A Survey of Deep Learning Approaches for OCR and Document Understanding

Nishant Subramani
Scale AI
Masakhane
nishant@nyu.edu

Alexandre Matton
Scale AI

Malcolm Greaves
Scale AI

Adrian Lam
Scale AI

Abstract

Documents are a core part of many businesses in many fields such as law, finance, and technology among others. Automatic understanding of documents such as invoices, contracts, and resumes is lucrative, opening up many new avenues of business. The fields of natural language processing and computer vision have seen tremendous progress through the development of deep learning such that these methods have started to become infused in contemporary document understanding systems. In this survey paper, we review different techniques for document understanding for documents written in English and consolidate methodologies present in literature to act as a jumping-off point for researchers exploring this area.

1 Introduction

Humans compose documents to record and preserve information. As information carrying vehicles, documents are written using different layouts to represent diverse sets of information for a variety of different consumers. In this work, we look at the problem of document understanding for documents written in English. Here, we take the term document understanding to mean the automated process of reading, interpreting, and extracting information from the written text and illustrated figures contained within a document's pages. From the perspective as practitioners of machine learning, this survey covers the methods by which we build models to automatically understand documents that were originally composed for human consumption. Document understanding models take in documents and segment pages of documents into useful parts (i.e. regions corresponding to a specific table or property), often using optical character recognition (OCR) (Mori et al., 1999) with some level of document layout analysis. These models use this information to understand the contents of the document at large, e.g. that this region or bounding box corresponds to an address. In this survey, we focus on these aspects of document understanding at a more granular level and discuss popular methods for these tasks. Our goal is to summarize the approaches present in modern document understanding and highlight current trends and limitations.

In Section 2, we discuss some general themes in modern NLP and document understanding and provide a framework for building end-to-end automated document understanding systems. Next, in Section 3, we look at the best methods for OCR encompassing both text detection (Section 3.1) and text transcription (Section 3.3). We take a broader view of the document understanding problem in section 4, presenting multiple approaches to document layout analysis: the problem of locating relevant information on each page. Following this, we discuss popular approaches for information extraction (Section 5).

2 Document Processing & Understanding

Document processing historically involved handcrafted rule-based algorithms (Lebourgeois et al., 1992; Ha et al., 1995; Amin and Shiu, 2001), but with the widespread success of deep learning (Collobert and Weston, 2008; Krizhevsky et al., 2012; Sutskever et al., 2014), computer vision (CV) and natural language processing (NLP) based methods have come to the fore. Advancements in object detection and image segmentation have led to systems that edge close to human performance on a variety of tasks (Redmon et al., 2016; Lin et al., 2017). As a result, these methods have been applied to a variety of other domains including NLP and speech (Gehring et al., 2017; Wu et al., 2018; Subramani and Rao, 2020). Since documents can be read and viewed as a visual information medium, many practitioners leverage computer vision techniques as well and use them for text detection and instance segmentation (Long et al., 2018; Katti et al., 2018). We cover specific methods to do these in Sections 3.1 and 4.1.

The widespread success and popularity of large pretrained language models such as ELMo and BERT have caused document understanding to shift towards using deep learning based models (Peters et al., 2018; Devlin et al., 2019). These models can be fine-tuned for a variety of tasks and have replaced word vectors as the de-facto standard for pretraining for natural language tasks. However, language models, both recurrent neural network based and transformer based (Vaswani et al., 2017), struggle with long sequences (Cho et al., 2014a; Subramani et al., 2019; Subramani and Suresh, 2020). Given that texts can be very dense and long in business documents, model architecture modifications are necessary. The most simple approach is to truncate documents into smaller sequences of 512 tokens such that pretrained language models can be used off-the-shelf (Xie et al., 2019; Joshi et al., 2019). Another approach that has gained traction recently is based on reducing the complexity of the self-attention component of transformer-based language models (Child et al., 2019; Beltagy et al., 2020; Katharopoulos et al., 2020; Kitaev et al., 2020; Choromanski et al., 2020).

All effective, modern, end-to-end document understanding systems present in the literature integrate multiple deep neural network architectures for both reading and comprehending a document's content. Since documents are made for humans, not machines, practitioners must combine CV as well as NLP architectures into a unified solution. While specific use cases will dictate the exact techniques used, a full end-to-end system employs:

- A computer-vision based document layout analysis module, which partitions each document page into distinct content regions. This model not only delineates between relevant and irrelevant regions, but also serves to categorize the type of content it identifies.

- An optical character recognition (OCR) model, whose purpose is to locate and faithfully transcribe all written text present in the document. Straddling the boundary between CV and NLP, OCR models may either use document layout analysis directly or solve the problem in an independent fashion.

- Information extraction models that use the output of OCR or document layout analysis to comprehend and identify relationships between the information that is being conveyed in the document. Usually specialized to a particular domain and task, these models provide the structure necessary to make a document machine readable, providing utility in document understanding.

In the following sections, we expand upon these concepts that constitute an end-to-end document understanding solution.

3 Optical Character Recognition

OCR has two primary components: text detection and text transcription. Generally, these two components are separate and employ different models for each task. Below, we discuss state-of-the-art methods for each of these components and show how a document can be processed through different generic OCR systems. See Figure 1 for details.
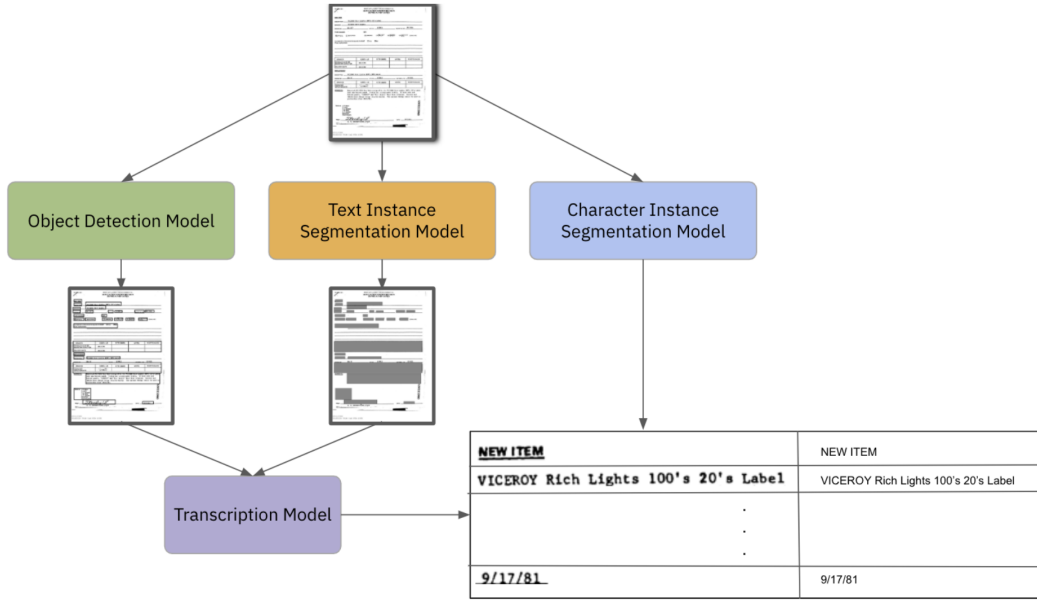
Figure 1: Here, we show the general OCR process. A document can take the left path and go through an object detection model, which outputs bounding boxes, and a transcription model that transcribes the text in each of those bounding boxes. If the document takes the middle path, the object passes through a generic text instance segmentation model that colors pixels black if they contain text and a text transcription model that transcribes the regions of text the instance segmentation model identifies. If the document takes the right path, the model goes through a character-specific instance segmentation model, which outputs which character a pixel corresponds to. All paths produce the same structured output. The document comes from FUNSD (Jaume et al., 2019).

3.1 Text Detection

Text detection is the task of finding text present in a page or image. The input, an image, is often represented by a three dimensional tensor, $C \times H \times W$, where C is the number of channels (often three, for red, green and blue), H is the height, and W is the width of the image. Text detection is a challenging problem because text comes in a variety of shapes and orientations and can often be distorted. We explore two common ways researchers pose the text detection problem: as an object detection task and as a instance segmentation task. A text detection model must either learn to output coordinates of bounding boxes around text (object detection), or a mask, where pixels with text are marked and pixels without are not (instance segmentation).

3.1.1 Text Detection as Object Detection

Traditionally, text detection revolved around hand-crafting features to detect characters (Matas et al., 2004; Lowe, 2004). Advances in deep learning, especially in object detection and semantic segmentation, have led to a change in how text detection is tackled. Using these well-performing object detectors from the traditional computer vision literature, such as the Single-Shot MultiBox Detector (SSD) and Faster R-CNN models (Liu et al., 2016; Ren et al., 2015), practitioners build efficient text detectors.

One of the first papers applying a regression-based detector for text is TextBoxes (Liao et al., 2016, 2018a). They added long default boxes that have large aspect ratios to SSD, in order to adapt the object detector to text. Several papers built on this work to make regression-based models resilient to orientations, like the Deep Matching Prior Network (DMPNet) and the Rotation-Sensitive Regression Detector (RRD) (Liu and Jin, 2017; Liao et al., 2018b). Other papers have a similar approach to the problem, but develop their own proposal network that is tuned towards text rather than towards natural images. For instance, Tian et al. (2016) combine convolutional networks with recurrent networks using a vertical anchor mechanism in their Connectionist Text Proposal Network to improve accuracy for horizontal text.

Object detection models are generally evaluated via an intersection over union (IoU) metric and an F1 score. The metric computes how much of a candidate bounding box overlaps with the ground truth bounding box (the intersection) divided by the total space occupied by both the candidate and ground truth bounding boxes (the union). Next, an IoU threshold τ is chosen to determine which predicted boxes count as true positives (IoU $\geq \tau$). The remainder are classified as false positives. Any box that the model fails to detect is classified as a false negative. Using those definitions, an F1 score is computed to evaluate the object detection model.

3.1.2 Text Detection as Instance Segmentation

Text detection in documents has its own unique set of challenges: notably, the text is usually dense and documents contain a lot more text than what is usually present in natural images. To combat this density problem, text detection can be posed as an ultra-dense instance segmentation task. Instance segmentation is the task of classifying each pixel of an image as specific, pre-defined categories.

Segmentation-based text detectors work at the pixel level to identify regions of text. These per-pixel predictions are often used to estimate probabilities of text regions, characters, and their relationships among adjacent characters in a unified framework. Practitioners use popular segmentation methods like Fully Convolutional Networks (FCN) to detect text (Long et al., 2015), improving upon object detection models, especially when text is misaligned or distorted. Several papers build on this segmentation foundation to output word bounding areas by extracting bounding areas directly from the segmentation output (Yao et al., 2016; He et al., 2017a; Deng et al., 2018). TextSnake extends this further by predicting the text region, center line, direction of text, and candidate radius from an FCN (Long et al., 2018). These features are then combined with a striding algorithm to extract the central axis points to reconstruct the text instance.

3.2 Word-level versus character-level

While most papers cited above try to directly detect words or even lines of words, some papers argue that character-level detection is an easier problem than general text detection because characters are less ambiguous than text lines or words. CRAFT uses an FCN model to output a two-dimensional Gaussian heatmap for each character (Baek et al., 2019). Characters that are close together are then grouped together in a rotated rectangle that has the smallest area possible to still encapsulate the set of characters. More recently, Ye et al. (2020) combine global, word-level, and character-level features obtained using Region Proposal Networks (RPN) to great success.

Most of the models described above were mainly developed for text scene detection, but can be easily adapted to document text detection to handle difficult cases like distorted text. We expect less distortion in documents than in natural images, but poorly scanned documents or documents with certain fonts could still pose these problems.

3.3 Text Transcription

Text transcription is the task of transcribing the text present in an image. The input, an image, is often a crop corresponding to either a character, word, or sequence of words, and has dimension $C \times H' \times W'$. A text transcription model must learn to ingest this cropped image and output a sequence of tokens belonging to some pre-specified vocabulary V. V often corresponds to a set of characters. For digit recognition for instance, this is the most intuitive approach (Goodfellow et al., 2013). Otherwise, V can also correspond to a set of words, similarly to a word-level language modeling problem (Jaderberg et al., 2014). In both cases, the problem can be framed as a multi-class classification problem with the number of classes equal to the size of the vocabulary V.

Word-level text transcription models require more data as the number of classes in the multi-class classification problem is much larger than for character-level. On one hand, predicting words instead of characters decreases the probability of making small typos (like replacing an "a" by an "o" in a word like "elephant"). On the other, limiting oneself to a word-level vocabulary means that it is not possible to transcribe words which are not part of this vocabulary. This problem doesn't exist at the character-level, as the number of characters is limited. As long as we know the language of the document, it is straightforward to build a vocabulary which contains all the possible characters. Subword units are a viable alternative (Sennrich et al., 2016), as they alleviate the issues present in both word and character level transcription.

Recently the research community has moved towards using recurrent neural networks, specifically recurrent models with LSTM or GRU units on top of a convolutional image feature extractor (Hochreiter and Schmidhuber, 1997; Cho et al., 2014b; Wang and Hu, 2017). To transcribe a token, two different decoding mechanisms are often used. One is standard greedy decoding or beam search using an attention-based sequence decoder with cross entropy loss (Bahdanau et al., 2014), exactly like decoding with a conditional language model. Sometimes images are poorly oriented or misaligned, reducing the effectiveness of standard sequence attention. To overcome this, He et al. (2018) uses attention alignment, encoding spatial information of characters directly, while Shi et al. (2016) use spatial attention mechanisms directly. The second way in which transcription decoding is often done is with connectionist temporal classification (CTC) loss (Graves et al., 2006), a common loss function in speech which models repeated characters in sequence outputs well.

The majority of text transcription models borrow from advances in sequence modeling for both text and speech and often can utilize these advancements well with only minor adjustments. As a result, practitioners seldom directly tackle this aspect relative to the other components of the document understanding task.

3.4 End-to-end models

End-to-end approaches combine text detection and text transcription in order to improve both components jointly (Li et al., 2017). For instance, if the text prediction has a very low probability, it means the detected box either did not capture the entire word or captured something that is not text. An end-to-end approach may be very effective in this case. Combining these two methods is fairly common and both Fast Oriented Text Spotting (FOTS) and TextSpotter with Explicit Alignment and Attention sequentially combine these models to train end-to-end (Liu et al., 2018; He et al., 2018). These approaches use shared convolutions as features to both text detection and recognition, and implement methods for complex orientations of text. Feng et al. (2019) introduce TextDragon, an end-to-end model that performs well on distorted text by utilizing a differentiable region of interest slide operator, which specializes in correcting distortions in regions of interest. Mask TextSpotter is another end-to-end model that combines region proposal networks for bounding boxes with text and character segmentation (Liao et al., 2020). These recent works show the power of end-to-end OCR solutions in reducing errors.

Yet, having separate text detection and text recognition models offers more flexibility. First, the two models can be trained separately. In the case where only a small dataset is available to train the whole OCR module, but a lot of text recognition data is easily accessible, it makes sense to leverage this big amount of data in the training of the recognition model. Moreover, with two separate models, it is easy to compute two separate sets of metrics and have a more complete understanding of where the bottleneck might be.

Hence, both two-model and end-to-end approaches are viable. Whether one is better than the other mainly depends on the data available and what one wants to achieve.

3.5 Datasets for Text Detection & Transcription

Most of the literature revolves around scene text detectio, rather than document text detection, and report results on those datasets. Some of the major ones are ICDAR (Karatzas et al., 2013, 2015), Total-Text (Ch'ng and Chan, 2017), CTW1500 (Yuliang et al., 2017), and SynthText (Gupta et al., 2016). Jaume et al. (2019) present FUNSD, a dataset for text detection, transcription, and document understanding with 199 fully annotated forms comprising of 31k word level bounding boxes.

4 Document Layout Analysis

Document layout analysis is the process of locating and categorizing regions of interest on a picture or scanned image of a page. Broadly, most approaches can be distilled into page segmentation and logical structural analysis (Binmakhashen and Mahmoud, 2019; Okun et al., 1999). Page segmentation methods focus on appearance and use visual cues to partition pages into distinct regions; the most common are text, figures, images, and tables. In contrast, logical structural analysis focuses on providing finer-grained semantic classifications for these regions, i.e. identifying a region of text that is a paragraph and distinguishing that from a caption or document title.
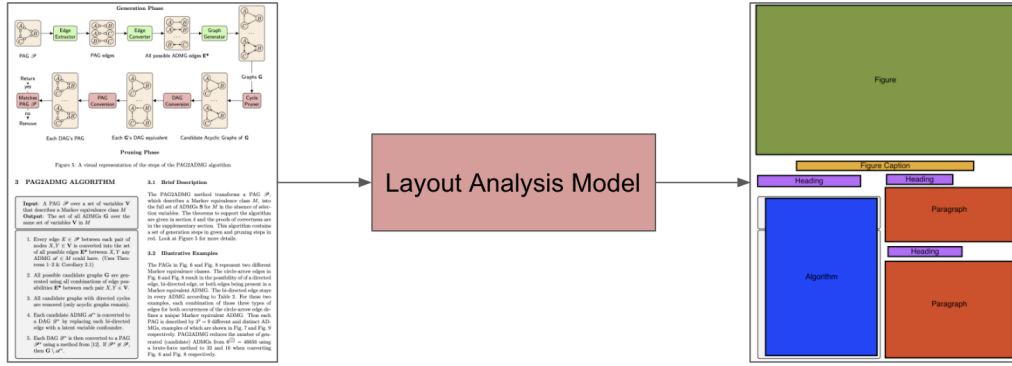
Figure 2: A document is passed through a generic layout analysis model, resulting in a layout segmentation mask with the following classes: figure (green), figure caption (orange), heading (purple), paragraph (red), and algorithm (blue). The document has been reproduced with permission (Subramani, 2016).

Research in methods for document layout analysis has a long history, both in academia and industry. [1] From the first pioneering heuristic approaches (Lebourgeois et al., 1992; Okun et al., 1999; Liang et al., 1997), to multi-stage classical machine learning systems (Qin et al., 2018; Wei et al., 2013; Eskenazi et al., 2017), the evolution of document layout analysis methods is now dominated by end-to-end differentiable methods (Yang et al., 2017; Binmakhashen and Mahmoud, 2019; Ares Oliveira et al., 2018; Agarwal et al., 2020; Monnier, 2020; Pramanik et al., 2020).

4.1 Instance Segmentation for Layout Analysis

When applied to the problem of layout analysis in business documents, instance segmentation methods predict per-pixel labels to categorize regions of interest. Such methods are flexible and easily adapt to the courser-grained task of page segmentation or the more-specific task of logical structural analysis.

In Yang et al. (2017), the authors describe an end-to-end neural network that combines both text and visual features in a encoder-decoder architecture that also incorporates an unsupervised pretraining network. During inference, their approach uses a downsampling cascade of pooling layers to encode visual information, which is fed into a symmetrical upsampling cascade for decoding. At each cascade level, the produced encoding is also directly passed into the respective decoding block, concatenating the down- and up-sampled representations. This architecture ensures that visual feature information at different levels of resolution is considered during the encoding and decoding process (Burt and Adelson, 1983). For the final decoding layer, localized text embeddings are supplied alongside the computed visual representation.

This U-Net inspired encoding-decoding architecture has been adopted for document layout analysis in several different approaches (Ronneberger et al., 2015). The method in Ares Oliveira et al. (2018), and later extended by Barman et al. (2020) via additional text embeddings, use convolution maxpooling layers with large filter sizes to feed the document image through a ResNet bottleneck (He et al., 2015). The representation is then processed by bilinear upsampling layers and smaller 1x1 and 3x3 convolution layers. Both works are used to perform layout analysis on historical documents and newspapers from multiple European languages, respectively. In Lee et al. (2019), the authors combine the U-Net architecture pattern with trainable multiplication layers. This layer type is specialized for extracting co-occurrence texture features from the network's convolution feature maps, which are effective for locating regions that have periodically repeating information, such as tables.

4.2 Addressing Data Scarcity and Alternative Approaches

Obtaining high quality training data for layout analysis is a labor intensive task that requires both mechanical precision and an understanding of the document's contents. As a consequence of the

---

[1]The first ISO standard that defined aspects of modern-day document layout analysis was drafted four decades ago: ISO 8613-1:1989

difficulties in layout annotation of documents from brand new domains, several approaches exist to either leverage structure in unlabeled data or use well-defined rule sets to generate synthetic labeled documents to further improve generalizability and performance of document layout analysis systems.

Masked language models such as BERT and RoBERTa have shown effective empirical performance on many downstream NLP tasks (Devlin et al., 2019; Liu et al., 2019b). Inspired by the pretraining strategy in BERT and RoBERTa, Xu et al. (2020) define a Masked Visual-Language Model, which randomly masks input tokens and uses the model to predict the masked tokens. Unlike BERT, their method provides the 2-D positional embedding of the token during this masked prediction task, which enables the model to combine both semantic and spatial relationships between textual elements. Mentioned earlier in section 4.1, Yang et al. (2017) introduce an auxiliary document image reconstruction task in their broader instance segmentation-based network. During training, this auxiliary module uses a separate upsampling decoder that, without the aid of skip connections, predicts the original pixel values from the encoded representation.

While pretraining lets practitioners gain more value from their unlabeled documents, this technique alone is not always sufficient to effectively surmount data scarcity concerns. Relying on the intuition that many business and academic documents have repeated patterns in both content as well as page-level organization, several approaches have emerged to manufacture synthetic, labeled data in order to provide data suitable for a pretraining-like routine (Zhong et al., 2019b). In Monnier (2020), the authors propose a three-stage method for synthesizing new labeled documents. First, they generate the document by randomly choosing the a document background from a set of nearly 200 known document backgrounds. Second, they use a grid based layout method to define both individual document element content and their respective sizes. Third, their process introduces corruptions, such as Gaussian blur and random image crops. This modular, rule-based synthetic document generation approach creates a heterogeneous dataset to make pretraining of layout analysis models more robust.

Alternatively, instead of defining rules to generate a heterogeneous set of documents, several synthesizing procedures take cues from data augmentation methods. Capobianco and Marinai (2017) and Journet et al. (2017) describe general-purpose toolkits that use an existing set of labeled documents to introduce deformations and perturbations in source images. Importantly, such changes to the training data are balanced so as to preserve the original semantic content while still exposing the model training to realistic errors that it must account for during inference on unseen data.

4.3   Datasets for Layout Analysis

Recently, there has been a deluge of datasets specifically targeting the document layout analysis problem. The International Conference on Document Analysis and Recognition (ICDAR) has produced several datasets from their various annual competitions; the most recent from 2017 and 2019 provide gold-standard data for document layout analysis and other document processing tasks (Gao et al., 2017; Clausner et al., 2019; Huang et al., 2019; Gao et al., 2019).

On the larger side, DocBank is a a collection of half a million document pages with token-level annotations suitable for training and evaluating document layout analysis systems (Li et al., 2020). The authors constructed this dataset using weak supervision (Hoffmann et al., 2011), matching data from the LaTeX source of known PDFs to form annotations. Similarly, Zhong et al. (2019b) created PubLayNet by automatically matching XML content representations for over one million PDFs on PubMed Central™, consiting of approximately 360 thousand document images. While not full document layout, Zhong et al. (2019a) have created PubTabNet from PubMed Central as well. Their data consists of 568 thousand table images alongside an HTML representations of content.

5   Information Extraction

The goal of information extraction for document understanding is to take documents that may have diverse layouts and extract information into a structured format. Examples include receipt understanding to identify item names, quantities, and prices and form understanding to identify different key-value pairs. Document extraction of information by humans goes beyond simply reading text on a page as it is often necessary to learn page layouts for complete understanding. As such, recent enhancements have extended text encoding strategies for documents by additionally encoding structural and visual information of text in a variety of ways.

5.1 2D Positional Embeddings

Multiple sequence tagging approaches have been proposed which augment current named entity recognition (NER) methods by embedding attributes of 2D bounding boxes and merging them with text embeddings to create models which are simultaneously aware of both context and spatial positioning when extracting information. Xu et al. (2020) embeds the pair of x, y coordinates that define a bounding box using two different embedding tables and pretrain a masked language model (LM). During pretraining, text is randomly masked but the 2D positional embeddings are retained. This model can then be fine-tuned on a downstream task. Alternatively, the bounding box coordinates can also be embedded using sin and cos functions like positional encoding methods (Vaswani et al., 2017; Hwang et al., 2020). Other features can also be embedded such as the line or sequence number (Hwang et al., 2019). In this scenario, the document is preprocessed to assign a line number to each individual token. Each token is then ordered from left to right and given a sequential position. Finally, both the line and sequential positions are embedded.

While these strategies have seen success, relying solely on the line number or bounding box coordinates can be misleading when the document has been scanned on an uneven surface, leading to curved text. Additionally, bounding box based embeddings still miss critical visual information such as typographical emphases (bold, italics) and images such as logos. To overcome these, a crop of the image corresponding to the token of interest can be embedded using a Faster R-CNN model to create token image embeddings which are combined with the 2D positional embeddings (Xu et al., 2020).

5.2 Image Embeddings

Information extraction for documents can also be framed as a computer vision challenge wherein the goal of the model is to semantically segment information or regress bounding boxes over the areas of interest. This strategy helps preserve the 2D layout of the document and allows models to take advantage of 2D correlations. While it is theoretically possible to learn strictly from the document image, directly embedding textual information into the image simplifies the task for models to understand the 2D textual relationships. In these cases, an encoding function is applied onto a proposed textual level (i.e. character, token, word) to create individual embedding vectors. These vectors are transposed into each pixel that comprises the bounding box corresponding to the embedded text, ultimately creating an image of $W \times H \times D$ where W is the width, H is the height, and D is the embedding dimension. Proposed variants are listed as following:

1. *CharGrid* embeds characters with a one-hot encoding into the image (Katti et al., 2018)
2. *WordGrid* embeds individual words using word2vec or FastText (Kerroumi et al., 2020)
3. *BERTgrid* finetunes BERT on task-specific documents and is used obtain contextual word-piece vectors (Denk and Reisswig, 2019)
4. *C+BERTgrid*, combines context-specific and character vectors (Denk and Reisswig, 2019)

When comparing the grid methods, C+BERTgrid has shown the best performance, likely due to its contextualized word vectors combined with a degree of resiliency to OCR errors. Zhao et al. (2019) proposes an alternative approach to directly apply text embeddings to the image. A grid is projected on top of the image and a mapping function assigns each token to a unique cell in the grid. Models then learn to assign each cell in the grid to a class. This method significantly reduces the dimensionality due to its grid system, while still retaining the majority of the 2D spatial relationships.

5.3 Documents as Graphs

Unstructured text on documents can also be represented as graph networks, where the nodes in a graph represent different textual segments. Two nodes are connected with an edge if they are cardinally adjacent to each other, allowing the relationship between words to be modeled directly (Qian et al., 2019). An encoder such as a BiLSTM encodes text segments into nodes (Qian et al., 2019). Edges can be represented as a binary adjacency matrix or a richer matrix, encoding additional visual information such as the distance between segments or shape of the source and target nodes (Liu et al., 2019a). A graph convolutional network is then applied at different receptive fields in a similar fashion to dilated convolutions (Yu and Koltun, 2015) to ensure that both local and global information can be learned (Qian et al., 2019). After this, the representation is passed to a sequence tagging decoder.

Documents can also be represented as a directed graph and a spatial dependency parser (Hwang et al., 2020). In this representation, nodes are represented by textual segments, but field nodes denoting the node type are used to initialize each DAG. In addition, two kinds of edges are defined:

1. Edges that group together segments belonging to the same category (STORENAME → Peet's → Coffee; a field node followed by two nodes representing a store name)

2. Edges that connect relationships between different groups (Peet's → 94107; a zipcode).

A transformer with an additional 2D positional embedding is used to spatially encode the text. After this, the task becomes to predict the relationship matrix for each edge type. This method can represent arbitrarily deep hierarchies and can be applied towards complicated document layouts.

5.4 Tables

Tabular data extraction remains a challenging aspect of information extraction due to their wide variety of formats and complex hierarchies. Table datasets typically have multiple tasks to perform (Shahab et al., 2010; Göbel et al., 2013; Gao et al., 2019; Zhong et al., 2019a). The first task is table detection which involves localizing the bounding box containing the table(s) inside the document. The next task is table structure recognition, which requires extracting the row, column, and cell information into a common format. This can be taken one step further to table recognition, which requires understanding both the structural information as well as the content by classifying cells within the table itself (Zhong et al., 2019a). As textual and visual features are equally important to properly extracting and understanding tables, many diverse methods have been proposed to perform this task.

One such proposal named TableSense performs both table detection and structure recognition (Dong et al., 2019a). TableSense uses a three stage approach: cell featurization, object detection with convolutional models, and an uncertainty-based active learning sampling mechanism. TableSense's proposed architecture for table detection performs significantly better than traditional methods in computer vision such as YOLO-v3 or Mask R-CNN (Redmon and Farhadi, 2018; He et al., 2017b). Since this approach does not work well for general spreadsheets, Dong et al. (2019b) extend upon the previous work by using a multitask framework to jointly learn table regions, structural components of spreadsheets, and cell types. They add an additional stage, which leverages language models to learn the semantic contents of table cells in order to flatten complex tables into a single standard format.

Wang et al. (2020) propose TUTA, which focuses on understanding the content within tables after the structure has been determined. The authors present three new objectives for language model pretraining for table understanding by using tree-based transformers. The objectives introduced for pretraining are designed to help the model understand tables at the token, cell, and table level. The authors mask a proportion of tokens depending on the table cell for the model to predict, randomly mask particular cell headers for the model to predict the header string based on its location, and provide the table with context such as table titles or descriptions that may or may not be associated for the model to identify which contextual elements are positively associated with the table. The transformer architecture is modified to reduce distractions from attention by limiting the attention connections to items based on a cell's hierarchical distance to another cell. Fine-tuning TUTA has demonstrated state of the art performance on multiple datasets for cell type classification.

6 Conclusion

Document understanding is a hot topic in industry and has immense monetary value. Most documents are private data corresponding to private contracts, invoices, and records. As a result, openly available datasets are hard to come by and have not been a focus for academia with respect to other application areas. The academic literature on methodologies to tackle document understanding is similarly sparse relative to areas with an abundance of publicly available data such as image classification and translation. However, the most effective approaches for document understanding make use of recent advancements in deep neural network modeling. End-to-end document understanding is achievable by creating an integrated system that performs layout analysis, optical character recognition, and domain-specific information extraction. In this survey, we attempt to consolidate and organize the methodologies which are present in literature in order to be a jumping-off point for scholars and practitioners alike who want to explore document understanding.

References

M. Agarwal, Ajoy Mondal, and C. Jawahar. 2020. Cdec-net: Composite deformable cascade network for table detection in document images. *ArXiv*.

A. Amin and R. Shiu. 2001. Page segmentation and classification utilizing bottom-up approach. *Int. J. Image Graph.*

Sofia Ares Oliveira, Benoit Seguin, and Frederic Kaplan. 2018. dhsegment: A generic deep-learning approach for document segmentation. *ICFHR*.

Youngmin Baek, Bado Lee, Dongyoon Han, Sangdoo Yun, and Hwalsuk Lee. 2019. Character region awareness for text detection. In *CVPR*.

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. In *ICLR*.

Raphaël Barman, Maud Ehrmann, Simon Clematide, Sofia Ares Oliveira, and Frédéric Kaplan. 2020. Combining visual and textual features for semantic segmentation of historical newspapers. *arXiv*.

Iz Beltagy, Matthew E. Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *ArXiv*.

Galal M. Binmakhashen and Sabri A. Mahmoud. 2019. Document layout analysis: A comprehensive survey. *ACM Comput. Surv.*

P. Burt and E. Adelson. 1983. The laplacian pyramid as a compact image code. *IEEE Transactions on Communications*.

S. Capobianco and S. Marinai. 2017. Docemul: A toolkit to generate structured historical documents. In *ICDAR*.

Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. 2019. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*.

Chee Kheng Ch'ng and Chee Seng Chan. 2017. Total-text: A comprehensive dataset for scene text detection and recognition. In *2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)*.

Kyunghyun Cho, B. V. Merrienboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014a. On the properties of neural machine translation: Encoder-decoder approaches. In *SSST@EMNLP*.

Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014b. Learning phrase representations using rnn encoder–decoder for statistical machine translation. In *EMNLP*.

Krzysztof Choromanski, Valerii Likhosherstov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Lukasz Kaiser, et al. 2020. Rethinking attention with performers. *arXiv preprint arXiv:2009.14794*.

C. Clausner, A. Antonacopoulos, and S. Pletschacher. 2019. Icdar2019 competition on recognition of documents with complex layouts - rdcl2019. *ICDAR*.

Ronan Collobert and Jason Weston. 2008. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *ICML*.

Dan Deng, Haifeng Liu, Xuelong Li, and Deng Cai. 2018. Pixellink: Detecting scene text via instance segmentation. In *AAAI*.

Timo I. Denk and C. Reisswig. 2019. Bertgrid: Contextualized embedding for 2d document representation and understanding. *ArXiv*, abs/1909.04948.

J. Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*.

Haoyu Dong, S. Liu, S. Han, Z. Fu, and D. Zhang. 2019a. Tablesense: Spreadsheet table detection with convolutional neural networks. In *AAAI*.

Haoyu Dong, Shijie Liu, Zhouyu Fu, Shi Han, and Dongmei Zhang. 2019b. Semantic structure extraction for spreadsheet tables with a multi-task learning architecture. In *Workshop on Document Intelligence at NeurIPS 2019*.

Sébastien Eskenazi, Petra Gomez-Krämer, and Jean-Marc Ogier. 2017. A comprehensive survey of mostly textual document segmentation algorithms since 2008. *Pattern Recognition*.

Wei Feng, Wenhao He, Fei Yin, Xu-Yao Zhang, and Cheng-Lin Liu. 2019. Textdragon: An end-to-end framework for arbitrary shaped text spotting. In *CVPR*.

L. Gao, Yilun Huang, H. Déjean, Jean-Luc Meunier, Q. Yan, Y. Fang, F. Kleber, and E. M. Lang. 2019. Icdar 2019 competition on table detection and recognition (ctdar). *ICDAR*.

L. Gao, X. Yi, Z. Jiang, L. Hao, and Z. Tang. 2017. Icdar2017 competition on page object detection. In *2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)*.

Jonas Gehring, M. Auli, David Grangier, Denis Yarats, and Yann Dauphin. 2017. Convolutional sequence to sequence learning. In *ICML*.

Ian J Goodfellow, Yaroslav Bulatov, Julian Ibarz, Sacha Arnoud, and Vinay Shet. 2013. Multi-digit number recognition from street view imagery using deep convolutional neural networks. In *ICLR*.

Alex Graves, Santiago Fernández, Faustino Gomez, and Jürgen Schmidhuber. 2006. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In *ICML*.

Ankush Gupta, Andrea Vedaldi, and Andrew Zisserman. 2016. Synthetic data for text localisation in natural images. In *CVPR*.

M. Göbel, T. Hassan, E. Oro, and G. Orsi. 2013. Icdar 2013 table competition. In *2013 12th International Conference on Document Analysis and Recognition*.

J. Ha, R. Haralick, and I. Phillips. 1995. Document page decomposition by the bounding-box project. *ICDAR*.

Dafang He, Xiao Yang, Chen Liang, Zihan Zhou, Alexander G Ororbi, Daniel Kifer, and C Lee Giles. 2017a. Multi-scale fcn with cascaded instance aware segmentation for arbitrary oriented word spotting in the wild. In *CVPR*.

Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. 2017b. Mask r-cnn. In *ICCV*.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2015. Deep residual learning for image recognition. In *CVPR*.

Tong He, Zhi Tian, Weilin Huang, Chunhua Shen, Yu Qiao, and Changming Sun. 2018. An end-to-end textspotter with explicit alignment and attention. In *CVPR*.

Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*.

Raphael Hoffmann, Congle Zhang, Xiao Ling, Luke Zettlemoyer, and Daniel S. Weld. 2011. Knowledge-based weak supervision for information extraction of overlapping relations. In *ACL*.

Z. Huang, Ke-Han Chen, Jianhua He, X. Bai, Dimosthenis Karatzas, S. Lu, and C. Jawahar. 2019. Icdar2019 competition on scanned receipt ocr and information extraction. *ICDAR*.

Wonseok Hwang, Seonghyeon Kim, Minjoon Seo, Jinyeong Yim, Seunghyun Park, Sungrae Park, Junyeop Lee, Bado Lee, and Hwalsuk Lee. 2019. Post-{ocr} parsing: building simple and robust parser via {bio} tagging. In *Workshop on Document Intelligence at NeurIPS 2019*.

Wonseok Hwang, Jinyeong Yim, Seunghyun Park, Sohee Yang, and Minjoon Seo. 2020. Spatial dependency parsing for semi-structured document information extraction. *arXiv*.

Max Jaderberg, K. Simonyan, A. Vedaldi, and Andrew Zisserman. 2014. Synthetic data and artificial neural networks for natural scene text recognition. *ArXiv*.

Guillaume Jaume, Hazim Kemal Ekenel, and Jean-Philippe Thiran. 2019. Funsd: A dataset for form understanding in noisy scanned documents. In *ICDARW*.

Mandar Joshi, Omer Levy, Daniel S. Weld, and Luke Zettlemoyer. 2019. Bert for coreference resolution: Baselines and analysis. In *EMNLP/IJCNLP*.

Nicholas Journet, Muriel Visani, Boris Mansencal, Kieu Van-Cuong, and Antoine Billy. 2017. Doccreator: A new software for creating synthetic ground-truthed document images. *Journal of Imaging*.

Dimosthenis Karatzas, Lluis Gomez-Bigorda, Anguelos Nicolaou, Suman Ghosh, Andrew Bagdanov, Masakazu Iwamura, Jiri Matas, Lukas Neumann, Vijay Ramaseshan Chandrasekhar, Shijian Lu, et al. 2015. Icdar 2015 competition on robust reading. In *ICDAR*.

Dimosthenis Karatzas, Faisal Shafait, Seiichi Uchida, Masakazu Iwamura, Lluis Gomez i Bigorda, Sergi Robles Mestre, Joan Mas, David Fernandez Mota, Jon Almazan Almazan, and Lluis Pere De Las Heras. 2013. Icdar 2013 robust reading competition. In *CVPR*.

Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. 2020. Transformers are rnns: Fast autoregressive transformers with linear attention. In *ICML*.

A. R. Katti, C. Reisswig, Cordula Guder, Sebastian Brarda, S. Bickel, J. Höhne, and Jean Baptiste Faddoul. 2018. Chargrid: Towards understanding 2d documents. In *EMNLP*.

Mohamed Kerroumi, Othmane Sayem, and Aymen Shabou. 2020. Visualwordgrid: Information extraction from scanned documents using a multimodal approach. *ArXiv*.

Nikita Kitaev, Łukasz Kaiser, and Anselm Levskaya. 2020. Reformer: The efficient transformer. In *ICLR*.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *NeurIPS*.

F. Lebourgeois, Z. Bublinski, and H. Emptoz. 1992. A fast and efficient method for extracting text paragraphs and graphics from unconstrained documents. In *ICPR*.

Joonho Lee, Hideaki Hayashi, Wataru Ohyama, and Seiichi Uchida. 2019. Page segmentation using a convolutional neural network with trainable co-occurrence features. In *ICDAR*.

Hui Li, Peng Wang, and Chunhua Shen. 2017. Towards end-to-end text spotting with convolutional recurrent neural networks. In *CVPR*.

Minghao Li, Yiheng Xu, Lei Cui, Shaohan Huang, Furu Wei, Zhoujun Li, and Ming Zhou. 2020. Docbank: A benchmark dataset for document layout analysis. *arXiv*.

Jisheng Liang, Richard Rogers, Robert M Haralick, and Ihsin T Phillips. 1997. Uw-isl document image analysis toolbox: An experimental environment. In *ICDAR*.

Minghui Liao, Guan Pang, Jing Huang, Tal Hassner, and Xiang Bai. 2020. Mask textspotter v3: Segmentation proposal network for robust scene text spotting. *arXiv*.

Minghui Liao, Baoguang Shi, and Xiang Bai. 2018a. Textboxes++: A single-shot oriented scene text detector. *IEEE Transactions on Image Processing*.

Minghui Liao, Baoguang Shi, Xiang Bai, Xinggang Wang, and Wenyu Liu. 2016. Textboxes: A fast text detector with a single deep neural network. In *AAAI*.

Minghui Liao, Zhen Zhu, Baoguang Shi, Gui-song Xia, and Xiang Bai. 2018b. Rotation-sensitive regression for oriented scene text detection. In *CVPR*.

Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. 2017. Feature pyramid networks for object detection. In *CVPR*.

Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. 2016. Ssd: Single shot multibox detector. In *ECCV*.

Xiaojing Liu, Feiyu Gao, Qiong Zhang, and Huasha Zhao. 2019a. Graph convolution for multimodal information extraction from visually rich documents. In *NAACL*.

Xuebo Liu, Ding Liang, Shi Yan, Dagui Chen, Yu Qiao, and Junjie Yan. 2018. Fots: Fast oriented text spotting with a unified network. In *CVPR*.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. Roberta: A robustly optimized bert pretraining approach. *arXiv*.

Yuliang Liu and Lianwen Jin. 2017. Deep matching prior network: Toward tighter multi-oriented text detection. In *CVPR*.

Jonathan Long, Evan Shelhamer, and Trevor Darrell. 2015. Fully convolutional networks for semantic segmentation. In *CVPR*.

Shangbang Long, Jiaqiang Ruan, Wenjie Zhang, Xin He, Wenhao Wu, and Cong Yao. 2018. Textsnake: A flexible representation for detecting text of arbitrary shapes. In *ECCV*.

David G Lowe. 2004. Distinctive image features from scale-invariant keypoints. *IJCV*.

Jiri Matas, Ondrej Chum, Martin Urban, and Tomás Pajdla. 2004. Robust wide-baseline stereo from maximally stable extremal regions. *Image and vision computing*.

Tom Monnier. 2020. docextractor: An off-the-shelf historical document element extraction. In *Semantic Scholar*.

Shunji Mori, Hirobumi Nishida, and Hiromitsu Yamada. 1999. *Optical character recognition*. John Wiley & Sons, Inc.

Oleg Okun, David Doermann, and Matti Pietikainen. 1999. Page segmentation and zone classification: The state of the art. Technical report, OULU UNIV (FINLAND) DEPT OF ELECTRICAL ENGINEERING.

Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *NAACL-HLT*.

Subhojeet Pramanik, Shashank Mujumdar, and Hima Patel. 2020. Towards a multi-modal, multi-task learning based pre-training framework for document representation learning. *arXiv*.

Yujie Qian, Enrico Santus, Zhijing Jin, Jiang Guo, and Regina Barzilay. 2019. GraphIE: A graph-based framework for information extraction. In *NAACL*.

W. Qin, R. Elanwar, and M. Betke. 2018. Laba: Logical layout analysis of book page images in arabic using multiple support vector machines. In *ASAR*.

Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. 2016. You only look once: Unified, real-time object detection. In *CVPR*.

Joseph Redmon and Ali Farhadi. 2018. Yolov3: An incremental improvement. *ArXiv*.

Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*.

Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *ACL*.

Asif Shahab, F. Shafait, T. Kieninger, and A. Dengel. 2010. An open approach towards the benchmarking of table structure recognition systems. In *DAS '10*.

Baoguang Shi, Xinggang Wang, Pengyuan Lyu, Cong Yao, and X. Bai. 2016. Robust scene text recognition with automatic rectification. *CVPR*.

Nishant Subramani. 2016. Pag2admg: An algorithm for the complete causal enumeration of a markov equivalence class. *arXiv preprint arXiv:1612.00099*.

Nishant Subramani, Samuel R. Bowman, and Kyunghyun Cho. 2019. Can unconditional language models recover arbitrary sentences? In *NeurIPS*.

Nishant Subramani and D. Rao. 2020. Learning efficient representations for fake speech detection. In *AAAI*.

Nishant Subramani and Nivedita Suresh. 2020. Discovering useful sentence representations from large pretrained language models. *ArXiv*.

Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *NeurIPS*.

Zhi Tian, Weilin Huang, Tong He, Pan He, and Yu Qiao. 2016. Detecting text in natural image with connectionist text proposal network. In *ECCV*.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NeurIPS*.

J. Wang and Xiaolin Hu. 2017. Gated recurrent convolution neural network for ocr. In *NeurIPS*.

Zhiruo Wang, Haoyu Dong, Ran Jia, Jiugang Li, Zhiyi Fu, Shi Han, and Dongmei Zhang. 2020. Structure-aware pre-training for table understanding with tree-based transformers. *ArXiv*.

Hao Wei, Micheal Baechler, Fouad Slimane, and Rolf Ingold. 2013. Evaluation of svm, mlp and gmm classifiers for layout analysis of historical documents. In *ICDAR*.

Xiang Wu, R. He, Z. Sun, and T. Tan. 2018. A light cnn for deep face representation with noisy labels. *IEEE Transactions on Information Forensics and Security*.

Qizhe Xie, Zihang Dai, E. Hovy, Minh-Thang Luong, and Quoc V. Le. 2019. Unsupervised data augmentation for consistency training. *ArXiv*.

Yiheng Xu, Minghao Li, Lei Cui, Shaohan Huang, Furu Wei, and Ming Zhou. 2020. Layoutlm pre-training of text and layout for document image understanding. *KDD*.

Xiao Yang, Ersin Yumer, Paul Asente, Mike Kraley, Daniel Kifer, and C Lee Giles. 2017. Learning to extract semantic structure from documents using multimodal fully convolutional neural networks. In *CVPR*.

Cong Yao, Xiang Bai, Nong Sang, Xinyu Zhou, Shuchang Zhou, and Zhimin Cao. 2016. Scene text detection via holistic, multi-channel prediction. *arXiv*.

Jian Ye, Z. Chen, J. Liu, and Bo Du. 2020. Textfusenet: Scene text detection with richer fused features. In *IJCAI*.

Fisher Yu and Vladlen Koltun. 2015. Multi-scale context aggregation by dilated convolutions. In *ICLR*.

Liu Yuliang, Jin Lianwen, Zhang Shuaitao, and Zhang Sheng. 2017. Detecting curve text in the wild: New dataset and new solution. *arXiv*.

Xiaohui Zhao, Endi Niu, Zhuo Wu, and Xiaoguang Wang. 2019. Cutie: Learning to understand documents with convolutional universal text information extractor. *ArXiv*.

Xu Zhong, Elaheh ShafieiBavani, and Antonio Jimeno-Yepes. 2019a. Image-based table recognition: data, model, and evaluation. In *ECCV*.

Xu Zhong, Jianbin Tang, and Antonio Jimeno Yepes. 2019b. Publaynet: largest dataset ever for document layout analysis. In *ICDAR*.