree construction based approach that addresses multiple subtasks concurrently, including page object detection (Detect), reading order prediction of identified objects (Order), and the construction of intended hierarchical structure (Construct). We present an effective end-to-end solution based on this framework to demonstrate its performance. To assess our approach, we develop a comprehensive benchmark called Comp-HRDoc, which evaluates the above subtasks simultaneously. Our end-to-end system achieves state-of-the-art performance on two large-scale document layout analysis datasets (PubLayNet and DocLayNet), a high-quality hierarchical document structure reconstruction dataset (HRDoc), and our Comp-HRDoc benchmark. The Comp-HRDoc benchmark is publicly available at <https://github.com/microsoft/CompHRDoc>.

**Keywords:** Document Layout Analysis, Table of Contents, Reading Order Prediction, Page Object Detection

---

## 1. Introduction

Document Structure Analysis (DSA) is a comprehensive process that identifies the fundamental components within a document, encompassing headings, paragraphs, lists, tables, and figures, and subsequently establishes the logical relationships and structures of these components. This process results in a structured

---

\*Corresponding author.

*Email addresses:* wangjiawei@mail.ustc.edu.cn (Jiawei Wang), hk970213@mail.ustc.edu.cn (Kai Hu), zhuoyao.zhong@gmail.com (Zhuoyao Zhong), kuangtongustc@gmail.com (Lei Sun), qianghuo@microsoft.com (Qiang Huo)

<sup>1</sup>Work done when Jiawei Wang and Kai Hu were interns, Zhuoyao Zhong and Lei Sun were FTEs at Multi-Modal Interaction Group, Microsoft Research Asia, Beijing, China.

representation of the document’s physical layout that accurately mirrors its logical structure, thereby enhancing the effectiveness and accessibility of information retrieval and processing. In a contemporary digital landscape, the majority of mainstream documents are structured creations, crafted using hierarchical-schema authoring software such as LaTeX, Microsoft Word, and HTML. Consequently, Hierarchical Document Structure Analysis (HDSA), which focuses on extracting and reconstructing the inherent hierarchical structures within these document layouts, has gained significant attention. However, despite its burgeoning popularity, HDSA poses a substantial challenge due to the diversity of document content and the intricate complexity of their layouts.

Over the past three decades, document structure analysis has garnered significant interest in the research community. Early research efforts primarily focused on physical layout analysis and logical structure analysis, employing various approaches such as knowledge-based [1], rule-based [2], model-based [3], and grammar-based [4] methods. However, these traditional methods face limitations in terms of effectiveness and scalability due to their susceptibility to noise, ambiguity, and difficulties in handling complex document collections. Furthermore, the absence of quantitative performance evaluations hinders the proper evaluation of these techniques. In the era of deep learning, a growing number of deep learning based approaches have been applied to the field of document structure analysis, leading to notable improvements in performance and robustness. However, these methods primarily focus on specific sub-tasks of DSA, such as Page Object Detection, Reading Order Prediction, and Table of Contents (TOC) Extraction, among others. Despite the substantial progress achieved in these individual sub-tasks, there remains a gap in the research community for a comprehensive end-to-end system or benchmark that addresses all aspects of document structure analysis concurrently. Filling this gap would significantly advance the field and encourage further research in this area.

Recently, hierarchical document structure analysis has gained traction with representative explorations like DocParser and HRDoc. DocParser [5] is the an end-to-end system for parsing document renderings into hierarchical document structures, encompassing all text elements, nested figures, tables, and table cell structures. Initially, the system employs Mask R-CNN [6] to detect all document entities within a document image. Subsequently, it devises a set of rules to predict two predefined relationships (i.e., “parent\_of” and “followed\_by”) between document entities to parse the complete physical structure of the document. However, the system does not take into account the logical structure of documents, such as the table of contents, and its reliance on a rule-based approach considerably limits its overall effectiveness and adaptability. On the other hand, HRDoc [7] proposed an encoder-decoder based hierarchical document structure parsing system (DSPS) to reconstruct the hierarchical structure of documents. This system employs a multi-modal bidirectional encoder and a structure-aware GRU decoder to predict the logical roles of the text-lines and the relationships between them. Although DSPS achieves significant performance improvements over DocParser and considers the logical structure of documents, it presumes that the reading order of the document

is provided, which is an essential aspect of document structure analysis that should not be overlooked. Furthermore, with the increase in text-lines within documents, the computational complexity of DSPS grows quadratically, presenting significant challenges when processing longer documents. Additionally, predicting relationships between line-level semantic units may result in the loss of broader contextual information, which is crucial for a comprehensive understanding of the document’s structure.

In this study, we propose a comprehensive approach to thoroughly analyzing hierarchical document structures using a tree construction based method. This method decomposes tree construction into three distinct stages, namely Detect, Order, and Construct, as illustrated in Fig. 1. Initially, given a set of document images, the Detect stage is dedicated to identifying all page objects and assigning a logical role to each object, thereby forming the nodes of the hierarchical document structure tree. Following this, the Order stage establishes the reading order relationships among these nodes, which corresponds to a pre-order traversal of the hierarchical document structure tree. Finally, the Construct stage identifies hierarchical relationships (e.g., Table of Contents) between semantic units to construct an abstract hierarchical document structure tree. By integrating the results of all three stages, we can effectively construct a complete hierarchical document structure tree, facilitating a more comprehensive understanding of complex documents.

To demonstrate its performance, we present an effective end-to-end solution based on this framework. For the Detect stage, we consider OCR’d text-lines as the basic semantic units and introduce a novel hybrid method, which combines a top-down model with a relation prediction model to simultaneously detect graphical page objects (e.g., tables, figures, etc.), group text-lines into text regions according to the intra-region reading order, and recognize the logical roles of text regions. Any top-down object detection or instance segmentation models can be directly applied to detecting graphical page objects, sharing a visual backbone network with the relation prediction model. Subsequently, we can cohesively formalize these three stages as relation prediction tasks by defining distinct types of relationships. They include the intra-region reading order relationships between text-lines to group text-lines into text regions, inter-region reading order relationships between text regions to generate the reading sequence of text regions, and TOC relationships between section headings to summarize the overall hierarchical document structure. To address these stages in a unified manner, we introduce a type of multi-modal transformer-based relation prediction models, which are designed to tackle all three stages. This novel relation prediction model approaches the relation prediction as a dependency parsing task, employing a multi-modal transformer encoder to model the interactions between input pairs via a global self-attention mechanism. Moreover, in response to the chain structure of reading order and the tree structure of table of contents, we design two structure-aware relation prediction models specifically tailored for these two structures, ensuring a more accurate and efficient analysis of these hierarchical relationships.

Throughout these three stages, several sub-tasks play an integral role in hierarchical document structure analysis. Consequently, during the performance evaluation phase, it is not sufficient to merely assess the

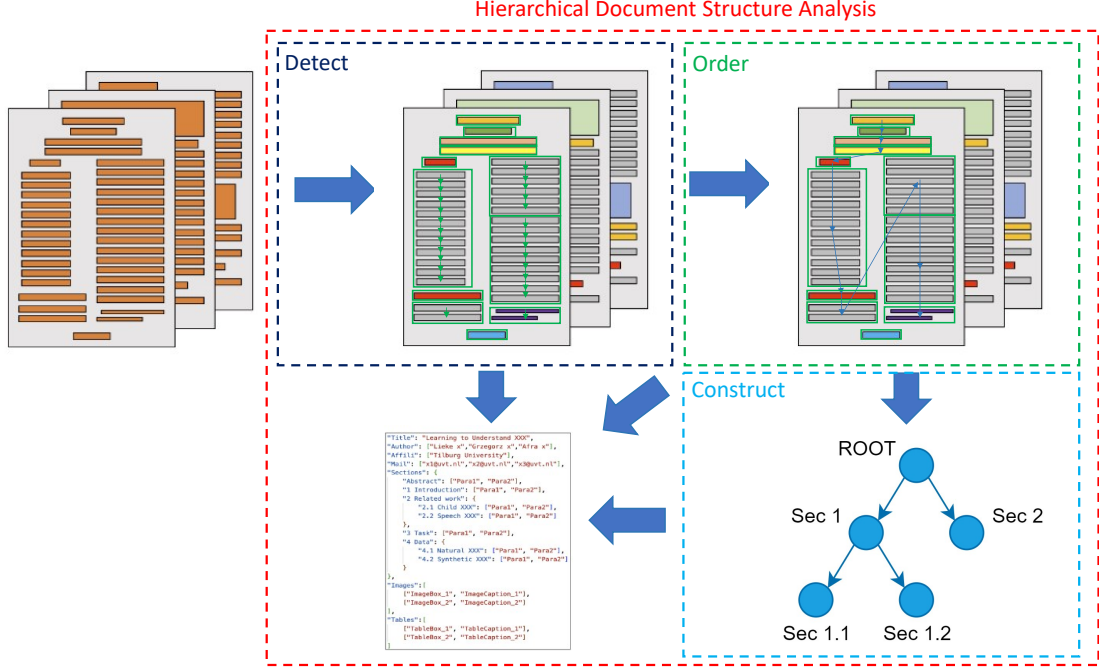


Figure 1: Overview of our tree construction based approach, named Detect-Order-Construct, for hierarchical document structure analysis.

overall accuracy of hierarchical document structure reconstruction, as done in HRDoc [7]. An exhaustive and thorough evaluation of each sub-task involved is equally important. Leveraging the HRDoc dataset, we establish a comprehensive benchmark, Comp-HRDoc, aimed at evaluating page object detection, reading order prediction, table of contents extraction, and hierarchical structure reconstruction concurrently. Extensive experimental results demonstrate that our proposed end-to-end system achieves state-of-the-art performance on two large-scale document layout analysis datasets (i.e., PubLayNet [8] and DocLayNet [9]), and a hierarchical document structure reconstruction dataset (i.e., HRDoc). Moreover, our proposed comprehensive benchmark, Comp-HRDoc, effectively illustrate the effectiveness and superiority of our approach across all sub-tasks.

The main contributions of this paper are as follows:

- Proposed a tree construction based approach, namely Detect-Order-Construct, for hierarchical document structure analysis. To exemplify the effectiveness of this framework, we devise an effective end-to-end solution by casting uniformly the three-stage tasks as relation prediction problems. Furthermore, we design multi-modal transformer-based relation prediction models with two structure-aware improvements for chain structures and tree structures respectively to enhance the overall system performance.

- Designed and established the first comprehensive benchmark, namely Comp-HRDoc, for the simultaneous evaluation of page object detection, reading order prediction, table of contents extraction, and hierarchical structure reconstruction.
- Our proposed end-to-end system achieves state-of-the-art performance on two large-scale document layout analysis datasets (i.e., PubLayNet and DocLayNet), a hierarchical document structure reconstruction dataset (i.e., HRDoc) and our comprehensive benchmark Comp-HRDoc.

Although a preliminary study of the Detect stage in our end-to-end system has been presented in our conference paper [10], this paper significantly extends it in the following aspects: (1) A tree construction based approach, namely Detect-Order-Construct, is proposed for hierarchical document structure analysis; (2) A comprehensive benchmark is designed and established to simultaneously evaluate page object detection, reading order prediction, table of contents extraction, and hierarchical structure reconstruction; (3) Experimental results on a public benchmark dataset HRDoc [7] and our new benchmark Comp-HRDoc are presented to compare our approach with other works more comprehensively.

## 2. Related Work

Since the 1980s, numerous studies have been conducted on document structure analysis, which can be categorized into physical structure analysis (or physical layout analysis) and logical structure analysis [11]. Physical layout analysis focuses on identifying homogeneous regions of interest, also known as page objects, while logical structure analysis aims to assign logical roles to these regions and determine their relationships. Early approaches to document structure analysis, mainly based on heuristic rules or grammar analysis, can be found in surveys [11, 12]. In the past decade, a growing body of research [8, 9, 13, 14] has focused on document layout analysis, specifically physical layout analysis and logical role classification, which is also known as page object detection [15]. To maintain clarity, we will consistently use the term **“page object detection”** throughout this article to refer to the document layout analysis task that incorporates both physical layout analysis and logical role classification. In addition to detecting page objects, numerous research studies have delved into the logical relationships between components within documents. These investigations have focused on aspects such as the reading order relationships and the organization of tables of contents. In this section, we primarily review and analyze recent developments in page object detection, reading order prediction, and hierarchical document structure reconstruction, providing an overview of the latest advancements and methodologies in these areas.

### 2.1. Page Object Detection

Page Object Detection, also known as POD [15], is a task that involves locating logical objects (e.g., paragraphs, tables, mathematical equations, graphics, and figures) within document pages. Deep learning-

based POD approaches can be broadly classified into three categories: object detection-based methods, semantic segmentation-based methods, and graph-based methods.

**Object detection-based methods.** These methods leverage the latest top-down object detection or instance segmentation frameworks to address the page object detection problem. Pioneering efforts by Yi et al. [16] and Oliveira et al. [17] adapted R-CNN [18] to identify and recognize page objects from document images. However, their performance was hindered by the limitations of traditional region proposal generation strategies. Subsequent research explored more sophisticated object detectors, such as Fast R-CNN [19], Faster R-CNN [20], Mask R-CNN [6], Cascade R-CNN [21], SOLOv2 [22], CondInst [23], YOLOv5 [24], and Deformable DETR [25] as investigated by Vo et al. [26], Zhong et al. [8], Saha et al. [27], Li et al. [28], Biswas et al. [29], Hu et al. [30], Pfizmann et al. [9], and Yang et al. [31], respectively. In addition, researchers have proposed effective techniques to enhance the performance of these detectors. For example, Zhang et al. [32] introduced a multi-modal Faster/Mask R-CNN model for page object detection that fused visual feature maps extracted by CNN with two 2D text embedding maps containing sentence and character embeddings. They also incorporated a graph neural network (GNN) based relation module to model the interactions between page object candidates. Shi et al. [33] proposed a novel lateral feature enhancement backbone network, while Yang et al. [31] employed Swin Transformer [34] as a more robust backbone network to boost the performance of Mask R-CNN and Deformable DETR for page object detection. Recently, Gu et al. [35], Li et al. [28], and Huang et al. [36] further improved the performance of Faster R-CNN, Mask R-CNN, and Cascade R-CNN-based page object detectors by pre-training the vision backbone networks on large-scale document images using self-supervised learning algorithms. Despite achieving state-of-the-art results on several benchmark datasets, these methods continue to face challenges in detecting small-scale text regions.

**Semantic segmentation based methods.** These methods, such as those proposed by Yang et al. [13], He et al. [37], Li et al. [38, 39], and Sang et al. [40], typically employ existing semantic segmentation frameworks, such as FCN [41], to initially generate a pixel-level segmentation mask. Subsequently, the pixels are merged to form distinct types of page objects. Yang et al. [13] introduced a multi-modal FCN for page object segmentation, which combined visual feature maps and 2D text embedding maps with sentence embeddings to enhance pixel-wise classification accuracy. He et al. [37] developed a multi-scale, multi-task FCN designed to concurrently predict a region segmentation mask and a contour segmentation mask. After refinement using a conditional random field (CRF) model, these two segmentation masks are processed by a post-processing module to obtain the final prediction results. Li et al. [39] integrated label pyramids and deep watershed transformation into the vanilla FCN structure to prevent the merging of adjacent page objects. Despite their advancements, the performance of existing semantic segmentation-based methods remains inferior to that of the other two categories of approaches when evaluated on recent document layout analysis benchmarks.

**Graph-based methods.** These approaches (e.g., [42–45]) represent each document page as a graph, where the nodes correspond to primitive page objects (e.g., words, text-lines, connected components), and the edges denote relationships between neighboring primitive page objects. The detection of page objects is then formulated as a graph labeling problem. Li et al. [42] employed image processing techniques to initially generate line regions, followed by the application of two CRF models to classify these regions into distinct types and predict whether pairs of line regions belong to the same instance, based on visual features extracted by CNNs. Subsequently, line regions that share the same class and instance are merged to form page objects. In their follow-up work [43], Li et al. replaced line regions with connected components as nodes and implemented a graph attention network (GAT) to enhance the visual features of both nodes and edges. Luo et al. [44] concentrated on the logical role classification task, proposing the use of multi-aspect graph convolutional networks (GCNs) to identify the logical role of each page object by leveraging syntactic, semantic, density, and appearance features. More recently, Wang et al. [45] focused on paragraph identification, developing a GCN-based approach to group text-lines into paragraphs. Liu et al. [46], Long et al. [47], and Xue et al. [48] further proposed a unified framework for text detection and paragraph (text-block) identification.

## 2.2. Reading Order Prediction

The objective of reading order prediction is to determine the appropriate reading sequence for documents. Generally, humans tend to read documents in a left-to-right and top-to-bottom manner. However, such simplistic sorting rules may prove inadequate when applied to complex documents with tokens extracted by OCR tools. Previous research has attempted to tackle the reading order issue using a variety of approaches. As categorized by Wang et al. [49], these methods can be broadly classified into rule-based sorting and machine learning-based sequence prediction, among others.

**Rule-based sorting.** Topological sorting, first introduced by Breuel [50], has been utilized for document layout analysis. In this method, partial orders are determined based on the x/y interval overlaps between text lines, enabling the generation of reading order patterns for multi-column text layouts. A bidimensional relation rule, proposed in [51], offers similar topological rules while also incorporating a row-wise rule by inverting the x/y axes from column-wise. In the same vein, an argumentation-based approach in [52] utilizes rules derived from relationships between text blocks. For text layouts with hierarchies and larger sizes, XY-Cut [53, 54] can serve as an efficient method to order all text blocks from top to bottom and left to right for specific layout types. Despite their effectiveness in certain scenarios, these rule-based methods can be prone to failure when confronted with out-of-domain cases.

**Machine learning-based sequence prediction.** Designed to learn from training examples across various domains, machine learning-based approaches aim to provide a general solution for reading order prediction. Ceci et al. [55] introduced a probabilistic classifier within the Bayesian framework, which is

capable of reconstructing single or multiple chains of layout components based on learned partial orders. Differently, an inductive logic programming (ILP) learning algorithm was applied in [56] to learn two kinds of predicates, *first\_to\_read/1* and *succ\_in\_reading/2*, thereby establishing an ordering relationship. In recent years, deep learning models have emerged as the leading solution for numerous machine learning challenges. Li et al. [57] proposed an end-to-end OCR text reorganizing model, using a graph convolutional encoder and a pointer network decoder to reorder text blocks. LayoutReader [58] introduced a benchmark dataset called ReadingBank, which contains reading order, text, and layout information, and employed a transformer-based architecture on spatial-text features to predict the reading order sequence of words. However, the decoding speed of these auto-regressive-based methods is limited when applied to rich text documents. Recently, Quir’os et al. [59] followed the idea of assuming a pairwise partial order at the element level from [50] and proposed two new reading-order decoding algorithms for reading order prediction on handwritten documents. They also provided a theoretical background for these algorithms. A significant limitation of this approach is that the partial order between two entities is determined solely by pair-wise spatial features, without considering the visual information and textual information.

### 2.3. Hierarchical Document Structure Reconstruction

The process of reconstructing a document’s hierarchical structure aims to recover its logical structure, which conveys semantic information beyond the character strings that comprise its contents. Table of Contents is a crucial component in reconstructing the hierarchical structure. Consequently, existing research studies on hierarchical structure reconstruction can be broadly categorized into two groups. The first group primarily focuses on extracting the table of contents within documents. The second group places emphasis on overall structure reconstruction of a document.

**Table of Contents.** Table of contents extraction is the task of restoring the structure of a document and recognizing the hierarchy of its sections. It is a challenging task due to the diversity of TOC styles and layouts. Early methods relied on heuristic rules derived from small data sets for specific domains, which were not effective in large-scale heterogeneous documents. Wu et al. [60] identified three basic TOC styles: “flat”, “ordered”, and “divided”. Based on these styles, they proposed an approach for TOC recognition that adaptively selects appropriate rules according to the basic TOC style features. However, this method assumes the existence of a Table of Contents page within the documents. Nguyen et al. [61] proposed a system that combines a TOC page detection method with a link-based TOC reconstruction method to address the TOC extraction problem. Cao et al. [62] developed a framework called Hierarchy Extraction from Long Document (HELD) to tackle the problem of TOC extraction in long documents. This approach sequentially inserts each section heading into the TOC tree at the correct position, considering sibling and parent information using LSTM [63]. Recently, Hu et al. [64] proposed an end-to-end model by using a multimodal tree decoder (MTD) for table of contents extraction. The MTD model fuses multimodal features

for each entity of the document and parses the hierarchical relationship by a tree-structured decoder.

**Overall Structure Reconstruction.** To reconstruct the overall structure of a document, it is critical to represent the structure and layout of the document. Intuitively, graph representation for document structure is most general and can encapsulate the relationship between regions and their properties. The graph representation, however, fails to capture the hierarchical nature of a document structure and layout. Also, it is hard to define a complete graph representation for a document. To accomplish this, one could use a rooted tree for representing document layout and logical structure [65]. One of the most powerful ways to express hierarchical structures is to use formal grammars [66]. The class of regular and context-free grammars are extremely useful in describing the structure of most documents. However, there could be multiple derivations corresponding to a particular sequence of terminals. This would mean multiple interpretations of the structure or layout. Tateisi et al. [67] proposed a stochastic grammar to integrate multiple evidences and estimate the most probable parse or interpretation of a given document. Despite its usefulness, stochastic grammars may lack the flexibility to model complex patterns and structures, particularly when handling highly diverse data. In recent years, some deep learning based methods are proposed for tree-based document structure reconstruction. Wang et al. [68] concentrated on form understanding task, treating the form structure as a tree-like hierarchy composed of text fragments. To predict the relationship between each pair of text fragments, they employed an asymmetric parameter matrix. However, this approach resulted in high computational complexity when dealing with documents containing a large number of text fragments. DocParser, as proposed by Rausch et al. [5], presented an end-to-end system designed to parse the complete physical structure of documents including all text elements, nested figures, tables, and table cell structures. This system employed rule-based algorithms for relation classification and inferred document structures in a holistic, principled manner. Nonetheless, the system did not consider the logical hierarchical structure of documents, such as the table of contents, and the reliance on a rule-based approach significantly constrained its overall effectiveness and adaptability. Recently, Ma et al. [7] introduced hierarchical reconstruction of document structures as a novel task and built a large-scale dataset, named HRDoc. Moreover, an encoder-decoder-based hierarchical document structure parsing system (DSPS) was proposed to reconstruct the hierarchical structure. While taking into account the logical structure of the document, this task presumes that the reading order is provided. Furthermore, DSPS directly predicts relationships between text-lines, resulting in low representational ability and high computational cost. In this work, we also consider overall structure reconstruction to be a recovery of the structure of the rooted tree of a document.

### 3. Problem Definition

The majority of document types, such as scientific papers, books, reports, and legal documents, typically exhibit a hierarchical document structure in a tree-like format. In this structure, the nodes within the

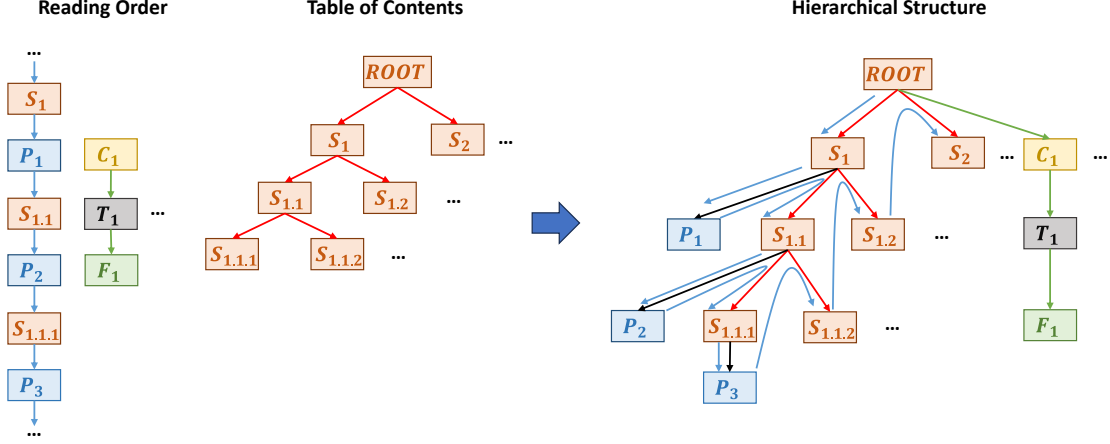


Figure 2: Hierarchical structure reconstruction of a document by integrating the Reading Order and Table of Contents. Blue arrows demonstrate the Text Region Reading Order Relationship, green arrows show the Graphical Region Relationship, and red arrows signify the TOC Relationship. The nodes “P”, “S”, “C”, “T” and “F” represent Paragraph, Section heading, Caption, Table and Footnote, respectively.

tree represent various page objects (e.g., section, paragraph, figure, caption) of the document, while the edges signify the hierarchical relationships and connections between these page objects. Given a multi-page document  $D$  comprised of  $D_1, D_2, \dots, D_n$ , where  $D_i$  represents an individual page within document  $D$ , the primary objective of hierarchical document structure analysis is to reconstruct its hierarchical structure tree  $H$ , consisting of both page objects and hierarchical relationships as follows:

**Page Objects** ( $O_i, i = 1, \dots, m$ ) refer to the various page objects within document  $D$ . Each page object is described by three attributes: 1) its logical role category  $\mathbf{c}_i \in C$  (e.g., title, section heading, table, figure, etc.); 2) its bounding box coordinates  $\mathbf{b}_i$ ; 3) its basic semantic units (not useful for graphical page objects and we use OCR’d text-lines as basic semantic units).

**Hierarchical Relationships** ( $R_{ij}, i, j = 1, \dots, m$ ) describe the relationships between page object pairs and are represented by triplets  $(O_i, r_{ij}, O_j)$ . Each triplet includes a subject page object  $O_i$ , an object page object  $O_j$ , and a relation type  $r_{ij} \in \Phi$ . Based on the categories of  $O_i$  and  $O_j$ , we define the following three relationship types: 1) **Text Region Reading Order Relationship** between main body text regions, 2) **Graphical Region Relationship** between caption, footnote and graphical page objects, i.e., tables or figures; 3) **Table of Contents Relationship** between section heading regions.

The combination of page objects and hierarchical relationships is sufficient to reconstruct the hierarchical tree  $H$  for a document, as illustrated in Fig. 2. Conversely, the hierarchy tree  $H$  can be used to extract various hierarchical relationships as needed, further emphasizing its importance in the process of hierarchical document structure analysis. For instance, the reading order sequence can be obtained by performing a pre-order traversal on the hierarchical tree  $H$ . Based on the problem description and objectives of hierarchical

document structure analysis, we divide it into following three distinct sub-tasks, which correspond to our proposed three-stage framework:

- **Page Object Detection** (Detect stage) aims to identify individual page object  $O_i$  (e.g., text regions, images, tables) within each page of the document  $D$  and assign a logical role to each detected page object (e.g., section headings, captions, footnotes).
- **Reading Order Prediction** (Order stage) focuses on determining the reading sequence of detected page objects based on their spatial arrangement within the document  $D$ . The reading order is represented as a permutation of the indices of the detected page objects.
- **Table of Contents Extraction** (Construct stage) aims to extract the table of contents within document  $D$ , which involves constructing a hierarchy tree that summarizes the overall hierarchical structure  $H$ . The hierarchy tree comprises a list of section headings and their hierarchical levels.

By integrating the results from all three sub-tasks, the hierarchical document structure tree  $H$  can be effectively reconstructed, offering a more comprehensive understanding of complex documents.

## 4. Methodology

### 4.1. Overview

Our newly proposed tree construction based approach for hierarchical document structure analysis, named Detect-Order-Construct, is illustrated in Fig. 1. This approach comprises three main components: 1) A **Detect** stage that identifies individual page objects within the document rendering and assigns a logical role to each detected page object (i.e., page object detection); 2) An **Order** stage responsible for determining the sequential order of the page objects (i.e., reading order prediction); and 3) A **Construct** stage that extracts the abstract hierarchy tree (i.e., table of contents extraction). By integrating the outputs from the previous tasks, we can effectively reconstruct a complete hierarchical document structure tree (i.e., hierarchical document structure reconstruction).

In our approach, we uniformly define the tasks of these three stages as relation prediction problems and present a type of multi-modal, transformer-based relation prediction models to tackle all tasks effectively. Our proposed relation prediction model approaches relation prediction as a dependency parsing task and incorporates structure-aware designs that align with the chain structure of reading order and the tree structure of table of contents. Utilizing our novel techniques and the proposed framework, we develop an effective end-to-end solution for hierarchical document structure analysis, which comprises three modules: the Detect module, the Order module, and the Construct module. We elaborate on the details of these three modules in Sections 4.2, 4.3, and 4.4, respectively.

Method	Recall	Prec	MT	ML	FP	FN	IDs	PM	MOTA	MOTP
Kalman-HA	28.5	70.0	32	334	2,033	28,320	685	837	19.2	60.9
Kalman-HA <sup>2</sup>	28.3	83.4	39	354	2,545	28,626	105	342	22.4	69.4
PDPA <sup>1</sup>	30.6	81.7	38	348	2,728	27,707	109	380	23.5	69.0
RNN-HA	37.5	85.2	50	287	1,984	24,420	515	883	24.0	68.7
RNN-LSTM	37.1	73.5	50	260	5,127	25,094	572	983	22.3	69.0

Table 1: Tracking results on the MOTChallenge training dataset. <sup>1</sup>Densets offline post-processing.

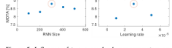


Figure 5: Influence of two regular hyper-parameters on the overall performance on the MOTChallenge benchmark, measured by MOTA. The optimal parameter is marked with a red circle. Note that the graph shows the performance of our prediction-based RNN block for only one target ( $N=1$ ), which explains the relatively low MOTA.

#### Training Data

It is well known that deep architectures require vast amounts of training data to avoid overfitting the model. Hager labeled datasets like ImageNet [Russakovsky et al., 2015] or Microsoft COCO [Lin et al., 2014] have enabled deep learning methods to realize their potential on tasks like image classification or pixel labelling. Unfortunately, mostly due to the very tedious and time-consuming task of video annotation, only very limited amount of labeled data for pedestrian tracking is publicly available today. We therefore resort to synthetic generation by sampling from a single generative trajectory model learned from real data. To that end, we first learn a trajectory model from task training sequences. For simplicity, we only estimate the mean and the variance of two features: the start location  $x$ , and the average velocity  $v$  from all annotated trajectories in that sequence. For each training sample we then generate regions  $\gamma$  tracks by sampling from a normal distribution with the learned parameters. Note that this simplified approach enables easy generation of arbitrary data, but does not accommodate any observations.

#### Implementation Details

We implemented our framework in Lua and Torch7. Both our entire code base as well as pre-trained models are publicly available.<sup>1</sup> Finding correct hyper-parameters for deep architectures still remains a non-trivial task [Gref et al., 2015]. In this section we will point out some of the most important parameters and implementation details. We follow most of the best practices found in the literature [Gref et al., 2015; Kargicli, Johnson, and Li, 2015], such as setting the initial weights for the larger parts higher (1) in our case, and also employ a standard grid search to find the best setting for the present task.

<sup>1</sup><https://github.com/sergey-antoniou/track-evaluation>

**Network size.** The RNN for state estimation and track management is trained with one layer and 500 hidden units. The data association is a more complex task, requiring more representational power. In that case, the LSTM module employed to learn the data association consists of two layers and 500 hidden units.

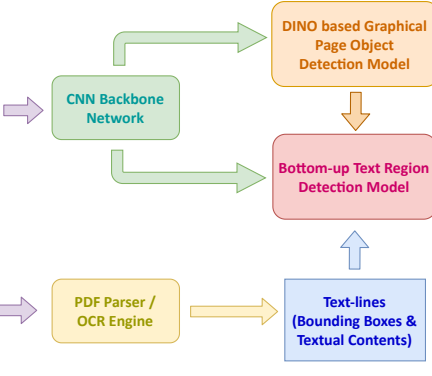
**Optimization.** We use the RMSprop (Tieleman and Hinton 2012) to minimize the loss. The learning rate is set initially to 0.0001 and is decreased by 1% every 20,000 iterations. We set the maximum number of iterations to 200,000, which is enough to reach convergence. The training of both modules takes approximately 20 hours on a GTX 1080. With a more accurate representation and the use of GPUs we believe that training can be sped up significantly.

**Data.** The RNN is trained with approximately 100K 20-frame long sequences. The data is divided into mini-batches of 10 samples per batch and normalized to the range  $[-0.5, 0.5]$ , w.r.t. the image dimensions. We experimented with the more popular row-major and unit variance data normalization but found that the first one based on the image size yields superior performance.

#### Experiments

To demonstrate the functionality of our approach, we first perform experiments on simulated data. Fig. 4 shows an example of the tracking results on synthetic data. Here, five targets with random birth and death times are generated in a rather cluttered environment. The initiation / termination indicators are illustrated in the bottom row.

We further test our approach on real-world data, using the MOTChallenge 2015 benchmark [Gref et al., 2015]. The pedestrian tracking dataset is a collection of 22 video sequences (1111 for training and testing, respectively) with a relatively high variation in target motion, camera motion, viewing angle and scene density. The evaluation is performed on a server using supervised ground truth. Next to precision and recall, we show the number of mostly missed trajectories (M<sub>1</sub>, Huang, and Nussim 2009), the number of false positives (FP), false negatives (FN) against identity swaps (IDS) and track fragmentations (FM). MOTA and MOTP are the widely used CLEAR metrics [Bennett and Smeets 2008] and summarize the tracking accuracy and precision, respectively. Across each task metric, indicate whether higher (or lower) values are better.



Method	Recall	Prec	MT	ML	FP	FN	IDs	PM	MOTA	MOTP
Kalman-HA	28.3	70.0	32	334	2,033	28,320	685	837	19.2	60.9
Kalman-HA <sup>2</sup>	28.3	83.4	39	354	2,545	28,626	105	342	22.4	69.4
PDPA <sup>1</sup>	30.6	81.7	38	348	2,728	27,707	109	380	23.5	69.0
RNN-HA	37.5	85.2	50	287	1,984	24,420	515	883	24.0	68.7
RNN-LSTM	37.1	73.5	50	260	5,127	25,094	572	983	22.3	69.0

Table 1: Tracking results on the MOTChallenge training dataset. <sup>1</sup>Densets offline post-processing.

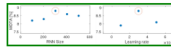


Figure 5: Influence of two regular hyper-parameters on the overall performance on the MOTChallenge benchmark, measured by MOTA. The optimal parameter is marked with a red circle. Note that the graph shows the performance of our prediction-based RNN block for only one target ( $N=1$ ), which explains the relatively low MOTA.

#### Training Data

It is well known that deep architectures require vast amounts of training data to avoid overfitting the model. Hager labeled datasets like ImageNet [Russakovsky et al., 2015] or Microsoft COCO [Lin et al., 2014] have enabled deep learning methods to realize their potential on tasks like image classification or pixel labelling. Unfortunately, mostly due to the very tedious and time-consuming task of video annotation, only very limited amount of labeled data for pedestrian tracking is publicly available today. We therefore resort to synthetic generation by sampling from a single generative trajectory model learned from real data. To that end, we first learn a trajectory model from task training sequences. For simplicity, we only estimate the mean and the variance of two features: the start location  $x$ , and the average velocity  $v$  from all annotated trajectories in that sequence. For each training sample we then generate regions  $\gamma$  tracks by sampling from a normal distribution with the learned parameters. Note that this simplified approach enables easy generation of arbitrary data, but does not accommodate any observations.

#### Implementation Details

We implemented our framework in Lua and Torch7. Both our entire code base as well as pre-trained models are publicly available.<sup>1</sup> Finding correct hyper-parameters for deep architectures still remains a non-trivial task [Gref et al., 2015]. In this section we will point out some of the most important parameters and implementation details. We follow most of the best practices found in the literature [Gref et al., 2015; Kargicli, Johnson, and Li, 2015], such as setting the initial weights for the larger parts higher (1) in our case, and also employ a standard grid search to find the best setting for the present task.

<sup>1</sup><https://github.com/sergey-antoniou/track-evaluation>

Figure 3: The overall architecture of our Detect module.

## 4.2. Detect Module

The proposed Detect module consists of three primary components: 1) A shared visual backbone network designed to extract multi-scale feature maps from input document images; 2) A top-down graphical page object detection model for detecting graphical page objects, such as tables, figures, and displayed formulas; 3) A bottom-up text region detection model that groups text-lines located outside graphical page objects into text regions, based on the intra-region reading order, and identifies the logical role of each text region. The overall architecture of the Detect module is illustrated in Fig. 3. In our conference paper [10], we selected a ResNet-50 network as the backbone network to generate multi-scale feature maps and the DINO [69] as the top-down graphical page object detector to localize these graphical objects. However, any suitable visual backbone network and object detection or instance segmentation model can be readily incorporated into our Detect module. In this paper, we primarily concentrate on the details of the newly proposed *Bottom-up Text Region Detection Model*.

A text region is a semantic unit of writing that comprises a group of text-lines arranged in natural reading order and associated with a logical label, such as paragraph, list/list-item, title, section heading, header, footer, footnote, and caption. Given a document page rendering  $D_i$  composed of  $n$  text-lines  $[t_1, t_2, \dots, t_n]$ , the objective of our bottom-up text region detection model is to group these text-lines into distinct text regions according to the intra-region reading order and to recognize the logical role of each text region. In this study, we assume that the bounding boxes and textual contents of text-lines have already been provided by a PDF parser or OCR engine. Based on the detection results of the top-down graphical page object detection model, we initially filter out those text-lines located inside graphical page objects and then

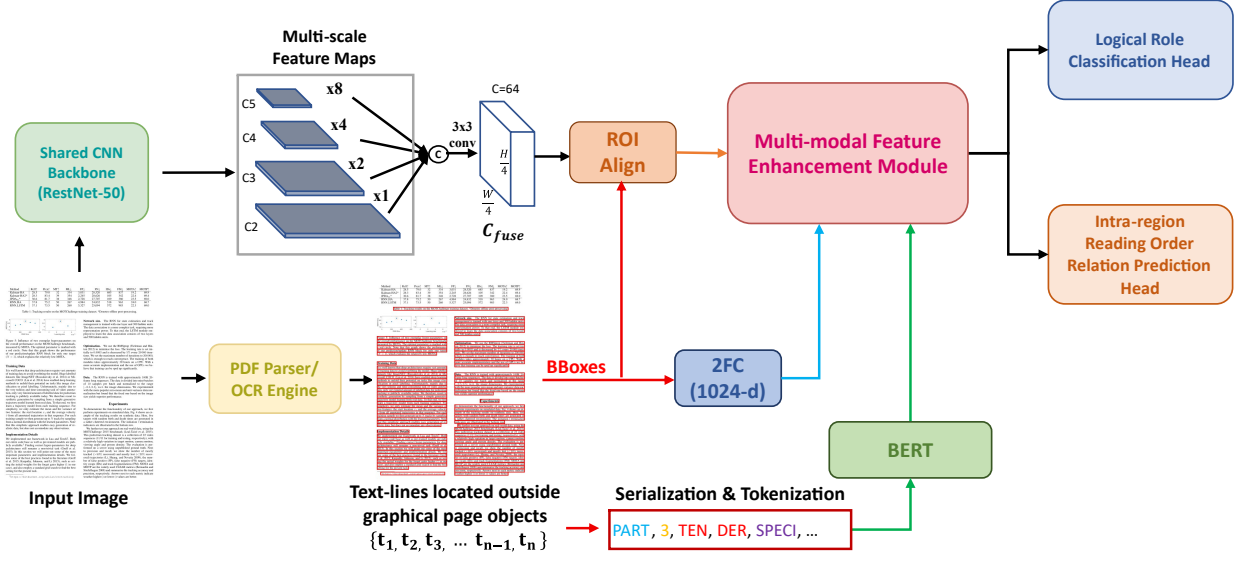


Figure 4: A schematic view of the proposed bottom-up text region detection model.

utilize the remaining text-lines as input. As depicted in Fig. 4, our bottom-up text region detection model consists of a multi-modal feature extraction module, a multi-modal feature enhancement module, and two prediction heads, i.e., an intra-region reading order relation prediction head and a logical role classification head. The detailed illustrations of the multi-modal feature enhancement module and the two prediction heads can be found in Fig. 5.

#### 4.2.1. Multi-modal Feature Extraction Module

In this module, we extract the visual embedding, text embedding, and 2D Positional Embedding for each text-line.

**Visual Embedding.** As shown in Fig. 4, we first resize  $C_4$  and  $C_5$  to the size of  $C_3$  and then concatenate these three feature maps along the channel axis, which are fed into a  $3 \times 3$  convolutional layer to generate a feature map  $C_{fuse}$  with 256 channels. For each text-line  $t_i$ , we adopt the RoIAlign algorithm [6] to extract  $7 \times 7$  feature maps from  $C_{fuse}$  based on its bounding box  $b_{t_i} = (x_i^1, y_i^1, x_i^2, y_i^2)$ , where  $(x_i^1, y_i^1)$ ,  $(x_i^2, y_i^2)$  represent the coordinates of its upper left and bottom right corners, respectively. The final visual embedding  $V_{t_i}$  of  $t_i$  can be represented as:

$$V_{t_i} = LN(ReLU(FC(ROIAlign(C_{fuse}, b_{t_i})))), \quad (1)$$

where FC is a fully-connected layer with 1,024 nodes and LN represents Layer Normalization [70].

**Text Embedding.** We leverage the pre-trained language model BERT [71] to extract the text embedding of each text-line. Specifically, we first serialize all the text-lines in a document image into a 1D sequence by reading them in a top-left to bottom-right order and tokenize the text-line sequence into a

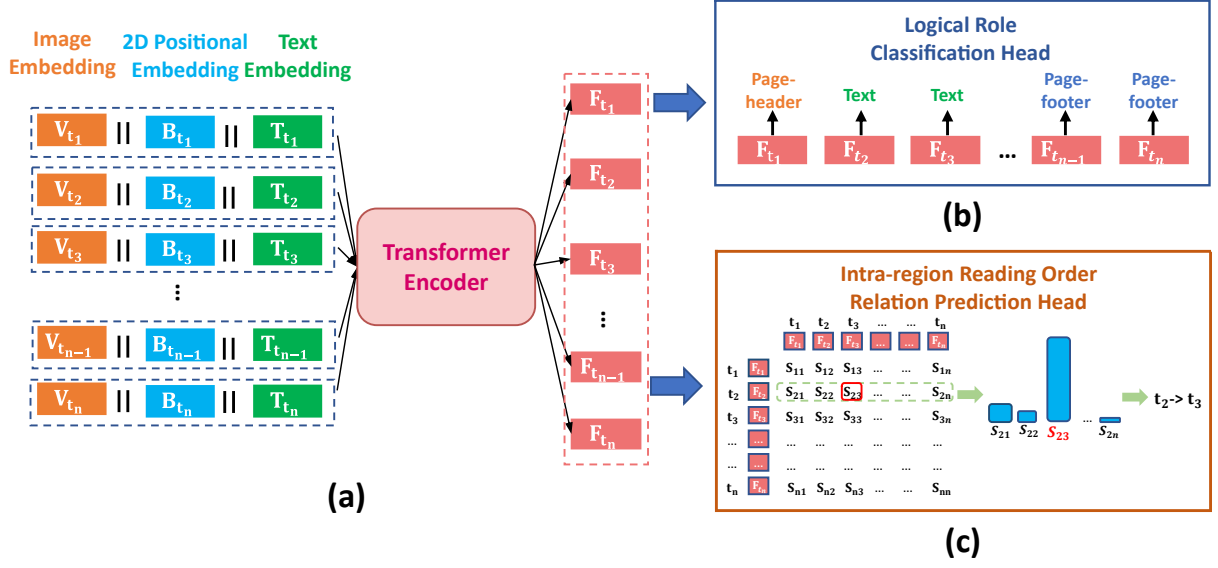


Figure 5: Illustration of (a) Multi-modal Feature Enhancement Module; (b) Logical Role Classification Head; (c) Reading Order Relation Prediction Head in bottom-up text region detection model.

sub-word token sequence, which is then fed into BERT to get the embedding of each token. After that, we average the embeddings of all the tokens in each text-line  $t_i$  to obtain its text embedding  $T_{t_i}$ , followed by a fully-connected layer with 1,024 nodes to make the dimension the same as that of  $V_{t_i}$ :

$$T_{t_i} = LN(ReLU(FC(T_{t_i}))). \quad (2)$$

**2D Positional Embedding.** For each text-line  $t_i$ , we encode its bounding box and size information as its 2D Positional Embedding  $B_{t_i}$ :

$$B_{t_i} = LN(MLP(x_i^1/W, y_i^1/H, x_i^2/W, y_i^2/H, w_i/W, h_i/H)), \quad (3)$$

where  $(w_i, h_i)$  and  $(W, H)$  represent the width and height of  $b_{t_i}$  and the input image, respectively. MLP consists of 2 fully-connected layers with 1,024 nodes, each of which is followed by ReLU.

For each text-line  $t_i$ , we concatenate its visual embedding  $V_{t_i}$ , text embeddings  $T_{t_i}$ , and 2D Positional Embedding  $B_{t_i}$  to obtain its multi-modal representation  $U_{t_i}$ .

$$U_{t_i} = FC(Concat(V_{t_i}, T_{t_i}, B_{t_i})), \quad (4)$$

where FC is a fully-connected layer with 1,024 nodes.

#### 4.2.2. Multi-modal Feature Enhancement Module

As shown in Fig. 5, we use a lightweight Transformer encoder to further enhance the multi-modal representations of text-lines by modeling their interactions with a self-attention mechanism. Each text-line

is treated as a token of the Transformer encoder and its multi-modal representation is taken as the input embedding:

$$F_t = \text{TransformerEncoder}(U_t), \quad (5)$$

where  $U_t = [U_{t_1}, U_{t_2}, \dots, U_{t_n}]$  and  $F_t = [F_{t_1}, F_{t_2}, \dots, F_{t_n}]$  are the input and output embeddings of the Transformer encoder,  $n$  is the number of the input text-lines. To save computation, here we only use a 1-layer Transformer encoder, where the head number, dimension of hidden state, and the dimension of feedforward network are set as 12, 768, and 2048, respectively.

#### 4.2.3. Intra-region Reading Order Relation Prediction Head

We propose to use a relation prediction head to predict intra-region reading order relationships between text-lines. Given a text-line  $t_i$ , if a text-line  $t_j$  is its succeeding text-line in the same text region, we define that there exists an intra-region reading order relationship  $(t_i \rightarrow t_j)$  pointing from text-line  $t_i$  to text-line  $t_j$ . If text-line  $t_i$  is the last (or only) text-line in a text region, its succeeding text-line is considered to be itself. Unlike many previous methods that consider relation prediction as a binary classification task [42, 45], we treat relation prediction as a dependency parsing task and use a softmax cross-entropy loss to replace the standard binary cross-entropy loss during optimization by following [72]. Moreover, we adopt a spatial compatibility feature introduced in [73] to effectively model spatial interactions between text-lines for relation prediction.

Specifically, we use a multi-class (i.e.,  $n$ -class) classifier to calculate a score  $s_{ij}$  to estimate how likely  $t_j$  is the succeeding text-line of  $t_i$  as follows:

$$f_{ij} = FC_q(F_{t_i}) \circ FC_k(F_{t_j}) + MLP(r_{b_{t_i}, b_{t_j}}), \quad (6)$$

$$s_{ij} = \frac{\exp(f_{ij})}{\sum_N \exp(f_{ij})}, \quad (7)$$

where each of  $FC_q$  and  $FC_k$  is a single fully-connected layer with 2,048 nodes to map  $F_{t_i}$  and  $F_{t_j}$  into different feature spaces;  $\circ$  denotes dot product operation; MLP consists of 2 fully-connected layers with 1,024 nodes and 1 node respectively;  $r_{b_{t_i}, b_{t_j}}$  is a spatial compatibility feature vector between  $b_{t_i}$  and  $b_{t_j}$ , which is a concatenation of three 6-d vectors:

$$r_{b_{t_i}, b_{t_j}} = (\Delta(b_{t_i}, b_{t_j}), \Delta(b_{t_i}, b_{t_{ij}}), \Delta(b_{t_j}, b_{t_{ij}})), \quad (8)$$

where  $b_{t_{ij}}$  is the union bounding box of  $b_{t_i}$  and  $b_{t_j}$ ;  $\Delta(\cdot, \cdot)$  represents the box delta between any two bounding boxes. Taking  $\Delta(b_{t_i}, b_{t_j})$  as an example,  $\Delta(b_{t_i}, b_{t_j}) = (d_{ij}^{x_{ctr}}, d_{ij}^{y_{ctr}}, d_{ij}^w, d_{ij}^h, d_{ji}^{x_{ctr}}, d_{ji}^{y_{ctr}})$ , where each dimension is given by:

$$\begin{aligned} d_{ij}^{x_{ctr}} &= (x_i^{ctr} - x_j^{ctr})/w_i, & d_{ij}^{y_{ctr}} &= (y_i^{ctr} - y_j^{ctr})/h_i, \\ d_{ij}^w &= \log(w_i/w_j), & d_{ij}^h &= \log(h_i/h_j), \\ d_{ji}^{x_{ctr}} &= (x_j^{ctr} - x_i^{ctr})/w_j, & d_{ji}^{y_{ctr}} &= (y_j^{ctr} - y_i^{ctr})/h_j, \end{aligned} \quad (9)$$

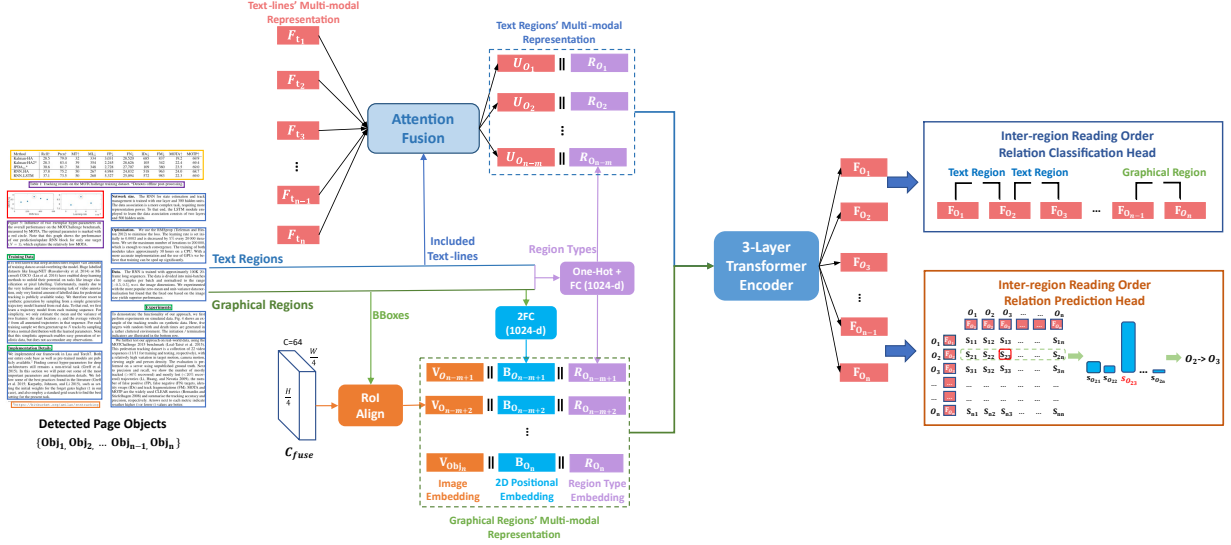


Figure 6: Architecture of our proposed Order module for reading order prediction.

where  $(x_i^{\text{ctr}}, y_i^{\text{ctr}})$  and  $(x_j^{\text{ctr}}, y_j^{\text{ctr}})$  are the center coordinates of  $b_{t_i}$  and  $b_{t_j}$ , respectively.

We select the highest score from scores  $[s_{ij}, j = 1, 2, \dots, n]$  and output the corresponding text-line as the succeeding text-line of  $t_i$ . To achieve higher relation prediction accuracy for the intra-region reading order relationship, which has a chain structure, we employ an additional relation prediction head to further identify the preceding text-line for each text-line. The prediction results from both relation prediction heads are then combined to obtain the final results. Based on the predicted intra-region reading order relationships, we group text-lines into text regions using a Union-Find algorithm. The bounding box of the text region is the union bounding box of all its constituent text-lines.

#### 4.2.4. Logical Role Classification Head

Given the enhanced multi-modal representations of text-lines  $F_t = [F_{t_1}, F_{t_2}, \dots, F_{t_n}]$ , we add a multi-class classifier to predict a logical role label for each text-line and determine the logical role of each text region by the plurality voting of all its constituent text-lines.

#### 4.3. Order Module

The Order module focuses on determining the reading sequence of graphical page objects and text regions identified by the Detect module within document  $D$ . Similar to the bottom-up text region detection model employed in the Detect module, we also utilize our proposed multi-modal, transformer-based relation prediction model to predict the inter-region reading order relationships among the recognized page objects. The Order module processes the detected page objects as input and employs an attention-based approach to integrate the features of text-lines belonging to the same text region, thereby achieving a more efficient

feature representation of the text region. Furthermore, we define two categories of inter-region reading order relationships: (1) Text region reading order relationships between main body text regions, (2) Graphical region reading order relationships between captions/footnotes and graphical page objects such as tables and figures. Consequently, we incorporate an additional inter-region reading order relation classification head to predict relation types. A detailed illustration of the Order module can be found in Fig. 6.

#### 4.3.1. Multi-modal Feature Extraction Module

Following Eqs. (1) and (3) as described in Section 4.2.1, we fuse the visual embedding and the 2D positional embedding to obtain a multi-modal representation  $U_{O_m}$  for each graphical page object  $O_m$  in a similar manner. For each detected text region page object  $O_n$  consisting of text-lines  $[t_{n_1}, t_{n_2}, \dots, t_{n_k}]$ , we propose an attention fusion model to integrate the features of text-lines  $[F_{t_{n_1}}, F_{t_{n_2}}, \dots, F_{t_{n_k}}]$  produced by Eq. (5), thereby forming a multi-modal representation  $U_{O_n}$  for this text region as follows:

$$\alpha_{t_{n_j}} = FC_1(\tanh(FC_2(F_{t_{n_j}}))), \quad (10)$$

$$w_{t_{n_j}} = \frac{\exp \alpha_{t_{n_j}}}{\sum_j \exp \alpha_{t_{n_j}}}, \quad (11)$$

$$U_{O_n} = \sum_j w_{t_{n_j}} F_{t_{n_j}}, \quad (12)$$

where both  $FC_1$  and  $FC_2$  are single fully-connected layers with 1,024 and 1 nodes, respectively. Furthermore, for each page object, we derive a region type embedding for each page object as follows:

$$R_{O_i} = LN(ReLU(FC(Embedding(r_{O_i})))), \quad (13)$$

where *Embedding* is an embedding layer with 1,024 hidden dimension and  $r_{O_i}$  is the logical role of the page object  $O_i$ .

Lastly, we concatenate each page object's multi-modal representation  $U_{O_i}$  and region type embedding  $R_{O_i}$  to obtain its final representation  $\hat{U}_{O_i}$  as follows:

$$\hat{U}_{O_i} = FC(Concat(U_{O_i}, R_{O_i})), \quad (14)$$

where  $FC$  is a fully-connected layer with 1,024 nodes.

#### 4.3.2. Multi-modal Feature Enhancement Module

As illustrated in Fig. 6, we adopt a similar approach to previous multi-modal feature enhancement module in the Group stage. In this case, we utilize a three-layer Transformer encoder to further improve the multi-modal representations of page objects by modeling their interactions using a self-attention mechanism. Each page object is treated as a token of the Transformer encoder, and its multi-modal representation serves as the input embedding:

$$F_O = TransformerEncoder(\hat{U}_O), \quad (15)$$

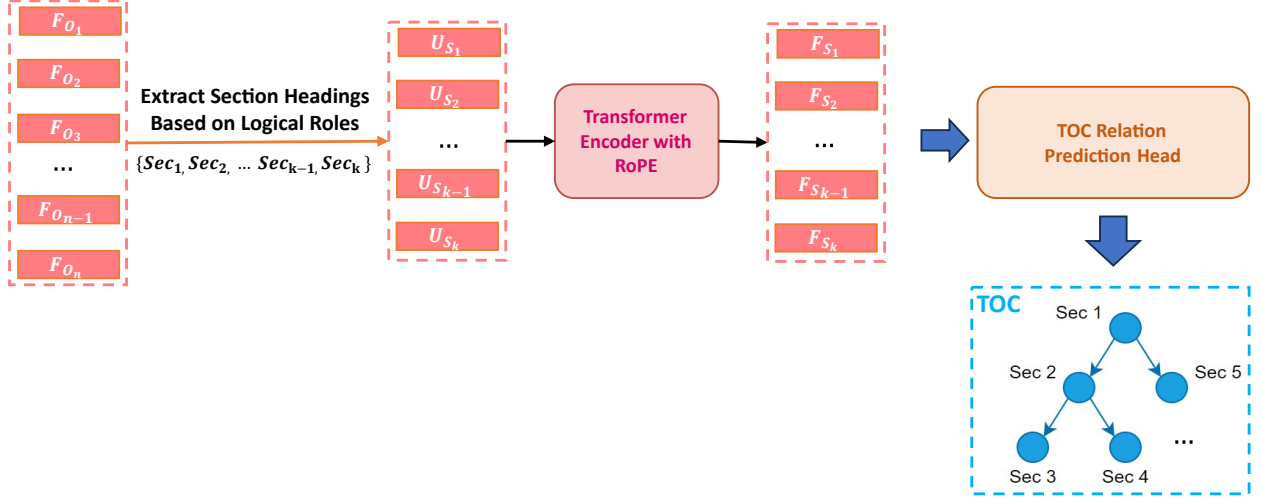


Figure 7: Illustration of the Construct module.

where  $\hat{U}_O = [\hat{U}_{O_1}, \hat{U}_{O_2}, \dots, \hat{U}_{O_n}]$  and  $F_O = [F_{O_1}, F_{O_2}, \dots, F_{O_n}]$  represent the input and output embeddings of the Transformer encoder, and  $n$  is the number of the input page objects. The hyperparameters of the transformer encoder are consistent with those in the Detect module, except for the layer number.

#### 4.3.3. Inter-region Reading Order Relation Prediction Head

Owing to the similarity between the inter-region reading order task of the Order module and the intra-region reading order task of the Detect module, we employ an identical structure for the inter-region reading order relation prediction head in both modules. Further details about this head can be found in Section 4.2.3.

#### 4.3.4. Inter-region Reading Order Relation Classification Head

We employ a multi-class classifier to compute the probability distribution across various classes in order to determine the relation type between page object  $O_i$  and page object  $O_j$ . It works as follows:

$$p_{ij} = BiLinear(FC_q(F_{O_i}), FC_k(F_{O_j})), \quad (16)$$

$$c_{ij} = \argmax(p_{ij}), \quad (17)$$

where both  $FC_q$  and  $FC_k$  represent single fully-connected layers with 2,048 nodes, which are used to map  $F_{O_i}$  and  $F_{O_j}$  into distinct feature spaces; *BiLinear* signifies the bilinear classifier; and *argmax* refers to identifying the index  $c_{ij}$  of the maximum value within the given probability distribution  $p_{ij}$  as the predicted relation type.

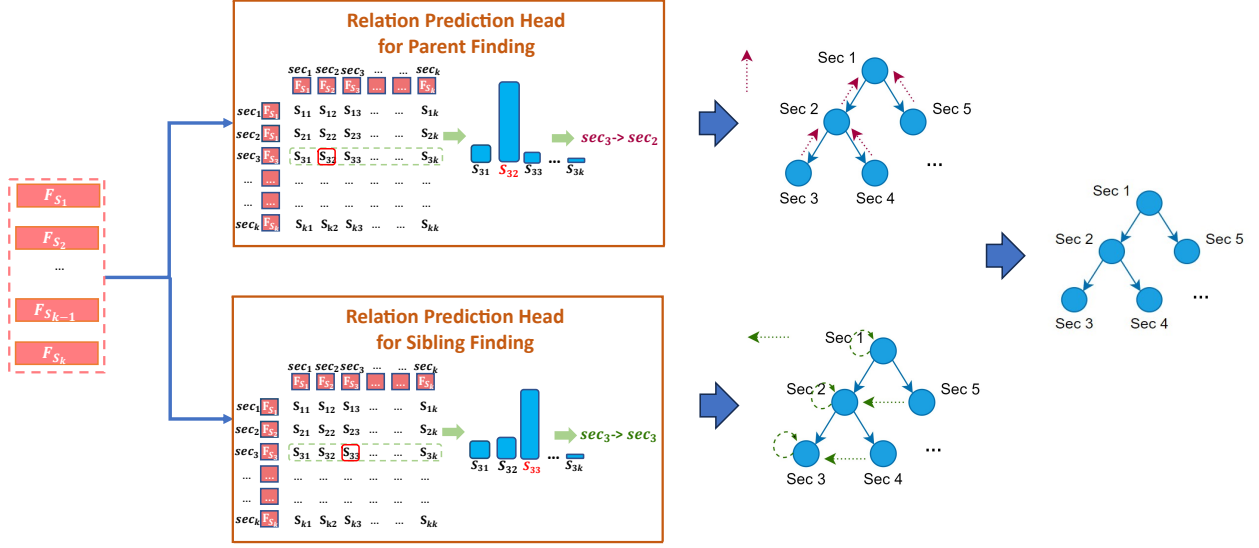


Figure 8: Illustration of TOC Relation Prediction Head.

#### 4.4. Construct Module

Given the detected section headings  $[sec_1, sec_2, \dots, sec_{k-1}, sec_k]$  arranged according to the predicted reading order sequence for document  $D$ , the goal of the Construct module is to generate a tree structure representing the hierarchical table of contents. As illustrated in Fig. 7, we extract the multi-modal representation  $F_{S_i}$  of each section heading  $sec_i$  from all page objects' multi-modal representation  $F_O$  based on the logical role. Subsequently, we input all section headings' representation  $U_S = [U_{S_1}, U_{S_2}, \dots, U_{S_k}]$  into a transformer encoder to further enhance the representations. However, unlike the transformer encoder employed in the Detect module and the Order module, both of which are order-agnostic, the input sequence  $U_S$  has the correct reading order predicted by the Order module, allowing us to add a positional encoding to convey the reading order information. To incorporate the relative position in the reading order sequence and accommodate a larger scale of page numbers in the document, we utilize an efficient positional encoding method called Rotary Positional Embedding (RoPE) [74]. RoPE encodes the absolute position using a rotation matrix and simultaneously includes the explicit relative position dependency in the self-attention formulation. Following the Multi-modal Feature Enhancement Module, we generate the enhanced representations  $F_S = [F_{S_1}, F_{S_2}, \dots, F_{S_k}]$  for section headings. Finally, we introduce a tree-aware TOC relation prediction head to predict the TOC relationships among these section headings. The specially designed relation prediction head is illustrated in Fig. 8.

##### 4.4.1. TOC Relation Prediction Head

During the generation of the ordered tree for Table of Contents, solely relying on the relationship features between child and parent nodes has proven to be insufficient. Some prior studies [7, 62, 64] have already

observed that incorporating information from sibling nodes can lead to an improved generation of the TOC. Inspired by these works, we propose two types of TOC relationships between section heads to further enhance the TOC generation process: parent-child relationships and sibling relationships.

The parent-child relationship is relatively straightforward: when a section heading  $sec_i$  serves as the parent node for another section heading  $sec_j$  within the TOC tree structure, we define a parent-child relationship ( $sec_j \rightarrow sec_i$ ) that points from  $sec_j$  to  $sec_i$ . Sibling relationships in a TOC tree are established as follows: if section heading  $sec_i$  acts as the left sibling of section heading  $sec_j$ , then a sibling relationship ( $sec_j \rightarrow sec_i$ ) is present. In cases where a section heading lacks a parent node or left sibling node, its parent-child or sibling relationship is defined as pointing to itself. This approach aims to provide a more comprehensive representation of the relationships among section heads, ultimately leading to a more accurate and robust TOC generation.

As illustrated in Fig. 8, our proposed TOC Relation Prediction Head comprises two distinct relation prediction heads for the parent-child and sibling relationships, respectively. Both relation prediction heads in our proposed module employ the same network structure. To elaborate, we use the relation prediction head for the parent-child relationship as an example. Specifically, we implement a multi-class (k-class) classifier to compute a score  $s_{ij}^p$ , which estimates the likelihood of  $sec_j$  being the parent node of  $sec_i$ . The calculation is as follows:

$$f_{ij} = FC_q(F_{S_i}) \circ FC_k(F_{S_j}), \quad (18)$$

$$s_{ij}^p = \frac{\exp(f_{ij})}{\sum_j \exp(f_{ij})}, \quad (19)$$

where each of  $FC_q$  and  $FC_k$  represents a single fully-connected layer with 2,048 nodes to map  $F_{S_i}$  and  $F_{S_j}$  into distinct feature spaces;  $\circ$  denotes the dot product operation. Similarly, we can obtain the score  $s_{ij}^s$  to estimate the likelihood of  $sec_j$  being the defined sibling node of  $sec_i$ . This unified network structure allows for efficient and effective prediction of relationships between section heads, contributing to the overall TOC generation process.

In a manner similar to the previously proposed reading order relation prediction head in Section 4.2.3, we treat relation prediction as a dependency parsing task and employ a softmax cross-entropy loss instead of the standard binary cross-entropy loss during the training phase. During the testing phase, we utilize serial decoding to integrate the outputs of the two relation prediction heads and introduce a tree structure constraint to enhance the final prediction results. Specifically, assuming that  $[sec_1, sec_2, \dots, sec_k]$  has been sorted according to the predicted reading order, we initialize a tree  $T$  containing only one root node,  $ROOT$ . Subsequently, we devise a tree insertion algorithm, as detailed in Algorithm 1, to insert each section heading in order, ultimately generating a complete table of contents tree. This approach ensures that the predicted relationships between section headings are consistent with the hierarchical tree structure, resulting in a more accurate and coherent TOC.

---

**Algorithm 1** Tree Insertion Algorithm

---

**Require:** Empty Tree  $T = \{ROOT\}$ , Ordered Section Headings  $[sec_1, sec_2, \dots, sec_k]$ ,

Parent Score Matrix  $\mathbf{s}^P$ , Sibling Score Matrix  $\mathbf{s}^S$

- 1: **for**  $i = 1$  **to**  $k$  **do**
  - 2:   Find the rightmost sub-tree of  $T$  and retrieve nodes  $[sec_{r_1}, sec_{r_2}, \dots, sec_{r_n}]$
  - 3:   Compute the parent score vector  $\mathbf{scores}_P = \mathbf{s}^P[sec_i, [sec_{r_1}, sec_{r_2}, \dots, sec_{r_n}]] \in \mathbb{R}^n$
  - 4:   Compute the sibling score vector  $\mathbf{scores}_S = \mathbf{s}^S[sec_i, [sec_{r_2}, \dots, sec_{r_n}, sec_i]] \in \mathbb{R}^n$
  - 5:   Compute the final score vector  $\mathbf{scores} = \mathbf{scores}_P \circ \mathbf{scores}_S \in \mathbb{R}^n$
  - 6:   Find the index  $m$  corresponding to the maximum score in  $\mathbf{scores}$
  - 7:   Insert  $sec_i$  as the right-most child of  $sec_{r_m}$
  - 8:   Update Tree  $T$
  - 9: **end for**
  - 10: **return** TOC Tree  $T$
- 

## 5. Experiments

### 5.1. Datasets and Evaluation Protocols

In our conference paper [10], we conducted experiments on two widely-recognized large-scale document layout analysis benchmarks, namely PubLayNet [8] and DocLayNet [9] to validate the effectiveness of our proposed Detect module. In this paper, we carry out extensive experiments on a high-quality public hierarchical document structure reconstruction benchmark, HRDoc [7], to validate the effectiveness of our proposed tree construction based framework. It is important to note that HRDoc solely provides annotations and benchmarks for the logical role classification task and the overall hierarchical structure reconstruction task. However, each sub-task plays a crucial role in hierarchical document structure analysis. Consequently, during the performance evaluation phase, conducting a thorough and rigorous assessment of each involved sub-task is essential.

To address this issue, we expand upon the foundation of HRDoc and develop a comprehensive benchmark called Comp-HRDoc for hierarchical document structure analysis, which simultaneously evaluates page object detection, reading order prediction, table of contents extraction, and hierarchical document structure reconstruction. It is worth noting that the logical role classification in HRDoc is actually text-line-level, which may not be a fair performance evaluation for top-down approaches. Therefore, we replace it with a more popular and significant subtask, termed page object detection, in our proposed benchmark. To the best of our knowledge, Comp-HRDoc is the first benchmark designed to assess such a diverse array of document structure analysis subtasks. Our proposed model has been rigorously evaluated on this benchmark, further demonstrating the superiority of our approach.

**PubLayNet** [8] is a large-scale dataset for document layout analysis released by IBM that contains 340,391, 11,858, and 11,983 document pages for training, validation, and testing, respectively. All the documents in this dataset are scientific papers publicly available on PubMed Central, and all the ground-truths are automatically generated by matching the XML representations and the content of the corresponding PDF files. It predefines 5 types of page objects, including Text (i.e., Paragraph), Title, List, Figure, and Table. The evaluation metric for PubLayNet is the COCO-style mean average precision (mAP) at multiple intersection over union (IoU) thresholds between 0.50 and 0.95 with a step of 0.05.

**DocLayNet** [9] is a challenging human-annotated document layout analysis dataset newly released by IBM that contains 69,375, 6,489, and 4,999 document pages for training, testing, and validation, respectively. It covers a variety of document categories, including financial reports, patents, manuals, laws, tenders, and scientific papers. It predefines 11 types of page objects, including Caption, Footnote, Formula, List-item, Page-footer, Page-header, Picture, Section-header, Table, Text (i.e., Paragraph), and Title. The evaluation metric for DocLayNet is also the COCO-style mean average precision (mAP), consistent with that of PubLayNet.

**HRDoc** [7] is a human-annotated dataset specifically designed to facilitate hierarchical document structure reconstruction. It features line-level annotations and cross-page relations, aiming to recover the semantic structure of PDF documents. In order to accommodate various layout types, the HRDoc dataset is divided into two parts. The first part, **HRDoc-Simple** (HRDS), consists of 1,000 documents exhibiting similar layouts. The second part, **HRDoc-Hard** (HRDH), encompasses 1,500 documents with diverse layouts. This heterogeneous collection of documents offers researchers an extensive resource to develop and assess algorithms for hierarchical document structure reconstruction in PDF documents.

Two evaluation tasks are associated with HRDoc, including semantic unit classification (i.e., logical role classification) and hierarchical structure reconstruction. For the semantic unit classification task, the F1 score for each logical role serves as the evaluation metric. Meanwhile, the hierarchical structure reconstruction task adopts the Semantic-TEDS [7] as its evaluation metric.

**Comp-HRDoc** is our proposed benchmark, specifically designed for comprehensive hierarchical document structure analysis. It encompasses tasks such as page object detection, reading order prediction, table of contents extraction, and hierarchical structure reconstruction. Comp-HRDoc is built upon the HRDoc-Hard dataset [7], which comprises 1,000 documents for training and 500 documents for testing. We retain all original images without modification and extend the original annotations to accommodate the evaluation of these included tasks.

In the page object detection task, we utilize the COCO-style segmentation-based mean average precision (mAP) evaluation metric rather than a box-based metric. This choice is due to the fact that a paragraph in HRDoc is considered a logical paragraph, which may span multiple columns in a multi-column page. Consequently, the paragraph’s segmentation is derived from the union of segmentations for all its text-lines.

Meanwhile, the segmentation of a graphical page object remains identical to its bounding box. Regarding the reading order prediction task, a document may encompass multiple reading order groups, such as multiple articles in newspapers, without an explicit reading order definition between them. This characteristic aligns with realistic user requirements. However, it renders some previous full ranking metrics [59] unsuitable for this situation. Furthermore, earlier reading order evaluation metrics primarily focused on the reading order of text units (e.g., text-lines) and neglected paragraph segmentation errors stemming from the Detect stage. As a result, they did not provide a comprehensive assessment of the reading order. In this paper, we propose a reading edit distance score (**REDS**) to evaluate the reading order task. Specifically, we primarily categorize reading order groups into two types: Text Region Reading Order Group and Graphical Region Reading Order Group, evaluating these two types of reading order groups independently. We define the basic evaluation units as text-lines and graphical page objects. To evaluate the paragraph segmentation error, we introduce a special tag  $i/p_i$  at paragraph ending positions within the reading order groups to serve as a marker for paragraph segmentation. We adopt the Levenshtein distance [75], which measures the minimum number of node operations (insertions, deletions, or substitutions) required to equalize two lists of nodes, to calculate the distance between two reading order groups. Given the presence of multiple groups, for each predicted reading order group, we compute the distance between it and all reading order groups in the ground truth. Subsequently, we utilize the Hungarian matching [76] to obtain the overall minimum distance  $D$  and multiply it by the normalization factor  $1/N$ , where  $N$  represents the number of basic units. Ultimately, we define  $1 - \frac{D}{N}$  as the evaluation score for the reading order prediction task. For both the table of contents extraction and hierarchical structure reconstruction tasks, we opt for the Semantic-TEDS [7] as their evaluation metric.

## 5.2. Implementation Details

We implement our approach using PyTorch v1.10, and all experiments are conducted on a workstation equipped with 8 Nvidia Tesla V100 GPUs (32 GB memory). It is crucial to mention that in PubLayNet, a list constitutes an entire object containing multiple list items with labels that are inconsistent with those of text or titles. To minimize ambiguity, we treat all lists as specific graphical page objects.

Since only the task of page object detection needs to be evaluated on PubLayNet and DocLayNet datasets, we only trained the Detect stage in our framework on these two datasets. In our experiments with PubLayNet and DocLayNet, we leverage three multi-scale feature maps  $\{C_3, C_4, C_5\}$  from the backbone network, along with the DINO-based graphical page object detection model, to identify graphical objects. In training, the parameters of the CNN backbone network are initialized with a ResNet-50 model [77] pretrained on the ImageNet classification task, while the parameters of the text embedding extractor are initialized with the pretrained BERT<sub>BASE</sub> model [71]. We optimize the models using the AdamW [78] algorithm with a batch size of 16 and trained for 12 epochs on PubLayNet and 24 epochs on DocLayNet.

The learning rate and weight decay are set to  $1e-5$  and  $1e-4$  for the CNN backbone network, and  $2e-5$  and  $1e-2$  for BERT<sub>BASE</sub>, respectively. The learning rate is divided by 10 at the 11<sup>th</sup> epoch for PubLayNet and 20<sup>th</sup> epoch for DocLayNet. Other hyperparameters of AdamW, including betas and epsilon, are set to (0.9, 0.999) and  $1e-8$ , respectively. We also adopt a multi-scale training strategy, randomly rescaling the shorter side of each image to lengths chosen from [512, 640, 768], ensuring the longer side does not exceed 800. During the testing phase, we set the shorter side of the input image to 640.

For HRDoc and Comp-HRDoc, we utilize four multi-scale feature maps  $\{C_2, C_3, C_4, C_5\}$  from the backbone network, in conjunction with the Mask2Former-based graphical page object detection model, to identify graphical objects. Given that hierarchical document structure analysis requires processing dozens of document pages, we choose the ResNet-18 model as the CNN backbone network to reduce GPU memory requirements. The parameters of the text embedding extractor are also initialized with the pretrained BERT<sub>BASE</sub> model. The models are optimized using the AdamW [78] algorithm with a batch size of 1 and trained for 20 epochs on HRDoc and Comp-HRDoc. The initial learning rate and weight decay are set to  $2e-4$  and  $1e-2$  for the CNN backbone network, and  $4e-5$  and  $1e-2$  for BERT<sub>BASE</sub>, respectively. After a warmup period (set to 2 epochs) during which it increases linearly from 0 to the initial learning rate set in the optimizer, the learning rate linearly decreases from the initial learning rate set in the optimizer to 0. For multi-scale training strategy, the shorter side of each image is randomly rescaled to a length chosen from [320, 416, 512, 608, 704, 800], ensuring that the longer side does not exceed 1024. During the testing phase, we set the shorter side of the input image to 512.

### 5.3. Comparisons with Prior Arts

The Detect module we proposed is a novel combination of top-down and bottom-up approaches for page object detection. Therefore, we first validate the effectiveness of our method on two large-scale document layout analysis datasets, i.e., DocLayNet and PubLayNet.

**DocLayNet.** We compare our proposed Detect module with the other most competitive methods, including Mask R-CNN, Faster R-CNN, YOLOv5, and DINO on DocLayNet. As shown in Table 1, our approach substantially outperforms the closest method YOLOv5 by improving mAP from 76.8% to 81.0%. Considering that DocLayNet is an extremely challenging dataset that covers a variety of document scenarios and contains a large number of text regions with fine-grained logical roles, the superior performance achieved by our proposed approach demonstrates the advantage of our approach.

**PubLayNet.** We also compare our approach with several state-of-the-art vision-based and multimodal methods on PubLayNet. The experimental results are presented in Table 2 and Table 3. We can see that our approach outperforms all these methods regardless of whether textual features are used in our bottom-up text region detection model.

To further validate the effectiveness of our proposed tree construction based framework for hierarchical

Table 1: Performance comparisons on the DocLayNet testing set (in %). The results of Mask R-CNN, Faster R-CNN, and YOLOv5 are obtained from [9].

	Human	Mask R-CNN	Faster R-CNN	YOLOv5	DINO	Ours
Caption	84-89	71.5	70.1	77.7	<b>85.5</b>	83.2
Footnote	83-91	71.8	73.7	<b>77.2</b>	69.2	69.7
Formula	83-85	63.4	63.5	<b>66.2</b>	63.8	63.4
List-item	87-88	80.8	81.0	86.2	80.9	<b>88.6</b>
Page-footer	93-94	59.3	58.9	61.1	54.2	<b>90.0</b>
Page-header	85-89	70.0	72.0	67.9	63.7	<b>76.3</b>
Picture	69-71	72.7	72.0	77.1	<b>84.1</b>	81.6
Section-header	83-84	69.3	68.4	74.6	64.3	<b>83.2</b>
Table	77-81	82.9	82.2	<b>86.3</b>	85.7	84.8
Text	84-86	85.8	85.4	<b>88.1</b>	83.3	84.8
Title	60-72	80.4	79.9	82.7	82.8	<b>84.9</b>
mAP	82-83	73.5	73.4	76.8	74.3	<b>81.0</b>

document structure analysis, we performed experiments with our method on both HRDoc and Comp-HRDoc datasets and made thorough comparisons with previous approaches.

**HRDoc.** As demonstrated in Table 4 and Table 5, we conducted separate performance evaluations for the two tasks in HRDoc, specifically semantic unit classification and hierarchical structure reconstruction. For semantic unit classification, it is evident that our proposed method achieves superior performance in the majority of categories, particularly in the *Fstl* (Firstline) and *Footn* (Footnote) classes, where our approach significantly surpasses previous methods. Although the DSPS Encoder is also a multimodal technique that integrates visual and linguistic information, its performance in the *Mail* category is notably inferior to that of Sentence-BERT. However, on HRDoc-Hard, our method attains an F1 score nearly 5% higher than the DSPS Encoder in this category. Regarding hierarchical structure reconstruction, our proposed tree construction based method markedly outperforms the DSPS Encoder. On HRDoc-Hard, we exceed its performance by 16.63% and 15.77% in Micro-STEDS and Macro-STEDS, respectively. Similarly, on HRDoc-Simple, we surpass the DSPS Encoder by 13.61% and 13.36% in Micro-STEDS and Macro-STEDS, respectively. It is important to highlight that our proposed method evaluates the performance based on the predicted reading order sequence, whereas the DSPS Encoder directly takes advantage of the ground-truth reading order.

**Comp-HRDoc.** As presented in Table 6, we conduct a comprehensive performance evaluation for all tasks in Comp-HRDoc, encompassing page object detection, reading order prediction, table of contents extraction, and hierarchical document reconstruction. We select previous state-of-the-art methods specifically designed for each task to be evaluated using our benchmark. Our proposed method is capable of handling

Table 2: Performance comparisons on the PubLayNet validation set (in %). Vision and Text stand for using visual and textual features, respectively.

Method	Modality	Text	Title	List	Table	Figure	mAP
Faster R-CNN [8]	Vision	91.0	82.6	88.3	95.4	93.7	90.2
Mask R-CNN [8]		91.6	84.0	88.6	96.0	94.9	91.0
Naik et al. [79]		94.3	88.7	94.3	97.6	96.1	94.2
Minouei et al. [80]		94.4	90.8	94.0	97.4	96.6	94.6
DiT-L [28]		94.4	89.3	96.0	97.8	<b>97.2</b>	94.9
SRRV [81]		95.8	90.1	95.0	97.6	96.7	95.0
DINO [69]		94.9	91.4	96.0	98.0	97.3	95.5
TRDLU [31]		95.8	92.1	97.6	97.6	96.6	96.0
UDoc [35]	Vision+Text	93.9	88.5	93.7	97.3	96.4	93.9
LayoutLMv3 [36]		94.5	90.6	95.5	97.9	97.0	95.1
VSR [32]		96.7	93.1	94.7	97.4	96.4	95.7
Ours	Vision	97.0	92.8	96.4	98.1	<b>97.4</b>	96.3
Ours	Vision+Text	<b>97.4</b>	<b>93.5</b>	<b>96.4</b>	<b>98.2</b>	97.2	<b>96.5</b>

all tasks concurrently and achieves significantly superior results in each of them. Specifically, for page object detection, our method surpasses Mask2former [82] by 14.52% in terms of segmentation-based mAP. Regarding reading order prediction, as previous methods rarely consider multiple reading order groups, we have enhanced the partial order-based algorithm proposed by Lorenzo et al. [59] to decode both categories of reading order groups simultaneously. We observe that in the more challenging category (i.e., text region reading order group), our method outperforms their approach by 15.78% in terms of previously defined REDS. For table of contents extraction, our method exceeds the Multimodal Tree Decoder (MTD) [64] by 18.50% and 16.87% in Micro-STEDS and Macro-STEDS, respectively. In hierarchical structure reconstruction, the evaluation of the DSPS Encoder [7] depends on the provided reading order ground-truth and bounding box ground-truth for graphical objects, while our method’s result is entirely independent of ground truth and is obtained through a comprehensive end-to-end evaluation. Under these conditions, our method still surpasses the DSPS Encoder by 14.68% and 13.94% in Micro-STEDS and Macro-STEDS, respectively. Because the Comp-HRDoc benchmark supports a holistic end-to-end evaluation process for hierarchical document structure analysis, it offers a better evaluation benchmark for universal layout analysis.

#### 5.4. Ablation Studies

We conducted a series of ablation experiments based on Comp-HRDoc to verify the impact of using different modules and modalities.

**Effectiveness of the hybrid strategy and multimodality in the Detect module.** In this section,

Table 3: Performance comparisons on the PubLayNet test set (in %). Vision and Text stand for using visual and textual features, respectively.

Method	Modality	Text	Title	List	Table	Figure	mAP
Faster R-CNN [8]	Vision	91.3	81.2	88.5	94.3	94.5	90.0
Mask R-CNN [8]		91.7	82.8	88.7	94.7	95.5	90.7
DocInsightAI [32]		94.5	88.3	94.8	95.8	97.5	94.2
SCUT [32]		94.3	89.7	94.3	96.6	97.7	94.5
SRK [32]		94.7	90.0	95.1	97.2	<b>98.0</b>	95.0
SiliconMinds [32]		96.2	89.8	94.6	97.0	97.6	95.0
VSR [32]	Vision+Text	<b>96.7</b>	92.3	94.6	97.0	97.9	95.7
Ours	Vision	95.0	96.4	95.2	97.0	97.8	96.3
Ours	Vision+Text	95.0	<b>96.6</b>	<b>95.7</b>	<b>97.3</b>	97.7	<b>96.5</b>

we first evaluate the effectiveness of the proposed hybrid strategy in the Detect module. To this end, we train two baseline models: 1) a Mask2Former baseline to detect both graphical page objects and text regions and 2) a hybrid model (denoted as Hybrid (V)) that leverages Mask2Former for graphical object detection and only uses visual and 2D position features for bottom-up text region detection. As shown in the first two rows of Table 7, compared with the Mask2Former-R50 model, the Hybrid-R18 (V) model can achieve comparable graphical page object detection results but much higher text region detection accuracy on Comp-HRDoc, leading to a 9.86% improvement in terms of segmentation-based mAP. In particular, the Hybrid-R18 (V) model can significantly improve small-scale text region detection performance, e.g., 84.67% vs. 68.97% for Page-footnote, 95.08% vs. 59.01% for Page-header and 95.93% vs. 62.68% for Page-footer. These experimental results clearly demonstrate the effectiveness of the proposed hybrid strategy that combines the best of both top-down and bottom-up methods. In addition, we also conducted an ablation experiment to explore the effectiveness of text modalities in the Detect module, as depicted in the last two rows of Table 7. We find that the hybrid model with text modality (denoted as Hybrid (V+T)) achieves much better performance in semantically sensitive categories, such as Author, Mail, and Affiliate, leading to a 4.66% improvement in terms of segmentation-based mAP. Notably, we have observed many cases of inconsistent paragraph annotations in HRDoc, which might be one of the reasons for the relatively lower performance in the *Para* (Paragraph) category. More ablation studies in the Detect module can be found in our conference paper [10].

**Effectiveness of multimodality in the Construct module.** In this study, we conducted an ablation experiment to investigate the effects of different modalities, specifically text and image modalities. To study the impact of section numbers on the task of table of contents extraction, we removed the section numbers from the text content of section headings and examined the resulting influence on the extraction

Table 4: Comparison results of different baseline models in the semantic unit classification task on HRDoc (in %). F1 means F1-score. The results of Cascade-RCNN, ResNet+RoIAlign, Sentence-Bert and DSPS Encoder are all obtained from [7].

Method	HRDoc-Hard F1 (%)															
	Title	Author	Mail	Affli	Sect	Fstl	Paral	Table	Fig	Cap	Equ	Foot	Head	Footn	Avg. F1 (%)	
															Micro	Macro
Cascade-RCNN	81.50	49.77	33.39	49.34	75.92	64.96	77.86	69.96	72.22	43.72	68.84	70.91	71.00	52.67	73.37	64.94
ResNet+RoIAlign	82.40	48.40	18.43	61.33	33.66	45.37	87.99	21.89	70.28	61.54	48.32	73.69	75.71	6.79	79.25	52.56
Sentence-Bert	95.85	89.92	<b>91.68</b>	<b>91.75</b>	94.26	88.68	96.77	76.96	91.67	91.99	93.94	94.68	92.65	62.61	94.68	89.53
DSPS Encoder	<b>97.71</b>	93.93	85.49	90.95	96.06	91.24	97.96	<b>100.0</b>	<b>100.0</b>	<b>97.32</b>	<b>97.92</b>	98.54	<b>97.83</b>	88.84	96.74	95.27
Ours	97.26	<b>94.22</b>	90.33	90.73	<b>96.25</b>	<b>94.09</b>	<b>98.55</b>	<b>100.0</b>	<b>100.0</b>	96.41	97.68	<b>98.57</b>	97.79	<b>90.75</b>	<b>97.59</b>	<b>95.90</b>

Method	HRDoc-Simple F1 (%)															
	Title	Author	Mail	Affli	Sect	Fstl	Paral	Table	Fig	Cap	Equ	Foot	Head	Footn	Avg. F1 (%)	
															Micro	Macro
Cascade-RCNN	78.83	72.74	64.54	70.13	91.35	87.53	89.7	89.3	73.87	64.87	83.87	87.5	-	79.32	88.30	80.85
ResNet+RoIAlign	93.67	82.53	81.33	84.39	37.09	38.39	91.86	58.44	48.53	70.75	26.89	98.33	-	49.76	85.61	66.30
Sentence-Bert	98.98	96.47	<b>98.95</b>	97.42	97.3	93.27	98.72	94.42	95.72	93.36	96.02	99.89	-	87.11	97.74	95.97
DSPS Encoder	99.43	98.83	96.45	97.33	<b>99.6</b>	98.22	<b>99.74</b>	<b>100.0</b>	99.95	<b>99.06</b>	<b>97.91</b>	<b>100.0</b>	-	99.15	99.52	98.90
Ours	<b>99.67</b>	<b>98.98</b>	98.78	<b>98.95</b>	99.39	<b>98.51</b>	<b>99.74</b>	<b>100.0</b>	<b>100.0</b>	98.03	97.07	<b>100.0</b>	-	<b>99.57</b>	<b>99.54</b>	<b>99.13</b>

Table 5: Comparison results of different models in the hierarchical document reconstruction task on HRDoc.

Method	Level	HRDoc-Simple		HRDoc-Hard	
		Micro-STEDS	Macro-STEDS	Micro-STEDS	Macro-STEDS
DocParser	Page	0.2361	0.2506	0.1873	0.2015
DSPS Encoder	Document	0.8143	0.8174	0.6903	0.6971
Ours	Document	<b>0.9504</b>	<b>0.9510</b>	<b>0.8566</b>	<b>0.8548</b>

process, as shown in Table 8. The experiment yields several significant insights. Firstly, the presence or absence of section numbers in the text content considerably affects the performance when considering the text modality. This observation highlights the pronounced relationship between text modality and section numbers in the extraction of tables of contents. Secondly, the findings indicate that the image modality, when used independently, performs admirably, achieving higher scores than relying solely on text modality. This demonstrates the robustness of the image modality. Lastly, the most favorable performance is achieved when both text and image modalities are incorporated into the methodology. This outcome underlines the necessity of employing a multimodal strategy to accomplish the most desirable results in extracting tables of contents.

**Effectiveness of various components in TOC Relation Prediction Head.** In this study, we perform an ablation experiment to comparatively assess the influence of individual components within our

Table 6: Comparison results of different models in tasks including page object detection, reading order prediction, table of contents extraction and hierarchical document reconstruction on Comp-HRDoc. The symbol <sup>†</sup> represents the results of our enhanced replication, whereas <sup>‡</sup> indicates that the evaluation of this result relies on the provided reading order ground-truth and bounding box ground-truth for graphical objects.

Methods	Page Object Detection	Reading Order Prediction		Table of Contents Extraction		Hierarchical Reconstruction	
	Segmentation mAP (%)	Text Region REDS	Graphical Region REDS	Micro-STEDS	Macro-STEDS	Micro-STEDS	Macro-STEDS
Mask2Former [82]	73.54	-	-	-	-	-	-
Lorenzo et al. <sup>†</sup> [59]	-	0.7741	0.8583	-	-	-	-
MTD [64]	-	-	-	0.6755	0.7099	-	-
DSPS Encoder <sup>‡</sup> [7]	-	-	-	0.5754	0.6230	0.6903	0.6971
Ours	<b>88.06</b>	<b>0.9319</b>	<b>0.8637</b>	<b>0.8605</b>	<b>0.8788</b>	<b>0.8371</b>	<b>0.8365</b>

Table 7: Ablation studies of hybrid strategy and multimodality in the Detect module on Comp-HRDoc (in %).

Method	Page Object Detection												
	Title	Author	Mail	Affili	Sect	Para	Table	Fig	Cap	Foot	Head	Footn	Seg mAP
Mask2Former-R18 [82]	84.41	70.92	51.95	62.13	76.51	71.31	79.78	86.23	71.77	56.73	54.89	66.50	69.42
Mask2Former-R50 [82]	85.37	74.67	66.66	69.21	78.18	74.62	80.49	86.55	76.10	62.68	59.01	68.97	73.54
Hybrid-R18 (V)	94.50	79.40	51.57	70.88	<b>89.65</b>	<b>83.39</b>	80.32	86.35	89.02	95.93	<b>95.08</b>	<b>84.67</b>	83.40
Hybrid-R18 (V + T)	<b>94.60</b>	<b>90.66</b>	<b>84.54</b>	<b>84.91</b>	89.48	81.72	<b>80.58</b>	<b>86.71</b>	<b>89.40</b>	<b>96.44</b>	94.96	82.91	<b>88.06</b>

proposed TOC relation prediction head. The experimental results are presented in Table 9. Firstly, when the Relation Prediction Head for Sibling Finding presented in Fig. 8 is removed, the performance slightly diminishes to 0.8545 and 0.8712 for Micro-STEDS and Macro-STEDS, respectively. Differently, a substantial performance decline is observed when the Tree Insert Algorithm (i.e., Algorithm 1) is omitted, with Micro-STEDS and Macro-STEDS scores of 0.7111 and 0.7652, respectively. Lastly, replacing the softmax cross-entropy loss with the standard binary cross-entropy loss also leads to a decrease in performance, with Micro-STEDS and Macro-STEDS scores of 0.7002 and 0.7475, respectively. These experimental results suggest that each component of our approach positively contributes to the overall performance. Especially, the Tree Insert Algorithm plays a critical role in enhancing the performance of the TOC extraction task.

### 5.5. Limitations of Our Approach

While our proposed end-to-end system demonstrates outstanding performance in a majority of tasks, as corroborated by prior experiments, it is not without limitations. For instance, we presume that the section headers supplied to the Construct module from previous stages are accurately recognized. Consequently, the recognition performance of section headings accounts for part of the Construct module’s bottleneck. Moreover, the information regarding section numbers is vital for harnessing the semantics of section headings

Table 8: Ablation studies of various modalities in the Construct module on Comp-HRDoc.

Modality			Micro-STEDS	Macro-STEDS	
Text		Image			
w/o Section Number	with Section Number				
✓	✓		0.6409	0.6834	
				0.8341	0.8528
		✓		0.8477	0.8685
✓	✓	✓	0.8436	0.8640	
		✓	✓	<b>0.8605</b>	<b>0.8788</b>

Table 9: Ablation studies of various components in TOC Relation Prediction Head on Comp-HRDoc.

Method	Level	Table of Contents Extraction	
		Micro-STEDS	Macro-STEDS
Ours	Document	<b>0.8605</b>	<b>0.8788</b>
- Sibling Finding	Document	0.8545	0.8712
- Tree Insert Algorithm	Document	0.7111	0.7652
- Softmax Cross Entropy Loss	Document	0.7002	0.7475

within our proposed system. Therefore, for documents lacking section numbers, our approach may not exhibit adequate robustness. Several failure examples are depicted in Fig. 9, with red boxes indicating incorrect predictions and green boxes signifying correct predictions. Note that these difficulties are common challenges faced by other state-of-the-art methods. Finding practical solutions to these problems will be the focus of our future work.

## 6. Conclusion and Future Work

In this study, we perform a thorough examination of various aspects of hierarchical document structure analysis (HDSA) and propose a tree construction based approach, named Detect-Order-Construct, to simultaneously address multiple crucial subtasks in HDSA. To showcase the effectiveness of this novel framework, we design an effective end-to-end solution and uniformly define the tasks of these three stages as relation prediction problems. Moreover, to comprehensively assess the performance of different approaches, we introduce a new benchmark, termed Comp-HRDoc, which concurrently evaluates page object detection, reading order prediction, table of contents extraction, and hierarchical structure reconstruction. As a result, our proposed end-to-end system attains state-of-the-art performance on two large-scale document layout analysis datasets (i.e., PubLayNet and DocLayNet), a hierarchical document structure reconstruction dataset (i.e., HRDoc), and our comprehensive benchmark (i.e., Comp-HRDoc).

A CONTEMPORARY OVERVIEW OF PROBABILISTIC LATENT VARIABLE MODELS	0
A Contemporary Overview of Probabilistic Latent Variable Models	1
Level: 1: 0	
Richard P. Feynman	2
Department of Structural Pathology, Massachusetts General Hospital, Harvard Medical School	3
Editor	5
Abstract	6
Level: 1: 0	
In this paper we provide a conceptual overview of latent variable models within a probabilistic modeling framework, an overview that emphasizes the compositional nature and the interconnectiveness of the seemingly disparate models commonly encountered in statistical inference.	7
Keywords: Latent Variable Models, Deep Generative Models, Probabilistic Models	8
Level: 0	
1. Introduction	9
In this paper, we will be concerned with building probabilistic latent variable models for data that populate an $N \times P$ matrix $\mathbf{Y}$ whose every element $y_{n,p} \in \mathbb{R}$ is real-valued measurement.	10
$\mathbf{Y} = \begin{bmatrix} y_{1,1} & y_{1,2} & \dots & y_{1,P} \\ y_{2,1} & y_{2,2} & \dots & y_{2,P} \\ \vdots & \vdots & \ddots & \vdots \\ y_{N,1} & y_{N,2} & \dots & y_{N,P} \end{bmatrix} \quad (1)$	11
In general, we will be limiting our focus mainly to the the multivariate statistics setting in which we treat an observation as a multivariate random vector $\mathbf{y}_n = (y_{n,1}, \dots, y_{n,P})$ consisting of $P$ features and the data as a set of $N$ observations $\mathbf{y} = \{\mathbf{y}_1, \dots, \mathbf{y}_n, \dots, \mathbf{y}_N\}$ . Accordingly, we can think of the data matrix $\mathbf{Y}$ as consisting of $N$ observations, each a $P$ -dimensional vector, that populate the rows of the matrix.	12
Level: 0	
2. Basic Definitions	13
In order to give a rigorous constructive account of probabilistic latent variable models, we begin with defining several foundational concepts that will later serve as our building blocks. In what follows, we attempt to convey the intuition at the expense of mathematical rigor. More rigorous definitions can be found in Durrett (2010) and Athreya and Lahiri (2006).	14
15	
Level: 1	
DETR and its variants	1
DETR [48] is a novel Transformer-based [49] object detection algorithm, which introduced the concept of object query and set prediction loss to object detection. These novel attributes make DETR get rid of many manually designed components in previous CNN-based object detectors like anchor design and non-maximum suppression (NMS). However, DETR has its own issues: 1) Slow training convergence; 2) Unclear physical meaning of object queries; 3) Hard to leverage high-resolution feature maps due to high computational complexity. Deformable DETR [15] proposed several effective techniques to address these issues: 1) Formulating queries as 2D anchor points; 2) Designing a deformable attention module that only attends to certain sampling points around a reference point to efficiently leverage multi-scale feature maps; 3) Proposing a two-stage DETR framework and an iterative bounding box refinement algorithm to further improve accuracy. Inspired by the concept of reference point in Deformable DETR, some follow-up works attempted to address the slow convergence issue by giving spatial priors to the object query. For instance, Conditional DETR [50] divided the cross-attention weights into two parts, i.e., content attention weights and spatial attention weights, and proposed a conditional spatial query to make each cross-attention head in each decoder layer focus on a different part of an object. Anchor DETR [51] generated object queries from 2D anchor points directly. DAB-DETR [52] proposed to use 4D anchor box coordinates to represent queries and dynamically update boxes in each decoder layer. SMCA [53] first predicted a reference 4D box for each query and then directly generated its related spatial cross-attention weights with a Gaussian prior in the transformer decoder. Inspired by two-stage Deformable DETR, Efficient DETR [54] took top-K scored proposals output from the first dense prediction stage and their encoder features as the reference boxes and object queries, respectively. Different from the above works, TSP [55] discarded the whole DETR decoder and proposed an encoder-only DETR. DN-DETR [56] pointed out that the bipartite matching algorithm used in Hungarian loss is another reason for slow convergence and proposed a denoising based training method to speed up DETR convergence. DINO [57] improved the performance and efficiency of DN-DETR further by using a contrastive way for denoising training, a mixed query selection method for anchor initialization, and a look-forward-twice scheme for box prediction.	2
Level: 0	
Methodology	3
Overview	4
As depicted in Fig. 3, the improved TSFormer model contains two key components: 1) A DQ-DETR based split module to predict all row and column separation lines from each input table image; 2) A relation network based cell merging module [14] to recover spanning cells. These two modules are attached to a shared convolutional feature map $P_1$ generated by a ResNet18-FPN backbone [58, 59]. Details of these two components will be described in Section 2.1 and Section 2.2, respectively.	5
Level: 1	
DQ-DETR based split module	6
In the split module, two parallel branches are attached to the shared feature map $P_1$ to predict row and column separators, respectively. Each branch comprises two modules: (1) A spatial CNN based feature enhancement module [14] to generate a content-enhanced feature map; (2) A DQ-DETR decoder to predict the positions of all separation lines. In subsequent sections, we will take the row separation line prediction branch as an example to introduce the details of these two modules.	7
Level: 1	
Spatial CNN based feature enhancement	8
Following RefNetNet [14], a spatial CNN based feature enhancement module is used to enhance the feature representation of each pixel on $P_1$ first. As shown in Fig. 3, we add a $3 \times 3$ convolutional layer and three repeated down-sampling blocks, each composed of a sequence of a $1 \times 2$ max-pooling layer, a $3 \times 3$ convolutional layer, and a ReLU activation function, after $P_1$ sequentially to generate a down-sample	9
10	

(a) Failure case due to incorrect recognition of section headings. (b) Failure case due to the lack of section number headings.

Figure 9: Some typical failure cases of Table of Contents extraction.

In future research, we aim to broaden the scope of our framework to encompass a wider range of real-life scenarios, including contracts, financial reports, and handwritten documents. Additionally, we recognize the importance of addressing documents with graph-based logical structures for more general applications. As such, we plan to explore more robust and effective approaches to handle these complex scenarios. Our ongoing efforts are dedicated to finding a comprehensive and universal document structure analysis solution.

## References

- [1] J. Kreich, A. Luhn, G. Maderlechner, An experimental environment for model based document analysis, in: Proceedings of the International Conference on Document Analysis and Recognition, 1991, pp. 50–58.
- [2] S. Tsujimoto, H. Asada, Understanding multi-articled documents, in: Proceedings of the International Conference on Pattern Recognition, 1990, pp. 551–556.
- [3] A. Yamashita, A model based layout understanding method for the document recognition system, in: Proceedings of the International Conference on Document Analysis and Recognition, 1991, pp. 130–140.
- [4] M. Krishnamoorthy, G. Nagy, S. Seth, M. Viswanathan, Syntactic segmentation and labeling of digitized pages from technical journals, IEEE Transactions on Pattern Analysis and Machine Intelligence 15 (7) (1993) 737–747.
- [5] J. Rausch, O. Martinez, F. Bissig, C. Zhang, S. Feuerriegel, Docparser: Hierarchical document structure parsing from renderings, in: Proceedings of the AAAI Conference on Artificial Intelligence, 2021, pp. 4328–4338.

- [6] K. He, G. Gkioxari, P. Dollár, R. Girshick, Mask r-cnn, in: *Proceedings of the International Conference on Computer Vision*, 2017, pp. 2961–2969.
- [7] J. Ma, J. Du, P. Hu, Z. Zhang, J. Zhang, H. Zhu, C. Liu, Hrdoc: Dataset and baseline method toward hierarchical reconstruction of document structures, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, 2023, pp. 1870–1877.
- [8] X. Zhong, J. Tang, A. J. Yepes, Publaynet: largest dataset ever for document layout analysis, in: *Proceedings of the International Conference on Document Analysis and Recognition*, 2019, pp. 1015–1022.
- [9] B. Pfitzmann, C. Auer, M. Dolfi, A. S. Nassar, P. W. Staar, Doclaynet: A large human-annotated dataset for document-layout analysis, *arXiv preprint arXiv:2206.01062* (2022).
- [10] Z. Zhong, J. Wang, H. Sun, K. Hu, E. Zhang, L. Sun, Q. Huo, A hybrid approach to document layout analysis for heterogeneous document images, in: *Proceedings of the International Conference on Document Analysis and Recognition*, 2023, pp. 189–206.
- [11] S. Mao, A. Rosenfeld, T. Kanungo, Document structure analysis algorithms: a literature survey, in: *Proceedings of Document Recognition and Retrieval X*, 2003, pp. 197–207.
- [12] Y. Y. Tang, S.-W. Lee, C. Y. Suen, Automatic document processing: a survey, *Pattern recognition* 29 (12) (1996) 1931–1952.
- [13] X. Yang, E. Yumer, P. Asente, M. Kraley, D. Kifer, C. Lee Giles, Learning to extract semantic structure from documents using multimodal fully convolutional neural networks, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2017, pp. 5315–5324.
- [14] M. Li, Y. Xu, L. Cui, S. Huang, F. Wei, Z. Li, M. Zhou, Docbank: A benchmark dataset for document layout analysis, in: *Proceedings of the International Conference on Computational Linguistics*, 2020, pp. 949–960.
- [15] L. Gao, X. Yi, Z. Jiang, L. Hao, Z. Tang, ICDAR2017 competition on page object detection, in: *Proceedings of the International Conference on Document Analysis and Recognition*, 2017, pp. 1417–1422.
- [16] X. Yi, L. Gao, Y. Liao, X. Zhang, R. Liu, Z. Jiang, Cnn based page object detection in document images, in: *Proceedings of the International Conference on Document Analysis and Recognition*, Vol. 1, 2017, pp. 230–235.
- [17] D. A. B. Oliveira, M. P. Viana, Fast cnn-based document layout analysis, in: *Proceedings of the International Conference on Computer Vision Workshops*, 2017, pp. 1173–1180.
- [18] R. Girshick, J. Donahue, T. Darrell, J. Malik, Rich feature hierarchies for accurate object detection and semantic segmentation, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2014, pp. 580–587.
- [19] R. Girshick, Fast r-cnn, in: *Proceedings of the International Conference on Computer Vision*, 2015, pp. 1440–1448.
- [20] S. Ren, K. He, R. Girshick, J. Sun, Faster r-cnn: Towards real-time object detection with region proposal networks, in: *Proceedings of the Advances in Neural Information Processing Systems*, 2015, pp. 91–99.
- [21] Z. Cai, N. Vasconcelos, Cascade r-cnn: High quality object detection and instance segmentation, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 43 (5) (2019) 1483–1498.
- [22] X. Wang, R. Zhang, T. Kong, L. Li, C. Shen, Solov2: Dynamic and fast instance segmentation, in: *Proceedings of the Advances in Neural Information Processing Systems*, Vol. 33, 2020, pp. 17721–17732.
- [23] Z. Tian, C. Shen, H. Chen, Conditional convolutions for instance segmentation, in: *Proceedings of the European Conference on Computer Vision*, 2020, pp. 282–298.
- [24] G. Jocher, et al., ultralytics/yolov5: v5.0 - YOLOv5-P6 1280 models, AWS, Supervise.ly and YouTube integrations (Apr. 2021).
- [25] X. Zhu, W. Su, L. Lu, B. Li, X. Wang, J. Dai, Deformable detr: Deformable transformers for end-to-end object detection, in: *Proceedings of the International Conference on Learning Representations*, 2021.
- [26] N. D. Vo, K. Nguyen, T. V. Nguyen, K. Nguyen, Ensemble of deep object detectors for page object detection, in:

- Proceedings of the International Conference on Ubiquitous Information Management and Communication, 2018, pp. 1–6.
- [27] R. Saha, A. Mondal, C. Jawahar, Graphical object detection in document images, in: Proceedings of the International Conference on Document Analysis and Recognition, 2019, pp. 51–58.
  - [28] J. Li, Y. Xu, T. Lv, L. Cui, C. Zhang, F. Wei, Dit: Self-supervised pre-training for document image transformer, in: Proceedings of the ACM International Conference on Multimedia, 2022, pp. 3530–3539.
  - [29] S. Biswas, A. Banerjee, J. Lladós, U. Pal, Docsegtr: an instance-level end-to-end document image segmentation transformer, arXiv preprint arXiv:2201.11438 (2022).
  - [30] K. Hu, Z. Zhong, L. Sun, Q. Huo, Mathematical formula detection in document images: A new dataset and a new approach, Pattern Recognition 148 (2024) 110212.
  - [31] H. Yang, W. Hsu, Transformer-based approach for document layout understanding, in: Proceedings of the International Conference on Image Processing, 2022, pp. 4043–4047.
  - [32] P. Zhang, C. Li, L. Qiao, Z. Cheng, S. Pu, Y. Niu, F. Wu, Vsr: a unified framework for document layout analysis combining vision, semantics and relations, in: Proceedings of the International Conference on Document Analysis and Recognition, 2021, pp. 115–130.
  - [33] C. Shi, C. Xu, H. Bi, Y. Cheng, Y. Li, H. Zhang, Lateral feature enhancement network for page object detection, IEEE Transactions on Instrumentation and Measurement 71 (2022) 1–10.
  - [34] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, B. Guo, Swin transformer: Hierarchical vision transformer using shifted windows, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 10012–10022.
  - [35] J. Gu, J. Kuen, V. I. Morariu, H. Zhao, N. Barmpalios, R. Jain, A. Nenkova, T. Sun, Unified pretraining framework for document understanding, arXiv preprint arXiv:2204.10939 (2022).
  - [36] Y. Huang, T. Lv, L. Cui, Y. Lu, F. Wei, Layoutlmv3: Pre-training for document ai with unified text and image masking, in: Proceedings of the ACM International Conference on Multimedia, 2022, pp. 4083–4091.
  - [37] D. He, S. Cohen, B. Price, D. Kifer, C. L. Giles, Multi-scale multi-task fcn for semantic page segmentation and table detection, in: Proceedings of the International Conference on Document Analysis and Recognition, Vol. 1, 2017, pp. 254–261.
  - [38] Y. Li, Y. Zou, J. Ma, Deeplayout: A semantic segmentation approach to page layout analysis, in: Proceedings of the International Conference on Intelligent Computing Methodologies, 2018, pp. 266–277.
  - [39] X.-H. Li, F. Yin, T. Xue, L. Liu, J.-M. Ogier, C.-L. Liu, Instance aware document image segmentation using label pyramid networks and deep watershed transformation, in: Proceedings of the International Conference on Document Analysis and Recognition, 2019, pp. 514–519.
  - [40] Y. Sang, Y. Zeng, R. Liu, F. Yang, Z. Yao, Y. Pan, Exploiting spatial attention and contextual information for document image segmentation, in: Proceedings of the Advances in Knowledge Discovery and Data Mining, 2022, pp. 261–274.
  - [41] J. Long, E. Shelhamer, T. Darrell, Fully convolutional networks for semantic segmentation, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2015, pp. 3431–3440.
  - [42] X. Li, F. Yin, C. Liu, Page object detection from pdf document images by deep structured prediction and supervised clustering, in: Proceedings of the International Conference on Pattern Recognition, 2018, pp. 3627–3632.
  - [43] X.-H. Li, F. Yin, C.-L. Liu, Page segmentation using convolutional neural network and graphical model, in: Proceedings of the International Workshop on Document Analysis Systems, 2020, pp. 231–245.
  - [44] S. Luo, Y. Ding, S. Long, J. Poon, S. C. Han, Doc-gcn: Heterogeneous graph convolutional networks for document layout analysis, in: Proceedings of the International Conference on Computational Linguistics, 2022, pp. 2906–2916.
  - [45] R. Wang, Y. Fujii, A. C. Popat, Post-ocr paragraph recognition by graph convolutional networks, in: Proceedings of the IEEE Winter Conference on Applications of Computer Vision, 2022, pp. 493–502.

- [46] S. Liu, R. Wang, M. Raptis, Y. Fujii, Unified line and paragraph detection by graph convolutional networks, in: Proceedings of the International Workshop on Document Analysis Systems, 2022, pp. 33–47.
- [47] S. Long, S. Qin, D. Panteleev, A. Bissacco, Y. Fujii, M. Raptis, Towards end-to-end unified scene text detection and layout analysis, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 1049–1059.
- [48] C. Xue, J. Huang, W. Zhang, S. Lu, C. Wang, S. Bai, Contextual text block detection towards scene text understanding, in: Proceedings of the European Conference on Computer Vision, 2022, pp. 374–391.
- [49] R. Wang, Y. Fujii, A. Bissacco, Text reading order in uncontrolled conditions by sparse graph segmentation, in: Proceedings of the International Conference on Document Analysis and Recognition, 2023, pp. 3–21.
- [50] T. M. Breuel, High performance document layout analysis, in: Proceedings of the Symposium on Document Image Understanding Technology, 2003, pp. 209–218.
- [51] M. Aiello, A. M. Smeulders, *Bidimensional relations for reading order detection* (2003).  
URL <https://research.rug.nl/en/publications/bidimensional-relations-for-reading-order-detection>
- [52] S. Ferilli, A. Paziienza, An abstract argumentation-based strategy for reading order detection, in: Proceedings of the AI\*IA Workshop on Intelligent Techniques, Vol. 1509, 2015.
- [53] J. Meunier, Optimized xy-cut for determining a page reading order, in: Proceedings of the International Conference on Document Analysis and Recognition, 2005, pp. 347–351.
- [54] Z. Gu, C. Meng, K. Wang, J. Lan, W. Wang, M. Gu, L. Zhang, Xylayoutlm: Towards layout-aware multimodal networks for visually-rich document understanding, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 4583–4592.
- [55] M. Ceci, M. Berardi, G. Porcelli, D. Malerba, A data mining approach to reading order detection, in: Proceedings of the International Conference on Document Analysis and Recognition, 2007, pp. 924–928.
- [56] D. Malerba, M. Ceci, Learning to order: A relational approach, in: Proceedings of the ECML/PKDD International Workshop on Mining Complex Data, Vol. 4944, 2007, pp. 209–223.
- [57] L. Li, F. Gao, J. Bu, Y. Wang, Z. Yu, Q. Zheng, An end-to-end OCR text re-organization sequence learning for rich-text detail image comprehension, in: Proceedings of the European Conference on Computer Vision, 2020, pp. 85–100.
- [58] Z. Wang, Y. Xu, L. Cui, J. Shang, F. Wei, Layoutreader: Pre-training of text and layout for reading order detection, in: Proceedings of the Conference on Empirical Methods in Natural Language Processing, 2021, pp. 4735–4744.
- [59] L. Quirós, E. Vidal, Reading order detection on handwritten documents, *Neural Computing and Applications* 34 (12) (2022) 9593–9611.
- [60] Z. Wu, P. Mitra, C. L. Giles, Table of contents recognition and extraction for heterogeneous book documents, in: Proceedings of the International Conference on Document Analysis and Recognition, 2013, pp. 1205–1209.
- [61] T. Nguyen, A. Doucet, M. Coustaty, Enhancing table of contents extraction by system aggregation, in: Proceedings of the International Conference on Document Analysis and Recognition, 2017, pp. 242–247.
- [62] R. Cao, Y. Cao, G. Zhou, P. Luo, Extracting variable-depth logical document hierarchy from long documents: Method, evaluation, and application, *Journal of Computer Science and Technology* 37 (3) (2022) 699–718.
- [63] S. Hochreiter, J. Schmidhuber, Long short-term memory, *Neural computation* 9 (8) (1997) 1735–1780.
- [64] P. Hu, Z. Zhang, J. Zhang, J. Du, J. Wu, Multimodal tree decoder for table of contents extraction in document images, in: Proceedings of the International Conference on Pattern Recognition, 2022, pp. 1756–1762.
- [65] G. Nagy, S. C. Seth, Hierarchical representation of optically scanned documents (1984) 347–349.
- [66] A. Conway, Page grammars and page parsing. a syntactic approach to document layout recognition, in: Proceedings of the International Conference on Document Analysis and Recognition, 1993, pp. 761–764.
- [67] Y. Tateisi, N. Itoh, Using stochastic syntactic analysis for extracting a logical structure from a document image, in:

- Proceedings of the IAPR International Conference on Pattern Recognition, 1994, pp. 391–394.
- [68] Z. Wang, M. Zhan, X. Liu, D. Liang, Docstruct: A multimodal method to extract hierarchy structure in document for general form understanding, in: Findings of the Association for Computational Linguistics, Vol. EMNLP 2020, 2020, pp. 898–908.
  - [69] H. Zhang, F. Li, S. Liu, L. Zhang, H. Su, J. Zhu, L. M. Ni, H. Shum, DINO: DETR with improved denoising anchor boxes for end-to-end object detection, in: Proceedings of the International Conference on Learning Representations, 2023.
  - [70] J. L. Ba, J. R. Kiros, G. E. Hinton, Layer normalization, arXiv preprint arXiv:1607.06450 (2016).
  - [71] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, in: Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, 2019, pp. 4171–4186.
  - [72] Y. Zhang, B. Zhang, R. Wang, J. Cao, C. Li, Z. Bao, Entity relation extraction as dependency parsing in visually rich documents, in: Proceedings of the Conference on Empirical Methods in Natural Language Processing, 2021, pp. 2759–2768.
  - [73] J. Zhang, M. Elhoseiny, S. Cohen, W. Chang, A. Elgammal, Relationship proposal networks, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2017, pp. 5678–5686.
  - [74] J. Su, Y. Lu, S. Pan, B. Wen, Y. Liu, Roformer: Enhanced transformer with rotary position embedding, arXiv preprint arXiv:2104.09864 (2021).
  - [75] V. I. Levenshtein, et al., Binary codes capable of correcting deletions, insertions, and reversals, in: Soviet physics doklady, Vol. 10, 1966, pp. 707–710.
  - [76] H. W. Kuhn, The hungarian method for the assignment problem, Naval research logistics quarterly 2 (1-2) (1955) 83–97.
  - [77] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2016, pp. 770–778.
  - [78] I. Loshchilov, F. Hutter, Decoupled weight decay regularization, arXiv preprint arXiv:1711.05101 (2017).
  - [79] S. Naik, K. A. Hashmi, A. Pagani, M. Liwicki, D. Stricker, M. Z. Afzal, Investigating attention mechanism for page object detection in document images, Applied Sciences 12 (15) (2022) 7486.
  - [80] M. Minouei, M. R. Soheili, D. Stricker, Document layout analysis with an enhanced object detector, in: Proceedings of the International Conference on Pattern Recognition and Image Analysis, 2021, pp. 1–5.
  - [81] H. Bi, C. Xu, C. Shi, G. Liu, Y. Li, H. Zhang, J. Qu, Srrv: A novel document object detector based on spatial-related relation and vision, IEEE Transactions on Multimedia 25 (2022) 3788–3798.
  - [82] B. Cheng, I. Misra, A. G. Schwing, A. Kirillov, R. Girdhar, Masked-attention mask transformer for universal image segmentation, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022, pp. 1280–1289.