



# TabAug: Data Driven Augmentation for Enhanced Table Structure Recognition

Umar Khan<sup>1</sup>(✉), Sohaib Zahid<sup>1</sup>, Muhammad Asad Ali<sup>1</sup>, Adnan Ul-Hasan<sup>1</sup>,  
and Faisal Shafait<sup>1,2</sup>

<sup>1</sup> Deep Learning Laboratory, National Center of Artificial Intelligence,  
Islamabad, Pakistan

{umar.khan1,sohaib.zahid,muhammad.asadali,adnan.ulhassan,  
faisal.shafait}@seecs.edu.pk

<sup>2</sup> School of Electrical Engineering and Computer Science,  
National University of Sciences and Technology (NUST), Islamabad, Pakistan

**Abstract.** Table Structure Recognition is an essential part of end-to-end tabular data extraction in document images. The recent success of deep learning model architectures in computer vision remains to be non-reflective in table structure recognition, largely because extensive datasets for this domain are still unavailable while annotating new data is expensive and time-consuming. Traditionally, in computer vision, these challenges are addressed by standard augmentation techniques that are based on image transformations like color jittering and random cropping. As demonstrated by our experiments, these techniques are not effective for the task of table structure recognition. In this paper, we propose TabAug, a re-imagined Data Augmentation technique that produces structural changes in table images through replication and deletion of rows and columns. It also consists of a data-driven probabilistic model that allows control over the augmentation process. To demonstrate the efficacy of our approach, we perform experimentation on ICDAR 2013 dataset where our approach shows consistent improvements in all aspects of the evaluation metrics, with cell-level correct detections improving from 92.16% to 96.11% over the baseline.

**Keywords:** Table Structure Recognition · Table augmentation · Data augmentation · Table data extraction · Probabilistic model · Data-driven model · Table segmentation · Deep learning

## 1 Introduction

Document structure analysis and parsing are some of the most crucial parts of document image processing for digitization and information extraction. One

---

U. Khan and S. Zahid—These authors have contributed equally.

© Springer Nature Switzerland AG 2021

J. Lladós et al. (Eds.): ICDAR 2021, LNCS 12822, pp. 585–601, 2021.

[https://doi.org/10.1007/978-3-030-86331-9\\_38](https://doi.org/10.1007/978-3-030-86331-9_38)


of the most important components of documents is tables. Tables are a structured way of representing data allowing for visual and logical grouping of data in a highly comprehensible manner. Tables in documents are frequently used to present key information such as financial records, receipts, and data forms. Extracting this information can be vital and beneficial to most businesses around the world. However, tabular data extraction can be a challenging task as more often this data is found as scanned document images that contain no structural or content metadata.

Tabular data extraction consists of three major tasks, Table Detection, Table Structure Recognition, and Semantic Understanding. Table detection is the task of locating tables in a given document, while table structure recognition aims towards the segmentation of tables into rows and columns for layout understanding. Semantic understanding involves the assignment of information such as row or column header (e.g., Unit price, Stock value, etc.) to a particular cell. Our focus in this work is the table structure recognition part, where the aim is to extract rows and columns in a given table image.

In practice, Convolutional Neural Network (CNN) is used as an effective tool for extracting meaningful features from visual information. Given that scanned document images contain only visual information, CNNs become increasingly important for Table Detection and Table Structure Recognition alike. State-of-the-art CNNs are efficient at extracting visual features from data; however, they can be highly data demanding in nature. This precondition of CNNs coupled with the unavailability of large Table Structure Recognition datasets presents a challenging problem as annotating large datasets can be both expensive and time-consuming. It is, therefore, crucial to focus on methods for improving the data-efficiency of deep learning models to achieve better results for Table Structure Recognition on smaller datasets.

Augmentation techniques are widely used to improve the data efficiency in Deep Neural networks. It is the practice of adding slight realistic changes to the original data to increase the diversity of the training data. This concept helps model regularize and avoid over-fitting in small datasets. There are several image transformation techniques such as scaling, rotating, smearing, etc. that are widely used in Computer Vision to augment data. For the purposes of comparison, we choose to apply Random cropping, (first employed in AlexNet [14]) and Color Jitter (used in re-implementation of ResNet [9] by Facebook AI Research) as the *standard* augmentation. Please refer to the bottom row of Fig. 2 for example of these transformations on ICDAR 2013 dataset. However, as shown by our experimentation in Sect. 4, they are not effective for table structure recognition as they do not alter the structure of a table but instead produce unrepresentative data, which results in a decreased performance.

To overcome these challenges, we propose a novel method of data augmentation based on two key operations of *Replication* and *Deletion* on rows and columns, illustrated in Fig. 1 and sample images of these illustrated operations can be seen in Fig. 2. We also introduce a data-driven probabilistic model for generating parameters that control the outlook of the augmented data. The aug-



**Fig. 1.** Visual depiction of proposed augmentation operations

mented variations generated by this technique better represent the real-world variation in tables. Please refer to Sect. 3 for further details about our methodology.

Our experiments on ICDAR 2013 dataset demonstrate that our proposed augmentation technique yields a higher data efficiency and out-performs both the non-augmented and standard image augmentation approaches. Please refer to Sect. 4 for further details on experimental evaluation.

These results set a strong foundation for more future work in this direction. To facilitate this, we have decided to make our implementation of structural data augmentation named *TabAug* and associated experimentation code open-source<sup>1</sup>. We have also developed the existing code to be fairly simple to plug into various existing models with support for T-Truth annotation format [20] with no training overhead.

The rest of paper is structured into following sections. Section 2, outlines literature review of relevant works in the domain of table structure recognition and tabular data extraction in general. In Sect. 3 we present our methodology and implementation of the proposed augmentation technique. In Sect. 4 we report on our experiments and comparative results. Finally, in Sect. 5 we conclude our research with future direction for our work.

## 2 Related Work

Tabular data extraction is an old and recognized problem with solutions maturing from over past 20 years. Starting with one of the earliest works in table detection in text files using heuristics was done by [24] in 1996 followed by Pyreddy et al. [15] in 1997. The following years saw a novel heuristic-based approach of detecting tables in document images from Keinenger et al. [11, 12] and [13] forming a combined table plotting and recognition system named TRECS. Zanibbi et al. [25] presented a comprehensive survey paper that focused on table recognition literature writing down extensive observations on various aspects of popular

<sup>1</sup> <https://github.com/sohaib023/splerge-tab-aug>.

methods available at the time. A novel approach of detecting tables from ruling lines was presented by Basilios et al. [6]. In 2010 Shahab et al. [20] presented comprehensive and rigorous evaluation metrics for table detection and structure recognition. Significant work was done by Shafait et al. [19] in developing table detection algorithm for multiple layouts, they also introduced more meaningful performance metrics for table detection.

Use of data-driven approaches in table detection started from Chen et al. [3] in 2011 where SVMs were used for table detection in handwritten documents with noise and artifacts. In 2013 Kasar et al. [10] also made use of SVMs for table region detection by classifying ruling lines. The following year Anukriti et al. [2] used Conditional Random Fields (CRFs) with encoded foreground block characteristics and the contextual information as features for learning layouts and labeling table and non-table regions in documents.

From 2015 onward, Deep Learning models have been used extensively for solving the detection and recognition problem. [1,8,18] and [22] made use of Faster R-CNN [17] with their own take on handcrafted features and pre-processing methodologies. These methods were successful in producing state-of-the-art results for table detection but failed to produce significant results in Structure Recognition. Schreiber et al. [18] mentions the unnatural approach of detecting rows and columns through Faster R-CNN [17] and instead proposes fine-grained image segmentation through FCN-X's architecture by Shelhamer et al. [21], however, their proposed technique heavily relies upon image stretching for expanding background pixels as a pre-processing methodology.

In more recent works from Qasim et al. [16] and Tensmeyer et al. [23] on table structure recognition, we see a more stable and natural approach towards the formulation of problem. Qasim et al. [16] made use of Graph Neural Networks for generating cell adjacency matrix for all existing OCR detected words and Tensmeyer et al. [23] formulating the problem of structure recognition as a combination of row, column splits and cell mergings for defining the structure of a table. Qasim et al. [16] mentions that lack of large datasets has been a major hindrance for Deep Learning methods in table structure recognition and makes use of synthetic data to show the effectiveness of their network. However, synthetic data generated randomly is limited in capturing the visual and general characteristics of the target dataset, due to which the model fails to perform on the target dataset despite the optimal results on the synthetic data. Similarly for the split model from Tensmeyer et al. [23], even though more data-efficient than Qasim et al. [16] GNN model, makes use of large proprietary dataset to train the network for an effective model to be tested on ICDAR 2013. Our experiments of training and testing of split model from Tensmeyer et al. [23] on ICDAR 2013 dataset reveals sub-optimal results than the network potential due to the lack of diverse and large training dataset.

With newer deep learning algorithms, we are seeing a trend of increase in performance on larger datasets; however, their results translate to sub-optimal results on smaller datasets thus limiting their potential use cases. Traditionally in computer vision, these problems are addressed with image transforma-

tion based data augmentation techniques defined as standard data augmentation for increasing data diversity in training data. Albeit successful in natural images, standard augmentation hold an insignificant and unstable impact on table images. In this paper, we propose a new re-imagined Tabular Data Augmentation inspired from [4] and [7] by replicating and removing table structure elements (rows and columns) to form augmented tables all while maintaining their visual artifacts, we have also introduced a data-driven probabilistic model (similar to [5]) for decision parameters of our method of Tabular Data Augmentation.

### 3 Methodology

We adopt an elementary approach of structural data augmentation for table structure recognition. We consider cells to be the building block of a table, that is, a table is defined as a combination of cells. Therefore, it should theoretically be possible to achieve a large number of combinations given a small number of cells. However, combining different cells into a table poses many limitations with regards to compatibility. First, widths and heights of cells can vary greatly, and thus combining them without due deliberation can lead to unnatural tables. Second, cells tend to inherit their formatting and styling from their adjacent cells. In a fully randomized collection of cells, one cell might have left-justified text while the next might have right-justified text, making for a confusing and unnatural table.

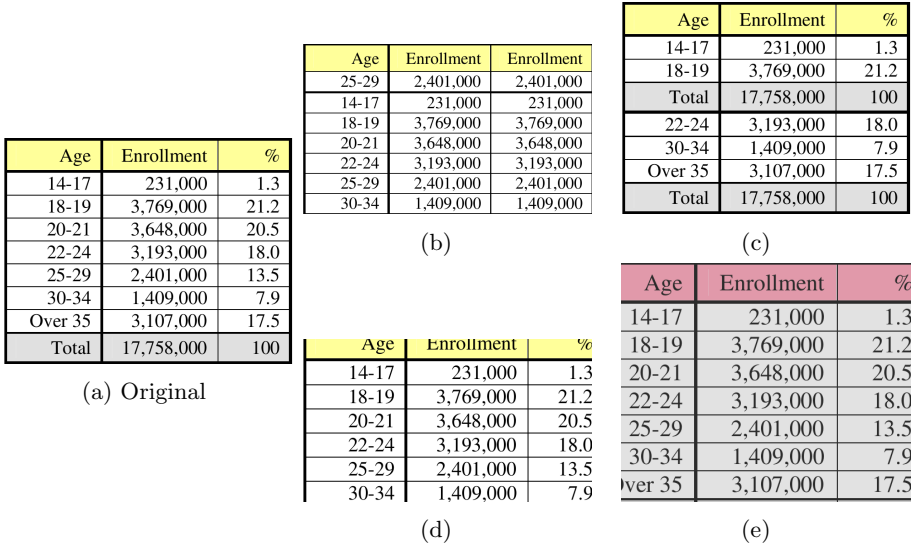
To overcome these limitations, we redefine a table to be a combination of rows and columns rather than cells. This helps keep intact the co-relation of cell styling within a row and column. To further simplify the problem, we limit these combinations to randomization of rows and columns within the same table. Thus, we can generate a randomized permutation of the rows and columns within a table to obtain a structurally augmented version of the table. This augmentation technique is then combined with a data-driven probabilistic sampler for maximizing its effectiveness.

#### 3.1 Augmentation Operations

To achieve randomization of rows and columns, we define four basic operations that we can apply to each table:

- Row deletion
- Column deletion
- Row replication
- Column replication

These operations are depicted visually in Fig. 1 for better understanding. These are the core atomic operations that we utilize during our augmentation pipeline. They can be applied sequentially in random orders on a table to achieve increasingly varying versions of the same table.



**Fig. 2.** Sub-figure 2a contains an original table sample from ICDAR 2013 dataset. Sub-figures 2b and 2c show variations of the table generated by TabAug. Sub-figures 2d and 2e show the results of image transformation based augmentations.

### 3.2 Augmentation Sub-operations

Each of the proposed augmentation operations consists of several sub-operations, which are further explained below. For simplicity, we will only consider Column deletion and replication, as Row deletion and replication can be directly inferred from it.

**Source Selection:** Before doing any replication or removal, we need to select a column on which we may apply the operation. We randomly select an index  $c$  in the range of  $\{1..C - 1\}$ , where  $C$  is the total number of columns in a given table. We purposefully leave out the  $0 - th$  index as the first column can have row headers (similarly column headers in the case of row selection), which should be retained in its location to maintain a natural table.


Furthermore, in case if spanning cells exist in the selected column, we must ensure that no partial/broken cells are copied.

$$\text{span}_{\min} = \min(\{\text{cells}_{r,c}^{\text{Start Column}} \forall r \in \{0..R - 1\}\}) \tag{1}$$

$$\text{span}_{\max} = \max(\{\text{cells}_{r,c}^{\text{End Column}} \forall r \in \{0..R - 1\}\}) \tag{2}$$

For the selected column  $c$ , the values of  $\text{span}_{\min}$  and  $\text{span}_{\max}$  are calculated using Eqs. 1 and 2. These values provide us with column indices for a self contained convex block (as depicted in Fig. 3a and 3b). In case if even after performing the above steps the selected block is non-convex (as depicted in Fig. 3c)

we abort the operation and return the table unaltered. Otherwise our selection is defined by  $c_{\max} = \text{span}_{\max}$  and  $c_{\min} = \text{span}_{\min}$  which is used by next sub-operations.



**Fig. 3.** A non-convex selection 3a cannot be replicated or deleted and hence it is expanded so as to make it convex 3b. Sub-figure 3c represents the case where expansion also results in a non-convex selection.

**Target Location Selection (only for Replication):** This sub-operation is similar to the previous one with few modifications. We select an index  $d$  in the range of  $\{1..C\}$  randomly, where  $C$  is the total number of columns in a given table. The target location for column replication will be at the start of  $d$ -th column. So we purposefully leave out the 0 - th index as we do not want to move the first column. Furthermore, if  $d = C$  then the target location is after the last column at the end of the image. We perform a check to see if the target location  $d$  is intersecting with a spanning cell. If so, placing a column at that location will cause the existing spanning cell to split apart, and hence we change  $d$  according to the equations below:

$$d = \begin{cases} \text{span}_{\min}, & \text{if } \text{abs}(d - \text{span}_{\min}) \leq \text{abs}(d - \text{span}_{\max}) \\ & \text{and } \text{span}_{\min} \neq 0 \\ \text{span}_{\max} + 1, & \text{otherwise} \end{cases} \quad (3)$$

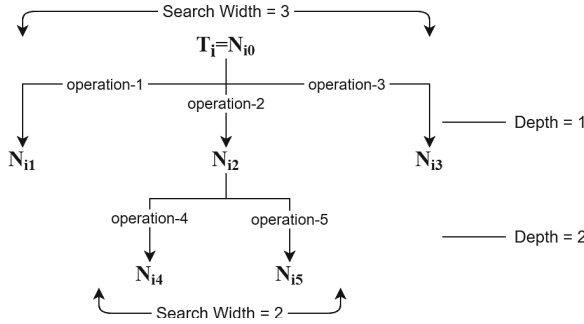
where  $\text{span}_{\min}$  and  $\text{span}_{\max}$  are calculated using Eqs. 1 and 2 by setting  $c = d$ .

Similar to the previous sub-operation, if after the correction of  $d$  it intersects with a spanning cell we abort the operation and return the table unaltered.

**Execution:** To execute the operation, whether deletion or replication, we need to alter both the table image and the ground-truth of the table, the specifics of which vary between the two operations. Following values are required in both operations:

$$x_{\min} = \text{columns}[c_{\min}].x1 \quad , \quad x_{\max} = \text{columns}[c_{\max}].x2 \quad , \quad w = x_{\max} - x_{\min}$$

where  $x1$  denotes starting x-coordinate of a column and  $x2$  denotes the ending x-coordinate of a column.



**Fig. 4.** An example of an augmentation search tree. Each of the nodes  $N_{ij}$  is obtained by applying an augmentation operation to its parent node, where the root of the tree is the original table  $T_i$ .

*Deletion:* Firstly, we remove the columns with indices in the range  $(c_{\min}, c_{\max})$  from the ground-truth, including all the contained cells and span information. For the columns and cells having indices greater than  $c_{\max}$ , we subtract  $w$  from their x-coordinates, to move them to the left.

For the image, we cut out the image from  $x_{\min}$  to  $x_{\max}$  and then move left the image pixels to the right side of  $x_{\max}$  by  $w$ .

*Replication:* Firstly, we calculate the x-coordinate of the target location where the replicated column is to be placed.

$$x_{\text{dst}} = \text{columns}[d].x1 \tag{4}$$

Then we copy the columns with indices in the range  $(c_{\min}, c_{\max})$  from the ground-truth, including all the contained cells and span information. For the copied columns and cells we offset their x-coordinates by  $x_{\text{dst}} - x_{\min}$ , so as to move it to the target location. Further, the columns and cells having indices greater than  $d$ , we add  $w$  to their x-coordinates, to move them to the right, clearing up space for the replicated column.

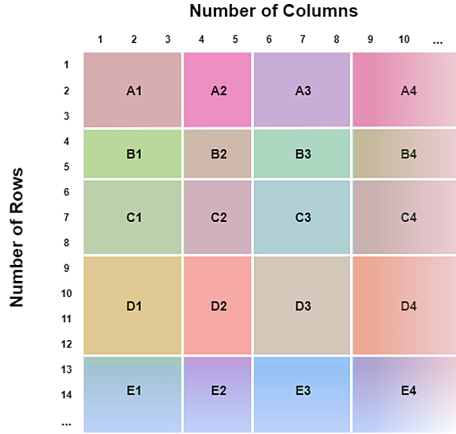
For the image, we move the image pixels to the right of  $x_{\text{dst}}$  by  $w$ , so as to clear up space for crop of the replicated columns. Then we copy the image pixels from  $x_{\min}$  to  $x_{\max}$  and move them by  $x_{\text{dst}} - x_{\min}$ , so as to move it to destination location  $x_{\text{dst}}$ .

### 3.3 Augmentation Pipeline

Having defined the augmentation operations, an effective pipeline must be established that can utilize the proposed operations for training a model, therefore we propose a pipeline in the following.

**Augmentation Tree Exploration:** Before training a model, we pre-compute a sample set  $N_i$  of augmented versions of a given table  $T_i$ . Each member of  $N_i$





**Fig. 5.** A grid enumerating all of the possible table categories. Each table is assigned one of these categories based on its number of rows and columns

is considered to be a node (represented as  $N_{ij}$ ), which is achievable by applying a series of augmentation operations (edges in the tree). An example of such a partial tree is shown in Fig. 4. By applying this tree search, we can extract a controlled sample set of all achievable augmented tables of a given root table  $T_i$ . To avoid inundation of samples, we apply pruning to the tree search to keep the number within a reasonable limit. First, only the nodes having tree depth greater than 5 and less than 10 are kept as a part of this sample set. Further, the maximum search width for each depth level is also restricted for pruning purposes. For depth = 1, maximum search width is 8, for depth = 2, it is 4, for depth = {3..5} it is 2, and finally for depth = {6..10} it is 1. Furthermore, out of the obtained sample set, any node is ignored that has a pixel height or width greater than 1.5 times the original table  $T_i$ , to ensure unnaturally large tables are not used for training.

**Categorization of Tables and Their Nodes:** As a pre-processing step, both the tables and their nodes are mapped into categories based on their number of rows and columns. The mapping table is of size  $5 \times 4$  as depicted in Fig. 5. An example of such a mapping is that a node with 4 rows and 5 columns will be assigned to  $B2$  category. Further, if the number of columns in  $N_{ij}$  are greater than 10, it is still assigned category  $\{A - E\}4$ . Similarly, if the number of rows in  $N_{ij}$  are greater than 14, it is assigned category  $E\{1 - 4\}$ .

After the completion of this step each  $T_i$  and  $N_{ij}$  will have a category associated with it which is referenced as  $T_i^{cat}$  and  $N_{ij}^{cat}$  respectively. It must be noted that the decision of category boundaries is purely empirical.

**Probability Based Selection:** For each table  $T_i$  we construct a probability distribution  $P_i$  using which  $T_i$ 's nodes will be sampled during training. Firstly,

a global frequency distribution of the tables over table categories is generated, giving us  $F^G$  a  $5 \times 4$  matrix. Further, for each table  $T_i$ , a frequency distribution of all  $N_{ij}$  over the table categories is generated, giving us  $F_i$ , which is another  $5 \times 4$  matrix. Lastly, we generate another  $5 \times 4$  matrix  $gauss_i$  which is a 2-D Gaussian centered around  $T_i^{cat}$ . The spread of this Gaussian distribution allows control over the diversity of the nodes  $N_{ij}$  that are sampled during training. These 3 matrices are multiplied to obtain  $P_i = Gauss_i * F^G * F_i$ .

During training, a random table category is selected using  $P_i$  as the probability distribution. Once the category is selected, one node is randomly sampled from all the nodes of  $T_i$  having the selected category. This selected node is passed for training.

## 4 Experiments and Results

To evaluate the efficacy of our proposed augmentation methodology, we train the Split model proposed by [23] on ICDAR 2013 dataset. We train the model using three different methodologies:

1. **Non-Augmented:** Images fed for training without any modification.
2. **Standard:** Basic image transformations, such as, hue saturation and brightness jitter in combination with image cropping.
3. **TabAug:** Our proposed augmentation methodology.

The dataset has a total of 128 pages and 156 tables, which is divided into train, test and validation splits with a proportion of  $0.72 : 0.2 : 0.08$  respectively. We train the model with 4 different percentages of the training set, that is, 25%, 50%, 75% and 100%. It must be noted that this is a ratio of the total training set utilized while training and not a ratio of training samples divided by total samples. Further, we repeat each experiment 3 times to get a better estimate of the average results and their deviation.

For training, we use a batch size of 1, and a learning rate of 0.00075. After every 15 epoch, we decay the learning rate by a factor of 0.8. Lastly, we train the models for a total of 27,500 iterations, which we empirically found to be enough for convergence of the model.

### 4.1 Ground Truth

For ICDAR 2013 dataset, ground-truths are provided as bounding boxes of cells, along with their starting and ending row/column indices. This ground-truth format was converted to T-Truth format [20], however, the resulting row and column separators were not aligned with the ruling lines of the table, hence, we re-annotated the dataset using T-Truth to align the separators and then cross-checked it to ensure that the annotations were coherent. The models are trained and tested on this re-annotated dataset, however, as they have been manually cross-checked, the evaluation is applicable to original annotations as well.

During the re-annotation phase, we upgraded the T-Truth annotation tool to fix several bugs and make it user-friendly. The upgraded version has been made publicly available<sup>2</sup>.

While T-Truth annotation format is used for the augmentation process, Split model [23] requires pixel-wise annotations. To generate such ground-truth, the separators of T-Truth annotations are expanded to the nearest words (horizontally in the case of column ground-truth, and vertically in the case of row ground-truth).

## 4.2 Performance Measures

We chose the evaluation metrics proposed by Shahab et al. [20] primarily for two reasons. First, these metrics are comprehensive, painting a complete picture of a model's performance. Second, these are general-purpose metrics and can be applied to any type of segment such as rows, columns, and cells.

In the evaluation metrics proposed by Shahab et al. [20], first the ground truth segments and the predicted segments are numbered. A correspondence matrix of shape  $n \times m$  is created where  $n$  is the number of ground truth segments while  $m$  is the number of predicted segments. Each entry  $[i, j]$  in the matrix stores the number of pixels that are intersecting between the given ground-truth segment  $G_i$  and predicted segment  $S_j$ , that is  $|G_i \cap S_j|$ . Further, the sum of  $i$ -th row in the correspondence matrix gives the total number of pixels in ground-truth segment  $G_i$  and the sum of  $j$ -th column gives the total number of pixels in predicted segment  $S_j$ . Once this correspondence matrix is generated, we define the following evaluation metrics using the threshold value of  $T = 0.1$  (consequently  $1 - T = 0.9$ ):

1. **Correct Detections:** The total number of ground-truth segments that have a one-to-one mapping with a predicted segment and a major overlap. Concretely, a given ground-truth segment  $G_i$  is considered to be a correct detection if it has a major overlap with a predicted segment  $S_j$  and  $S_j$  does not have a significant overlap with any other ground-truth segment ( $G_k; \forall k \neq i$ ). That is:

$$\frac{|G_i \cap S_j|}{G_i} > 1 - T \text{ and } \frac{|G_k \cap S_j|}{G_k} < T \ ; \forall k \neq i$$

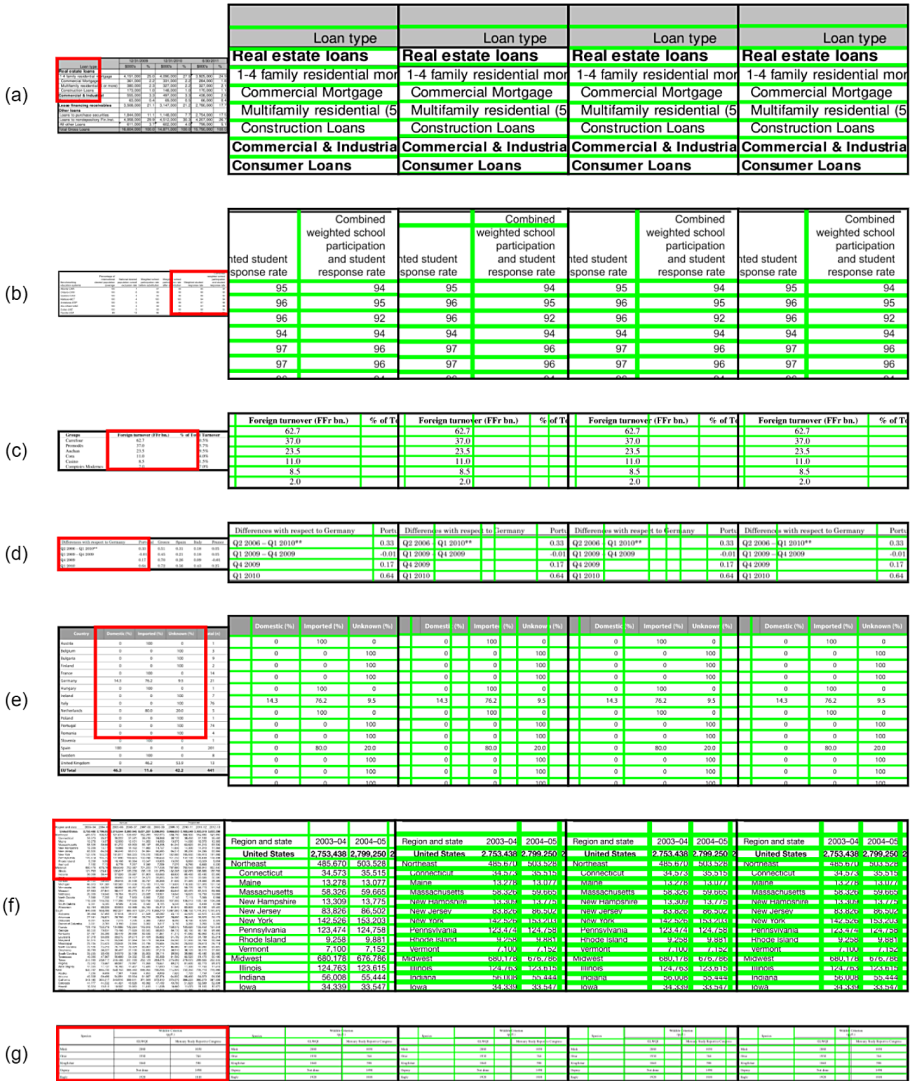
2. **Over-Segmentations:** The total number of ground-truth segments that have a significant overlap with more than one predicted segment. That is, a ground-truth segment  $G_i$  is over-segmented if:

$$T < \frac{|G_i \cap S_j|}{G_i} < 1 - T \text{ and } T < \frac{|G_i \cap S_k|}{G_i} < 1 - T \text{ where; } k \neq j$$

3. **Under-Segmentations:** Total number of predicted segments that have a significant overlap with more than one ground-truth segments. That is, a

<sup>2</sup> <https://github.com/sohaib023/T-Truth>.

Original Sample Ground-Truth Non-Augmented Standard TabAug



**Fig. 6.** Sample outputs of all three approaches for a comparative analysis. Each row provides a separate test sample. The red boxes highlight the image region that is displayed for visualization of ground-truth and the predictions. Ground-truth is displayed as blue lines while predictions are displayed as green lines. (Color figure online)

predicted segment  $S_j$  is under-segmented if:

$$T < \frac{|G_i \cap S_j|}{G_i} < 1 - T \text{ and } T < \frac{|G_k \cap S_j|}{G_k} < 1 - T \text{ where; } k \neq i$$

**Table 1.** Results of evaluating the trained models on 20% of ICDAR 2013 dataset reserved for testing.

		Non-augmented	Standard	TabAug
Row	Correct (%)	96.44 $\pm$ 1.13	97.86 $\pm$ 0.80	<b>97.86 <math>\pm</math> 0.35</b>
	Over-segmented (%)	2.64 $\pm$ 0.54	1.71 $\pm$ 0.46	<b>1.71 <math>\pm</math> 0.35</b>
	Under-segmented (%)	0.71 $\pm$ 0.56	0.21 $\pm$ 0.18	<b>0.21 <math>\pm</math> 0.18</b>
Column	Correct (%)	92.12 $\pm$ 1.11	86.38 $\pm$ 1.54	<b>94.44 <math>\pm</math> 0.25</b>
	Over-segmented (%)	4.84 $\pm$ 0.76	5.20 $\pm$ 1.11	<b>3.77 <math>\pm</math> 0.76</b>
	Under-segmented (%)	1.43 $\pm$ 0.25	3.95 $\pm$ 1.01	<b>0.90 <math>\pm</math> 0.25</b>
Cell	Correct (%)	92.16 $\pm$ 3.84	82.12 $\pm$ 6.76	<b>96.11 <math>\pm</math> 1.61</b>
	Over-segmented (%)	1.58 $\pm$ 0.25	1.13 $\pm$ 0.28	<b>1.02 <math>\pm</math> 0.37</b>
	Under-segmented (%)	3.37 $\pm$ 1.99	6.90 $\pm$ 2.21	<b>1.46 <math>\pm</math> 0.72</b>

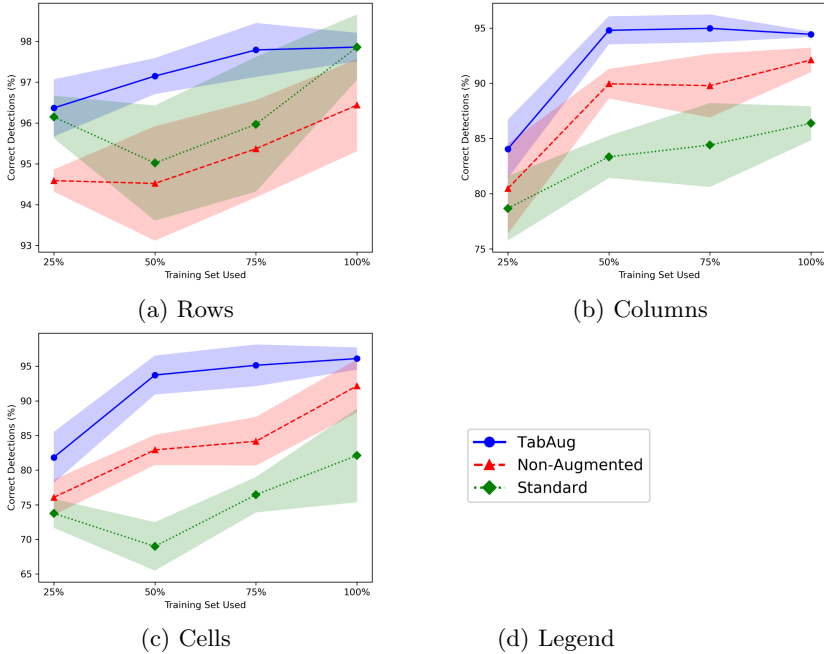
All of these evaluation metrics are normalized by division with the total number of ground-truth segments, which are presented as percentages in the results. We intentionally leave out the Partial Detections, Missed Segments, and False Positive Detections from our evaluation metrics, as due to the problem formulation of Split Model [23] they always evaluate to zero.

### 4.3 Results and Analysis

We evaluate all the trained models on a consistent 20% test split of ICDAR 2013 dataset. An evaluation of the performance measures described in Sect. 4.2 is provided in Table 1. Our proposed approach outperforms all other approaches in all performance measures for each of the Cell, Row, and Column recognition. Further, the models trained with the proposed approach show less standard deviation in the evaluation metrics, which is indicative of more stability in training. Moreover, as depicted in Table 2 and Fig. 7 our proposed approach consistently shows improvements across a range of training dataset sizes.

The standard augmentation approach manages to improve upon the evaluation metrics for Row, however, happens to be deficient in Cell and Column evaluations. This is explained by the fact that Row segmentation, especially in ICDAR 2013 dataset, is an easier task with less visual variability and complexity. Hence, it can be learned mostly based on white-spaces, for which image transformations provide an ample amount of augmentation. Column segmentation contains many complex scenarios and requires learning of more abstract and structural features, for which image transformations turn out to be insufficient.

We have further shown a comparison of the predictions generated by the models trained on each of the three approaches in Fig. 6. In Fig. 6a, we see the case where both standard augmentation and TabAug approaches successfully segment the rows, however, the non-augmented approach under-segments a row



**Fig. 7.** Models were trained on a varying percentage of the total training dataset and then evaluated on a consistent 20% test split. The graphs show the progression of correct detection with an increase in training data for all three approaches.

in the header region. Conversely, in Fig. 6b, we see a case of over-segmentation by the non-augmented approach. In Fig. 6c, non-augmented and standard augmentation approaches fail to recognize the boundaries of the columns correctly, due to the white-spaces that exist between the protruding words of the header. However, TabAug robustly recognizes the boundaries of the column and predicts a single column separator. In Fig. 6d, we see a scenario where the TabAug approach fails, as it confuses vertically consistent breaks in the text as a column breakage. Regardless, we see that it still does a decent job at making logically sound column segmentations as compared to segmentations from the other two approaches. Figure 6e depicts another sample that is correctly segmented by TabAug, however causes significant confusion in column segmentation for the other two approaches. Finally, in Fig. 6g the first column is over-segmented by all of the approaches, as the column header has no overlap with its content below. This demonstrates a hard sample that requires a higher cognitive understanding for correct prediction.

**Table 2.** Correct detection percentages achieved by models trained on different percentages of the training dataset evaluated on a consistent 20% test split of ICDAR 2013 dataset.

	Training data used	Correct detections (%)		
		Non-augmented	Standard	TabAug
Row	25%	94.59 ± 0.27	96.15 ± 0.52	<b>96.37 ± 0.70</b>
	50%	94.52 ± 1.40	95.02 ± 1.41	<b>97.15 ± 0.44</b>
	75%	95.37 ± 1.19	94.73 ± 3.38	<b>97.79 ± 0.66</b>
	100%	96.44 ± 1.13	97.86 ± 0.80	<b>97.86 ± 0.35</b>
Column	25%	80.46 ± 4.08	78.67 ± 2.92	<b>84.05 ± 2.65</b>
	50%	89.96 ± 1.34	83.34 ± 1.91	<b>94.80 ± 1.27</b>
	75%	89.79 ± 2.88	84.41 ± 3.80	<b>94.98 ± 1.26</b>
	100%	92.12 ± 1.11	86.38 ± 1.54	<b>94.44 ± 0.25</b>
Cell	25%	76.07 ± 2.56	73.76 ± 2.08	<b>81.84 ± 3.68</b>
	50%	82.92 ± 2.17	68.99 ± 3.50	<b>93.73 ± 2.81</b>
	75%	84.17 ± 3.51	76.45 ± 2.55	<b>95.14 ± 3.01</b>
	100%	92.16 ± 3.84	82.12 ± 6.76	<b>96.11 ± 1.61</b>

## 5 Conclusion

In this paper, we presented TabAug, a novel augmentation technique capable of producing structural changes in a table through replication and deletion of rows and columns. A data-driven probabilistic model is used in conjunction with the augmentation technique to control the augmentation process. Following the promising results of Split-model [23] trained on the publicly available ICDAR 2013 dataset using TabAug, we believe our work provides a strong foundation for numerous future extensions. In future, we plan to explore ideas for cross-table augmentation through statistical layout matching.

**Acknowledgement.** This work has been partially funded by the Higher Education Commission of Pakistan’s grant for National Center of Artificial Intelligence (NCAI).

## References

1. Arif, S., Shafait, F.: Table detection in document images using foreground and background features. *Digital Image Comput. Tech. Appl.* **2018**, 1–8 (2018)
2. Bansal, A., Harit, G., Dutta Roy, S.: Table extraction from document images using fixed point model. In: *ICVGIP 2014: Proceedings of the 2014 Indian Conference on Computer Vision Graphics and Image Processing*, pp. 1–8 (2014)
3. Chen, J., Lopresti, D.: Table detection in noisy off-line handwritten documents. In: *2011 International Conference on Document Analysis and Recognition, Beijing, China*, pp. 399–403 (2011)

4. Dwibedi, D., Misra, I., Hebert, M.: Cut, paste and learn: surprisingly easy synthesis for instance detection. In: 2017 IEEE International Conference on Computer Vision (ICCV), pp. 1310–1319 (2017)
5. Fang, H., Sun, J., Wang, R., Gou, M., Li, Y., Lu, C.: InstaBoost: boosting instance segmentation via probability map guided copy-pasting. In: 2019 IEEE/CVF International Conference on Computer Vision (ICCV), pp. 682–691 (2019)
6. Gatos, B., Danatsas, D., Pratikakis, I., Perantonis, S.J.: Automatic table detection in document images. In: Singh, S., Singh, M., Apte, C., Perner, P. (eds.) ICAPR 2005. LNCS, vol. 3686, pp. 609–618. Springer, Heidelberg (2005). <https://doi.org/10.1007/11551188.67>
7. Ghiasi, G., et al.: Simple copy-paste is a strong data augmentation method for instance segmentation. ArXiv (2020)
8. Gilani, A., Qasim, S.R., Malik, I., Shafait, F.: Table detection using deep learning. In: 14th International Conference on Document Analysis and Recognition, pp. 771–776 (2017)
9. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 770–778 (2016)
10. Kasar, T., Barlas, P., Adam, S., Chatelain, C., Paquet, T.: Learning to detect tables in scanned document images using line information. In: Twelfth International Conference on Document Analysis and Recognition, pp. 1185–1189 (2013)
11. Kieninger, T., Dengel, A.: A paper-to-HTML table converting system. In: Proceedings of Document Analysis Systems, pp. 356–365 (1998)
12. Kieninger, T., Dengel, A.: Table recognition and labeling using intrinsic layout features. In: International Conference on Advances in Pattern Recognition, pp. 307–316 (1999)
13. Kieninger, T., Dengel, A.: Applying the T-Recs table recognition system to the business letter domain. In: International Conference on Document Analysis and Recognition, p. 0518 (2001)
14. Krizhevsky, A., Sutskever, I., Hinton, G.E.: ImageNet classification with deep convolutional neural networks. In: Pereira, F., Burges, C.J.C., Bottou, L., Weinberger, K.Q. (eds.) Advances in Neural Information Processing Systems, vol. 25 (2012)
15. Pyreddy, P., Croft, W.B.: TINTI: a system for retrieval in text tables TITLE2: Technical report, University of Massachusetts, USA (1997)
16. Qasim, S.R., Mahmood, H., Shafait, F.: Rethinking table recognition using graph neural networks. In: 2019 International Conference on Document Analysis and Recognition (ICDAR), pp. 142–147 (2019)
17. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: towards real-time object detection with region proposal networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **39**, 1137–1149 (2015)
18. Schreiber, S., Agne, S., Wolf, I., Dengel, A., Ahmed, S.: DeepDeSRT: deep learning for detection and structure recognition of tables in document images. In: Fourteenth International Conference on Document Analysis and Recognition, vol. 1, pp. 1162–1167 (2017)
19. Shafait, F., Smith, R.: Table detection in heterogeneous documents. In: Proceedings of the 9th IAPR International Workshop on Document Analysis Systems, pp. 65–72. Document analysis systems (2010)
20. Shahab, A., Shafait, F., Kieninger, T., Dengel, A.: An open approach towards the benchmarking of table structure recognition systems. In: Document Analysis Systems, pp. 113–120 (2010)



21. Shelhamer, E., Long, J., Darrell, T.: Fully convolutional networks for semantic segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* **39**(4), 640–651 (2017)
22. Siddiqui, S., Malik, M., Agne, S., Dengel, A., Ahmed, S.: DeCNT: deep deformable CNN for table detection. *IEEE Access* **6**, 74151–74161 (2018)
23. Tensmeyer, C., Morariu, V.I., Price, B., Cohen, S., Martinez, T.: Deep splitting and merging for table structure decomposition. In: 2019 International Conference on Document Analysis and Recognition (ICDAR), pp. 114–121 (2019)
24. Tupaj, S., Shi, Z., Chang, D.H.: Extracting tabular information from text files. In: EECS Department, Tufts University (1996)
25. Zanibbi, R., Blostein, D., Cordy, J.: A survey of table recognition. *IJDAR* **7**, 1–16 (2004)