arXiv:2110.00061v1 [cs.LG] 30 Sep 2021

Scientific evidence extraction

Brandon Smock **Rohith Pesala** **Robin Abraham**
Microsoft
Redmond, WA
brsmock,ropesala,robin.abraham@microsoft.com

Abstract

Recently, interest has grown in applying machine learning to the problem of table structure inference and extraction from unstructured documents. However, progress in this area has been challenging both to make and to measure, due to several issues that arise in training and evaluating models from labeled data. This includes challenges as fundamental as the lack of a single definitive ground truth output for each input sample and the lack of an ideal metric for measuring partial correctness for this task. To address these we propose a new dataset, PubMed Tables One Million (PubTables-1M), and a new class of metric, *grid table similarity* (GriTS). PubTables-1M is nearly twice as large as the previous largest comparable dataset, can be used for models across multiple architectures and modalities, and addresses issues such as ambiguity and lack of consistency in the annotations. We apply DETR [1] to table extraction for the first time and show that object detection models trained on PubTables-1M produce excellent results out-of-the-box for all three tasks of detection, structure recognition, and functional analysis. We describe the dataset in detail to enable others to build on our work and combine this data with other datasets for these and related tasks. It is our hope that PubTables-1M and the proposed metrics can further progress in this area by creating a benchmark suitable for training and evaluating a wide variety of models for table extraction. Data and code will be released at https://github.com/microsoft/table-transformer.

1 Introduction

A table is a compact, structured representation for storing data and communicating it in documents and other manners of presentation, such as PDF or images. In its presented form, however, a table may not and often does not explicitly represent its logical structure. This is an important problem as a significant amount of data is communicated through documents, but without structure information this data is unable to be used in further applications.

The problem of inferring a table's structure from its presentation and converting it into a structured form is called table extraction (TE). TE is challenging for automated systems [2–5] and even for people [6] due to the wide variety of formats, styles, and structures found in presented tables. Recently, there has been a shift in the research literature from traditional rule-based methods [7–9] for TE to data-driven methods based on deep learning (DL) [2, 10, 11]. The primary advantage of DL methods is that they can learn to be more robust to the wide variety of table presentation formats. However, these methods require a large amount of data and still rely significantly on additional rules, hand-engineered components, or special training procedures to achieve good performance.

Recent datasets for table structure recognition (TSR) [4, 3, 11], while large, have several limitations, including in some cases lack of location information for cells, compatibility with only specific model architectures, and lack of guarantees for data quality and consistency. A more fundamental issue is that for a given input table, there may not be only one way to annotate its structure [6]. Yet these

Preprint. Under review.

datasets have been used for model training and evaluation as if each annotation is the only correct output, which leads to inconsistent feedback during training and noise during evaluation.

Another challenge for model evaluation in this area is the lack of an ideal metric for partial correctness. Several metrics have been proposed for evaluating the performance of TSR methods [12, 3, 13, 4]. While it is useful to evaluate TSR from multiple perspectives, these metrics lack a theoretical motivation, evaluate tables in ways that do not preserve their topological structure, and have different forms that lack an obvious connection between them, making them difficult to compare.

To address these and other issues, we introduce a new dataset, PubMed Tables One Million (PubTables-1M), and a new class of evaluation metric for TSR, *grid table similarity* (GriTS).

- PubTables-1M is the largest dataset of its kind, containing nearly 1M tables from the PubMed Central Open Access[1] (PMCOA) database. It is nearly twice as large as the current largest similar dataset and nearly nine times as large as the most comparable dataset. It contains both PDF and image bounding box annotations for table detection (TD), TSR, and functional analysis (FA), useful for any model whose data can be derived from PDF documents.
- PubTables-1M is the first attempt to create a TE dataset with unambiguous ground truth, making it more suitable than previous datasets for benchmarking progress. We introduce a *canonicalization* procedure whose goal is to ensure each table has an unambiguous structure interpretation. We also process and filter the data to ensure it has consistent, high-quality annotations for table content.
- Unlike previous metrics, GriTS evaluates a table in its natural matrix form. It also can evaluate multiple aspects of TSR within the same formulation, making it easier to make comparisons across different forms of output.
- We apply the Detection Transformer (DETR) [1] for the first time to the tasks of TD, TSR, and FA, and demonstrate how with our data all three tasks can be addressed within an object detection framework out-of-the-box without any custom components or training procedures.
- Finally, we plan to release all data and code for training and evaluation.

2 Background

Wang [14] distinguishes between a table in three forms, which we summarize here as:

1. Abstract table: a data structure that represents information in terms of a set of values, uniquely indexed by a multi-dimensional hierarchical system of keys.
2. Grid table: an abstract table with a two-dimensional arrangement of keys and values into cells occupying ordered rows and columns.
3. Presentation table: a visualization of a grid table with typography, spacing, and style.

A grid table is composed of cells, with each cell containing content. Each intersection of a row and a column forms a *grid cell*. A cell that spans multiple rows or columns is called a *spanning cell*, and its content is considered to be repeated at each grid cell location that it spans.

Generally, TE is considered the problem of inferring a table's grid form from its presentation. TE can be broken into three subproblems [15]: *table detection* (TD), which locates the table; *table structure recognition* (TSR), which recognizes the structure of a table in terms of rows, columns, and cells; and *functional analysis* (FA), which recognizes the keys and values of the table. In this paper we address all three subproblems, but give particular attention to training and evaluating methods for TSR.

The output of a TSR model can be evaluated from three perspectives: *cell topology recognition*, which considers just the layout of the cells in a grid; *cell content recognition*, which considers both cell topology and the text content of each cell; and *cell location recognition*, which considers both cell topology and the absolute coordinates of each cell within a document. For evaluation, all three perspectives are useful. Cell content recognition is most aligned with the end goal of TE but for PDF and image input it can be dependent on the quality of OCR. Cell location recognition does not depend on OCR, but not every TSR method reports cell locations. Cell topology recognition is free of OCR

[1] https://www.ncbi.nlm.nih.gov/pmc/tools/openftlist/

Table 1: Comparison of recent large datasets for table structure recognition.

Name	Format	# Tables	Cell Topology	Cell Content	Cell Location	Row and Column Location	Canonical Ground Truth	Content Consistency Verification
TableBank[4]	Image	145k	✓					
SciTSR[16]	Image	15k	✓	✓				
PubTabNet[3, 11]	Image	568k	✓	✓	✓†			
FinTabNet[11]	Image, PDF	113k	✓	✓	✓†			
PubTables-1M (ours)	**Image, PDF**	**948k**	✓	✓	✓	✓	✓	✓

† Cell bounding boxes are given for non-blank cells only and exclude any non-text portion of a cell.

and is applicable to all TSR methods, but is not anchored to the actual content of the cells either by text content or location. Thus, a high score on a cell topology metric would be necessary but not sufficient for performing well at TE.

3 Related Work

Datasets Several large datasets have been introduced recently for TE [17, 18, 4, 3, 11]. To attain a large number of labeled examples, these datasets are typically crowd-sourced from thousands of authors. PubTabNet, for example, is created using an automated alignment procedure [18] to match the same table in unaligned pairs of PDF and XML versions of the same scientific articles from the PMCOA database. Table 1 shows a comparison of large datasets for TSR. Among previous datasets, PubTabNet is the largest with 568k tables, although no test set has been released for benchmarking. In terms of usability, FinTabNet is the most widely applicable, as it annotates source PDF documents rather than rendered images; and both FinTabNet and the updated version of PubTabNet [11] have the most complete annotations, as they both contain location information for cells.

However, both FinTabNet and PubTabNet are missing bounding boxes for rows, columns, and blank cells; and in these datasets a cell's bounding box covers only its text portion, which ignores the role of the non-text portion of the cell. A more fundamental issue not addressed by these datasets is the potential for noisy, ambiguous ground truth. For instance, the original structure annotations from crowd-sourced datasets are often labeled inconsistently with respect to each other, due to a problem we refer to as *oversegmentation* (see Section 4). Further, errors or other inconsistencies in the original text content annotations may exist, and these remain in the final annotations due to the lack of a verification step to guarantee that all cell content annotations match the corresponding text in the PDF document. These issues affect the reliability of both model training and performance evaluation and limit the kinds of approaches that can be applied to TSR using this data.

Metrics The most straightforward metric for TSR is accuracy, which is the percentage of predicted tables where all predicted cells match exactly with the ground truth. However, this metric can be sensitive to things such as label noise and ambiguity, and it gives no credit for table predictions that differ slightly from the ground truth. As a result, a number of metrics have been proposed for measuring *partial* correctness for table predictions.

Göbel et al. [12] propose a content metric based on precision and recall for all pairs of adjacent cell content. Li et al. [4] propose a topology metric that evaluates HTML output using the 4-gram BLEU score. Zhong et al. [3] propose a content metric that is a modified tree-edit distance (TEDS) on a custom HTML tagset with a text content score. Gao et al. [13] propose a location version of the metric proposed by Göbel et al. [12], which evaluates precision and recall for pairs of adjacent cells whose intersection-over-union (IoU) with a ground truth cell is above different thresholds.

However, one problem with comparing tables as HTML sequences or trees is that a seemingly small difference such as a missing tag or difference in a cell's row span may produce significant cascading effects on a table's overall structure. Similarly, model output in forms such as HTML is not guaranteed to be able to be converted unambiguously into a two-dimensional grid table—for instance if different rows in the HTML sequence have different numbers of columns—yet may be penalized only a small amount for this depending on the choice of metric. Thus the main problem with all of these metrics is that there may be a mismatch between the scale of the differences as measured in these alternate table representations and their effects as measured within a table's full two-dimensional structure. The metrics proposed by Zhong et al. [3] and Li et al. [4] also do not

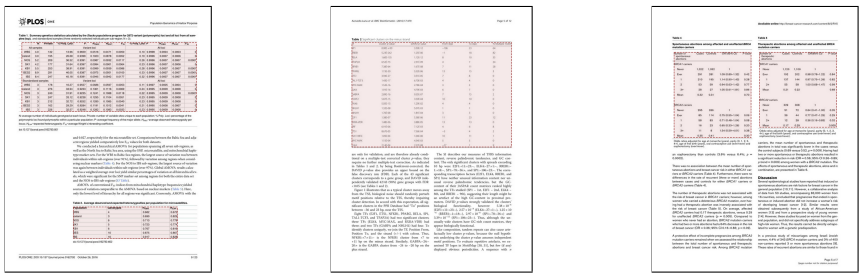

Figure 1: Examples of pages with table bounding box annotations (shaded red) in PubTables-1M.

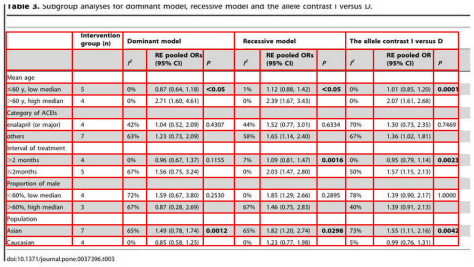

Table 3. Subgroup analyses for dominant model, recessive model and the allele contrast I versus D.

	Intervention group (n)	Dominant model			Recessive model			The allele contrast I versus D		
		I^2	RE pooled ORs (95% CI)	P	I^2	RE pooled ORs (95% CI)	P	I^2	RE pooled OR (95% CI)	P
Mean age										
≤60 y, low median	5	0%	0.87 (0.64, 1.18)	**<0.05**	1%	1.12 (0.88, 1.42)	**<0.05**	0%	1.01 (0.85, 1.20)	**0.0001**
>60 y, high median	4	0%	2.71 (1.60, 4.61)		0%	2.39 (1.67, 3.43)		0%	2.07 (1.61, 2.68)	
Category of ACEIs										
enalapril (or major)	4	42%	1.04 (0.52, 2.09)	0.4307	44%	1.52 (0.77, 3.01)	0.6334	70%	1.30 (0.73, 2.35)	0.7469
others	7	63%	1.23 (0.73, 2.09)		58%	1.65 (1.14, 2.40)		67%	1.36 (1.02, 1.81)	
Interval of treatment										
>2 months	4	0%	0.96 (0.67, 1.37)	0.1155	7%	1.09 (0.81, 1.47)	**0.0016**	0%	0.95 (0.79, 1.14)	**0.0023**
≤2months	5	67%	1.56 (0.75, 3.24)		0%	2.03 (1.47, 2.80)		50%	1.57 (1.15, 2.13)	
Proportion of male										
<60%, low median	4	72%	1.59 (0.67, 3.80)	0.2530	0%	1.85 (1.29, 2.66)	0.2895	78%	1.39 (0.90, 2.17)	1.0000
≥60%, high median	3	67%	0.87 (0.28, 2.69)		67%	1.46 (0.75, 2.83)		40%	1.39 (0.91, 2.13)	
Population										
Asian	7	65%	1.49 (0.78, 1.74)	**0.0012**	65%	1.82 (1.20, 2.74)	**0.0298**	73%	1.55 (1.11, 2.16)	**0.0042**
Caucasian	4	0%	0.85 (0.58, 1.25)		0%	1.23 (0.77, 1.98)		5%	0.99 (0.76, 1.31)	

doi:10.1371/journal.pone.0037396.t003

(a) Original structure annotation

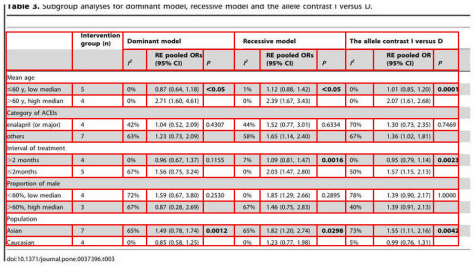

Table 3. Subgroup analyses for dominant model, recessive model and the allele contrast I versus D.

	Intervention group (n)	Dominant model			Recessive model			The allele contrast I versus D		
		I^2	RE pooled ORs (95% CI)	P	I^2	RE pooled ORs (95% CI)	P	I^2	RE pooled OR (95% CI)	P
Mean age										
≤60 y, low median	5	0%	0.87 (0.64, 1.18)	**<0.05**	1%	1.12 (0.88, 1.42)	**<0.05**	0%	1.01 (0.85, 1.20)	**0.0001**
>60 y, high median	4	0%	2.71 (1.60, 4.61)		0%	2.39 (1.67, 3.43)		0%	2.07 (1.61, 2.68)	
Category of ACEIs										
enalapril (or major)	4	42%	1.04 (0.52, 2.09)	0.4307	44%	1.52 (0.77, 3.01)	0.6334	70%	1.30 (0.73, 2.35)	0.7469
others	7	63%	1.23 (0.73, 2.09)		58%	1.65 (1.14, 2.40)		67%	1.36 (1.02, 1.81)	
Interval of treatment										
>2 months	4	0%	0.96 (0.67, 1.37)	0.1155	7%	1.09 (0.81, 1.47)	**0.0016**	0%	0.95 (0.79, 1.14)	**0.0023**
≤2months	5	67%	1.56 (0.75, 3.24)		0%	2.03 (1.47, 2.80)		50%	1.57 (1.15, 2.13)	
Proportion of male										
<60%, low median	4	72%	1.59 (0.67, 3.80)	0.2530	0%	1.85 (1.29, 2.66)	0.2895	78%	1.39 (0.90, 2.17)	1.0000
≥60%, high median	3	67%	0.87 (0.28, 2.69)		67%	1.46 (0.75, 2.83)		40%	1.39 (0.91, 2.13)	
Population										
Asian	7	65%	1.49 (0.78, 1.74)	**0.0012**	65%	1.82 (1.20, 2.74)	**0.0298**	73%	1.55 (1.11, 2.16)	**0.0042**
Caucasian	4	0%	0.85 (0.58, 1.25)		0%	1.23 (0.77, 1.98)		5%	0.99 (0.76, 1.31)	

doi:10.1371/journal.pone.0037396.t003

(b) After canonicalization

Figure 2: The same table annotations before and after canonicalization. Canonicalization merges the oversegmented cells in rows 3, 6, 9, 12, and 15, and the cells at the top of columns 1 and 2.

isolate TSR from functional analysis, taking into account aspects of both in their evaluation. These issues motivate us in Section 6 to propose new metrics for TSR with a clearer theoretical grounding that preserve a table's true topological structure and are natural to use in combination.

4 PubTables-1M Dataset

The source data for creating PubTables-1M are pairs of PDF and XML versions of the same document from the PMCOA dataset. Roughly the same text appears in both, but the text in the PDF has spatial location $[x_{\text{min}}, y_{\text{min}}, x_{\text{max}}, y_{\text{max}}]$, while the text in the XML appears inside semantically labeled tags. We use the Needleman-Wunsch algorithm [19] to align the text from both sources, connecting each XML tag to its spatial location.

Canonicalization To remedy the issue of inconsistency and ambiguity in the annotations, we propose to convert each table annotation into a *canonical* form. This is similar to that defined by Seth et al. [20], who describe a set of permissible tilings of a table into cells. However, ours is motivated from the goal of ensuring each presentation table has a *unique interpretation*, which is a way of favoring one particular segmentation of a table into rows, columns, and cells over other possibilities.

To do this, canonicalization assumes that the row and column headers in a table correspond in their abstract representation to trees. For an interpretation of the headers in a table to be unambiguous, there should be a one-to-one correspondence between its header cells and logical tree nodes. Tables in the PMCOA dataset, however, are often annotated with header cells that are *oversegmented*, meaning spanning cells are split into their individual grid cells (see Figure 2). Permitting such annotations in the ground truth leads to ambiguity in the TSR task, as annotators may differ in their choices about when to split certain cells, creating multiple possible ground truths. Canonicalization is a procedure that attempts to eliminate this ambiguity by identifying and consolidating oversegmented header cells into a one-to-one correspondence with their abstract tree nodes. For details, please see the Appendix.

Quality control To create a large dataset, PubTables-1M is drawn from crowd-sourced data that is repurposed for the task of TE and uses a potentially error-prone automated alignment procedure. Therefore, it is crucial to include a step that filters out potential errors from either of these steps and provides a measurable guarantee of quality. First, we discard tables with rows that overlap or

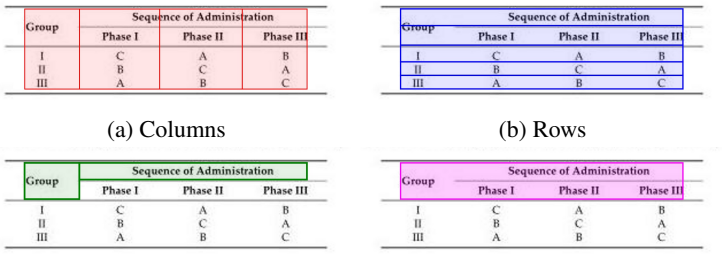

(a) Columns (b) Rows

(c) Spanning cells (d) Column header

Figure 3: An example table with dilated bounding box annotations for different object classes.

with columns that overlap, which we found is likely due to alignment mistakes. Second, we remove outliers by counting the number of objects in a table (defined in Section 5) and removing tables with more than 100. In all, less than 0.1% of tables are discarded as outliers.

Finally to filter out mistakes made both by the original annotators and our automated processing, we compare the edit distance between the non-whitespace text for every cell in the original XML annotations with the text extracted from the PDF inside the cell's bounding box. We filter out any tables for which the normalized edit distance between these averaged over every cell is above 0.05. We do not force the text from each to be *exactly* equal, as the PDF text can differ even when everything is annotated correctly, due to things like word wrapping, which may add hyphens that are not in the source annotations. When the annotations do slightly differ from their corresponding PDF text, we choose to consider the PDF text to be the ground truth. PubTables-1M is the first dataset that does this verification step (which previous datasets could not do because they do not associate a bounding box with every cell) and provides a measurable assurance of consistency for the ground truth.

Dataset splits and statistics In all we yield 947,642 annotated tables, of which 52.7% are complex. Prior to canonicalization, only 40.1% of the tables in the set were considered complex by the original annotators. In total, canonicalization adjusts the annotations in some way for 328,421 tables (34.7%), or 65.8% of the complex tables in the final set. We split the data randomly into train, validation, and test sets at the document level using an 80/10/10 split. For TSR, this results in 758,849 tables for training; 94,959 for validation; and 93,834 for testing. For each document, we note if all tables in the XML version of the document are present in the final set of annotations. While every table in the set can be used for training TSR models, only tables from documents with all of their tables annotated can be used for table detection. For TD, there are 460,589 fully-annotated pages containing tables for training; 57,591 for validation; and 57,125 for testing. Examples of pages with their table bounding box annotations are given in Figure 1. Note that tables that span multiple pages are considered outside the scope of this work.

5 Model

We model all three tasks of TD, TSR, and FA as object detection with images as input. For TD, we use two object classes: *table* and *table rotated*. The *table rotated* class corresponds to tables that are rotated counterclockwise 90 degrees, which is often the case for very wide tables.

TSR and FA model We use a novel approach that models TSR and FA jointly using six object classes: *table*, *table column*, *table row*, *table column header*, *table projected row header*, and *table spanning cell*. Four of these classes are illustrated in Figure 3. An example of a projected row header, also known as a *projected multi-level row header* [21] or a *section header* [22], can be seen in row 3 of the table in Figure 2. We are the first to explicitly model projected row headers for TE. The intersection of each pair of *table column* and *table row* objects can be considered to form a seventh implicit class, *table grid cell*. These objects model a table's hierarchical structure through physical overlap and model sequential ordering through their relative vertical and horizontal positioning.

Dilated bounding boxes For training the TSR and FA model, instead of using the original bounding boxes, which tightly surround each object, we instead use *dilated* bounding boxes. To create dilated

bounding boxes, for each pair of adjacent row bounding boxes and adjacent column bounding boxes, we expand their boundaries until they meet halfway, which fills the empty space in between them. Similarly we expand the objects from the other classes so their boundaries match the adjustments made to the rows and columns they occupy. After, there are no gaps or overlap between rows, between columns, or between cells.

DETR To demonstrate the proposed dataset and the object detection modeling approach, we apply for the first time the Detection Transformer (DETR) [1] to all three TE tasks. DETR has the advantage over other object detection approaches that it is capable of modeling global context for objects and does not perform an early-stage non-maximum suppression step that would prevent it from outputting objects from different classes with the same bounding box. We train one DETR model for TD and one DETR model for both TSR and FA. For comparison, we also train a Faster R-CNN [23] model for the same tasks. All models use a ResNet-18 backbone. Our goal is to avoid custom engineering the models and training procedure for each task, using default settings wherever possible and allowing primarily the data to drive the result. For more details on the architectures and training procedure, please see the Appendix.

6 Proposed Metrics

To address the weaknesses of prior evaluation metrics, we propose a new family of related metrics we refer to as *grid table similarity* (GriTS). Unlike previous metrics, GriTS evaluates the topological representation of a table as a two-dimensional grid, or matrix.

2D-LCS As a starting point for these metrics, we first consider the generalization of longest common subsequence to two dimensions, which is called two-dimensional longest common substructure (2D-LCS) [24]. Let $\mathbf{M}[R, C]$ be a matrix with $R = [r_1, \ldots, r_m]$ representing its rows and $C = [c_1, \ldots, c_n]$ representing its columns. 2D-LCS operates on two matrices, $\mathbf{A}$ and $\mathbf{B}$, and determines the largest two-dimensional substructure, $\tilde{\mathbf{M}}$, the two have in common. In other words, $\tilde{\mathbf{M}} = \mathbf{A}[R'_A, C'_A] = \mathbf{B}[R'_B, C'_B]$, where $R' \mid R$ is a subsequence of rows R, and $C' \mid C$ is a subsequence of columns C. Taking inspiration from the standard F-score, we can define a similarity measure based on this as $S(\mathbf{A}, \mathbf{B}) = \frac{2|\tilde{\mathbf{M}}|^{-1}}{|\mathbf{A}|^{-1}+|\mathbf{B}|^{-1}}$, where $|\mathbf{M}_{m\times n}| = m \cdot n$.

2D-MSS An extension to this is to relax the exact match constraint, and instead determine the two most *similar* two-dimensional substructures, $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{B}}$. We define this by replacing equality between entries $\mathbf{A}_{i,j}$ and $\mathbf{B}_{k,l}$ with some choice of similarity function between them $f(\mathbf{A}_{i,j}, \mathbf{B}_{k,l})$, which maps to the range $[0, 1]$. We call this two-dimensional most similar substructures (2D-MSS).

Grid table similarity (GriTS) GriTS is 2D-MSS with a particular choice of similarity function and a particular matrix of entries to compare. Given a similarity function $f()$ and choice of matrices $\mathbf{A}$ and $\mathbf{B}$ we define GriTS_f as:

$$\text{GriTS}_f(\mathbf{A}, \mathbf{B}) = \max_{R'_A, C'_A, R'_B, C'_B} \frac{2 \cdot \left(\sum_{i,j} f(\mathbf{A}[R'_A, C'_A]_{i,j}, \mathbf{B}[R'_B, C'_B]_{i,j})\right)^{-1}}{|\mathbf{A}|^{-1} + |\mathbf{B}|^{-1}}, \tag{1}$$

$$= \frac{2 \cdot \left(\sum_{i,j} f(\tilde{\mathbf{A}}_{i,j}, \tilde{\mathbf{B}}_{i,j})\right)^{-1}}{|\mathbf{A}|^{-1} + |\mathbf{B}|^{-1}}. \tag{2}$$

Letting $\mathbf{A}$ be the ground truth matrix and $\mathbf{B}$ be the predicted matrix, we can also define the following quantities, which we interpret as recall and precision: $\text{GriTS-Recall}_f(\mathbf{A}, \mathbf{B}) = \frac{\sum_{i,j} f(\tilde{\mathbf{A}}_{i,j}, \tilde{\mathbf{B}}_{i,j})}{|\mathbf{A}|}$ and $\text{GriTS-Precision}_f(\mathbf{A}, \mathbf{B}) = \frac{\sum_{i,j} f(\tilde{\mathbf{A}}_{i,j}, \tilde{\mathbf{B}}_{i,j})}{|\mathbf{B}|}$, where GriTS is then interpreted as the F-score.

One advantage of GriTS is we can use the same formulation for all aspects of TSR. We define one type of matrix and similarity for cell location recognition, $\text{GriTS}_{\text{Loc}}$, one for cell content recognition, $\text{GriTS}_{\text{Cont}}$, and one for cell topology recognition, $\text{GriTS}_{\text{Top}}$. The matrices used for each metric are visualized in Figure 4. For cell location, $\mathbf{A}_{i,j}$ contains the bounding box of the cell at position (i, j),

Group	Sequence of Administration	Sequence of Administration	Sequence of Administration
Group	Phase I	Phase II	Phase III
I	C	A	B
II	B	C	A
III	A	B	C

(a) $\text{GriTS}_{\text{Cont}}$

[0, 0, 1, 2]	[0, 0, 3, 1]	[-1, 0, 2, 1]	[-2, 0, 1, 1]
[0, -1, 1, 1]	[0, 0, 1, 1]	[0, 0, 1, 1]	[0, 0, 1, 1]
[0, 0, 1, 1]	[0, 0, 1, 1]	[0, 0, 1, 1]	[0, 0, 1, 1]
[0, 0, 1, 1]	[0, 0, 1, 1]	[0, 0, 1, 1]	[0, 0, 1, 1]
[0, 0, 1, 1]	[0, 0, 1, 1]	[0, 0, 1, 1]	[0, 0, 1, 1]

(b) $\text{GriTS}_{\text{Top}}$

[136.42, 477.25, 160.62, 501.45]	[185, 477.25, 470.89, 487.22]	[185, 477.25, 470.89, 487.22]	[185, 477.25, 470.89, 487.22]
[136.42, 477.25, 160.62, 501.45]	[185, 491.48, 271.9, 501.45]	[284.5, 491.48, 371.39, 501.45]	[384, 491.48, 470.89, 501.45]
[136.42, 505.82, 160.62, 515.72]	[185, 505.82, 271.9, 515.72]	[284.5, 505.82, 371.39, 515.72]	[384, 505.82, 470.89, 515.72]
[136.42, 515.73, 160.62, 525.63]	[185, 515.73, 271.9, 525.63]	[284.5, 515.73, 371.39, 525.63]	[384, 515.73, 470.89, 525.63]
[136.42, 525.64, 160.62, 535.53]	[185, 525.64, 271.9, 535.53]	[284.5, 525.64, 371.39, 535.53]	[384, 525.64, 470.89, 535.53]

(c) $\text{GriTS}_{\text{Loc}}$

Figure 4: Ground truth matrices for different GriTS metrics for the table in Figure 3. Each matrix entry corresponds to one grid cell. Entries that correspond to spanning cells are shaded darker.

and we use IoU to compute similarity between bounding boxes. For cell content, $\mathbf{A}_{i,j}$ contains the text content of the cell located at position (i, j), and we use normalized longest common subsequence (LCS) to compute similarity between text.

For cell topology, we use the same similarity function as cell location but on bounding boxes with size and relative position given in the grid coordinate system. For the cell at position (i, j), let $\alpha_{i,j}$ be its rowspan, let $\beta_{i,j}$ be its colspan, let $\rho_{i,j}$ be the minimum row it occupies, and let $\theta_{i,j}$ be the minimum column it occupies. Then for cell topology recognition, $\mathbf{A}_{i,j}$ contains the bounding box $[\theta_{i,j} - j, \rho_{i,j} - i, \theta_{i,j} - j + \beta_{i,j}, \rho_{i,j} - i + \alpha_{i,j}]$. Note that for any cell with rowspan of 1 and colspan of 1, this box is $[0, 0, 1, 1]$.

Factored 2D-MSS Computing the 2D-LCS of two matrices is NP-hard [24]. This suggests that all metrics for TSR may end up an approximation to what could be considered the ideal metric. We propose a heuristic approach to determine 2D-MSS by factoring the problem and determining the optimal 1D subsequences of rows and of columns independently from each matrix. This procedure uses dynamic programming (DP) in a nested manner, which is run twice: once to determine the most similar rows and once to determine the most similar columns between the two matrices. The nested DP procedure is $O(|\mathbf{A}| \cdot |\mathbf{B}|)$. Because the outcome of the procedure is a selection of rows and columns for each matrix, it still yields a valid 2D substructure of each—these just may not be the most similar substructures possible. It follows that the similarity computed by this procedure is a lower bound on the true similarity between $\mathbf{A}$ and $\mathbf{B}$.

7 Experiments

Metrics To validate that the proposed metrics are well-behaved and improve upon previous metrics, we perform experiments evaluating GriTS on the actual ground truth (GT) versus versions of the GT that are corrupted in straightforward ways. To produce a corrupted version of the GT, we select each row and each column from the actual GT with probability x, in their original order, discarding the rest. Note that every row-column pair (or, cell) is kept with probability x^2. For these experiments, we evaluate the recall, precision[2], and F-score of both GriTS and the adjacency metric proposed by Göbel et al. [12], which we refer to as Adj_{Cont}.

In the first experiment, we vary x from $[0, 1]$ in increments of 0.1. We report the resulting values of the metrics in Figure 5. As can be seen, because the relative ordering of the cells that remain in the corrupted GT is unchanged, all three GriTS metrics have a recall roughly equal to the probability of each cell being present, x^2, and a precision close to 1, which is intuitive and matches the intended behavior. On the other hand, removing rows and columns has a larger effect on Adj_{Cont} because it is based on adjacency relationships, which are significantly affected by this operation.

In the second experiment, we fix x to be 0.5 (x^2 to be 0.25) and instead vary how we sample the 50% of rows and columns. We test sampling rows and columns using three different schemes: 1) *first*, where we select the first 50% of both rows and columns, 2) *alternating*, where we select either

[2]For continuity, in the case of there being no positive predictions (empty output), we define precision to be 1.

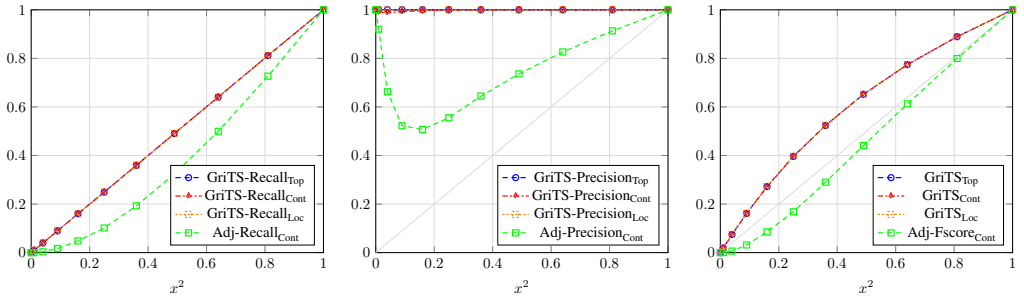

(a) Recall (b) Precision (c) F-score

Figure 5: Comparison of the proposed GriTS metrics versus Adj_{Cont} [12] when evaluating corrupted ground truth, where in the corrupted ground truth we keep each true row and column (in order) with probability x, discarding the rest.

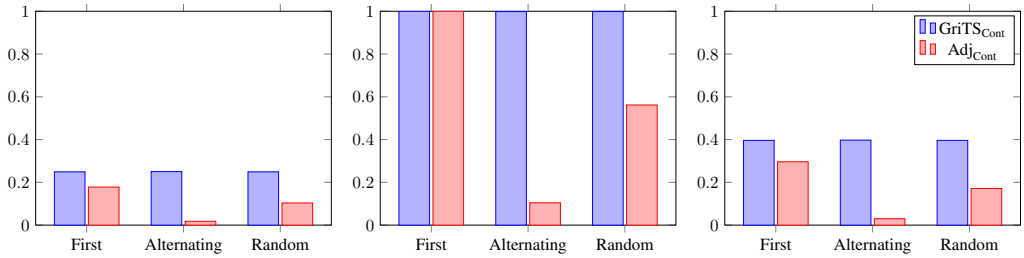

(a) Recall (b) Precision (c) F-score

Figure 6: Comparison of $\text{GriTS}_{\text{Cont}}$ versus Adj_{Cont} [12] when evaluating corrupted ground truth, where in the corrupted ground truth we keep each true row and column (in order) with probability 0.5, discarding the rest, and vary the scheme used to sample the rows and columns.

the odd or even-numbered rows and columns, and 3) *random*, where we randomly select 50% of the rows and columns, as in the first experiment. We show in Figure 6 the recall, precision, and F-score of $\text{GriTS}_{\text{Cont}}$ and Adj_{Cont} for each sampling scheme. As can be seen, for the same reason above, GriTS is essentially invariant to how we sample rows and columns, while Adj_{Cont} exhibits significantly different behavior based on which rows and columns are sampled, which is undesirable and contributes noise to the evaluation.

Model Evaluation In the next set of experiments, we train models on object detection data derived from PubTables-1M. All models reported in these experiments use a ResNet-18 backbone.

For TD, we train three models: DETR and Faster R-CNN trained on the full training set, and DETR trained on 10% of the data randomly sampled from the training set (DETR-100k). We report the results in Table 2. As can be seen, for TD DETR slightly outperforms Faster R-CNN on AP_{50} but significantly outperforms on AP. We interpret this to mean that while both models are able to learn to detect tables, DETR precisely localizes tables much better than Faster R-CNN. This also shows that the dataset is not trivial to learn, despite the strong performance of DETR. Also, while DETR is trained on 10x the data as DETR-100k, this results in only a slight gain in AP. This suggests that the large size of the full PubTables-1M dataset is not critical for achieving great performance on TD.

For TSR and FA, we train four models: DETR and Faster R-CNN on the full training set, DETR on 10% of the training data randomly sampled (DETR-100k), and DETR on the original, non-canonical (NC) annotations (DETR-NC). We report the results using object detection metrics for the models trained on canonical data in Table 2, which measures performance jointly on TSR and FA, and report results for all models using TSR-only metrics in Table 3. Table 3 includes the table content accuracy metric (Acc_{Cont}), which is the percentage of tables whose text content matches the ground truth exactly for every cell. We also break the results down by table category (simple, complex) and evaluate DETR-NC on both the original NC test data as well as the canonical test data.

Table 2: Test performance of models on PubTables-1M using object detection metrics.

Task	Model	AP	AP_{50}	AP_{75}	AR
TD	Faster R-CNN	0.825	0.985	0.927	0.866
	DETR-100k	0.957	**0.995**	**0.989**	0.977
	DETR	**0.966**	**0.995**	0.988	**0.981**
TSR + FA	Faster R-CNN	0.722	0.815	0.785	0.762
	DETR-100k	0.888	0.963	0.932	0.931
	DETR	**0.912**	**0.971**	**0.948**	**0.942**

Table 3: Test performance of the TSR + FA models on PubTables-1M on the proposed GriTS metrics.

Model	Table Category	Acc_{Cont}	$GriTS_{Top}$	$GriTS_{Cont}$	$GriTS_{Loc}$	$GriTS_{RawLoc}$ †	Adj_{Cont}
Non-Canonical Test Data							
DETR-NC	Simple	0.8678	0.9872	0.9859	0.9821	0.9753	0.9801
	Complex	0.5360	0.9600	0.9618	0.9444	0.9330	0.9505
	All	0.7336	0.9762	0.9761	0.9668	0.9582	0.9681
Canonical Test Data							
DETR-NC	Simple	0.9349	0.9933	0.9920	0.9900	0.9848	0.9865
	Complex	0.2712	0.9257	0.9290	0.9044	0.8933	0.9162
	All	0.5851	0.9576	0.9588	0.9449	0.9366	0.9494
DETR-100k	Simple	**0.9495**	0.9948	**0.9938**	0.9920	0.9897	0.9830
	Complex	0.6869	0.9726	0.9736	0.9620	0.9537	0.9629
	All	0.8111	0.9831	0.9832	0.9762	0.9707	0.9754
Faster R-CNN	Simple	0.0867	0.8682	0.8571	0.6869	0.6108	0.8024
	Complex	0.1193	0.8556	0.8507	0.7518	0.6988	0.7734
	All	0.1039	0.8616	0.8538	0.7211	0.6572	0.7871
DETR	Simple	0.9468	**0.9949**	**0.9938**	**0.9922**	**0.9900**	**0.9893**
	Complex	**0.6944**	**0.9752**	**0.9763**	**0.9654**	**0.9604**	**0.9667**
	All	**0.8138**	**0.9845**	**0.9846**	**0.9781**	**0.9744**	**0.9774**

† For the definition of $GriTS_{RawLoc}$, please see the Appendix.

As can be seen, DETR trained on the canonical data produces strong results for TSR and FA, outperforming the other models when evaluated on all tables. Comparing DETR-NC evaluated on NC ground truth versus DETR evaluated on canonical ground truth, we see that using canonical data improves performance on the metrics across all table types. This is even more apparent for the table accuracy metric, for which the use of canonical data is responsible for a jump in performance from 0.5360 to 0.6944 for complex tables.

To consider the positive impact that canonicalization has just on facilitating a more reliable evaluation, we compare DETR-NC evaluated on canonical data versus NC data. Even though it is trained on NC data, the metrics for simple tables are much higher when DETR-NC is evaluated on canonical data (0.9349 accuracy) than when it is evaluted on NC data (0.8678 accuracy). This shows even more clearly that the canonical data is less noisy and contributes to a more reliable evaluation.

Comparing DETR-NC evaluated on NC ground truth versus DETR-100k evaluated on canonical ground truth, we see even more the positive impact that canonicalization has, as we see an improvement in performance even with 10x less training data. This suggests that in this case having cleaner data is more important than having more data.

Overall the results confirm that canonical data significant improves performance for TSR models. Even setting aside the argument that NC data is noisier and more ambiguous, it is also a *different* way of annotating the data, using oversegmentation, which is less useful because it does not correspond to a table's true logical structure. This difference is apparent when we compare DETR-NC versus DETR, both evaluated on canonical test data. DETR-NC performs much worse on complex tables due in large part to the less precise scheme it learns for understanding the spanning cells in their headers.

8 Conclusion

In this paper we introduced a new dataset, PubMed Tables One Million (PubTables-1M), the largest of its kind, and *grid table similarity* (GriTS), a new class of evaluation metric for table structure recognition (TSR). We proposed a novel *canonicalization* procedure for TSR annotations and showed this leads to both more reliable evaluation and a significant improvement in model performance. GriTS evaluates model predictions in a table's natural matrix form, improving over previous metrics and making it more suitable for reliable evaluation. Finally, we cast table detection, TSR, and functional analysis within an object detection framework and trained DETR on them for the first time, demonstrating excellent performance is possible using our data with minimal customization for these tasks. While we do not believe this work raises any issues regarding negative impacts to society, we welcome a discussion on any potential impacts raised by others.

9 Future Work

We hope the proposed dataset, metrics, and methods will aid progress by making it easier and more meaningful to compare different methods for TE. While we demonstrated the importance of canonicalization to improving performance for TSR, we did so only for data from a single domain. We hope to apply canonicalization to data from additional domains, such as the financial documents in FinTabNet. Also, TE is often just one stage in larger pipelines for document understanding and information retrieval. Developing end-to-end systems in these areas is an important problem with its own challenges, including the composability of both models and datasets for these tasks. PMCOA is a rich resource for these systems, and we hope that releasing a large pool of high-quality annotated tables can further progress on end-to-end systems and ensures that others can create combined datasets large enough to facilitate a diverse range of these tasks.

Acknowledgments

We would like to thank Pramod Sharma, Natalia Larios Delgado, Joseph N. Wilson, Mandar Dixit, and John Corring for helpful discussions and feedback while preparing this manuscript.

References

[1] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European Conference on Computer Vision*, pages 213–229. Springer, 2020.

[2] Sebastian Schreiber, Stefan Agne, Ivo Wolf, Andreas Dengel, and Sheraz Ahmed. DeepDeSRT: Deep learning for detection and structure recognition of tables in document images. In *2017 14th IAPR international conference on document analysis and recognition (ICDAR)*, volume 1, pages 1162–1167. IEEE, 2017.

[3] Xu Zhong, Elaheh ShafieiBavani, and Antonio Jimeno Yepes. Image-based table recognition: data, model, and evaluation. *arXiv preprint arXiv:1911.10683*, 2019.

[4] Minghao Li, Lei Cui, Shaohan Huang, Furu Wei, Ming Zhou, and Zhoujun Li. Tablebank: Table benchmark for image-based table detection and recognition. In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 1918–1925, 2020.

[5] Shubham Singh Paliwal, D Vishwanath, Rohit Rahul, Monika Sharma, and Lovekesh Vig. Tablenet: Deep learning model for end-to-end table detection and tabular data extraction from scanned document images. In *2019 International Conference on Document Analysis and Recognition (ICDAR)*, pages 128–133. IEEE, 2019.

[6] Jianying Hu, Ramanujan Kashi, Daniel Lopresti, George Nagy, and Gordon Wilfong. Why table ground-truthing is hard. In *Proceedings of Sixth International Conference on Document Analysis and Recognition*, pages 129–133. IEEE, 2001.

[7] Wolfgang Gatterbauer, Paul Bohunsky, Marcus Herzog, Bernhard Krüpl, and Bernhard Pollak. Towards domain-independent information extraction from web tables. In *Proceedings of the 16th international conference on World Wide Web*, pages 71–80, 2007.

[8] Ermelinda Oro and Massimo Ruffolo. TREX: An approach for recognizing and extracting tables from PDF documents. In *2009 10th International Conference on Document Analysis and Recognition*, pages 906–910. IEEE, 2009.

[9] Alexey O Shigarov. Table understanding using a rule engine. *Expert Systems with Applications*, 42(2): 929–937, 2015.

[10] Devashish Prasad, Ayan Gadpal, Kshitij Kapadni, Manish Visave, and Kavita Sultanpure. CascadeTabNet: An approach for end to end table detection and structure recognition from image-based documents. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 572–573, 2020.

[11] Xinyi Zheng, Douglas Burdick, Lucian Popa, Xu Zhong, and Nancy Xin Ru Wang. Global table extractor (GTE): A framework for joint table identification and cell structure recognition using visual context. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 697–706, 2021.

[12] Max Göbel, Tamir Hassan, Ermelinda Oro, and Giorgio Orsi. A methodology for evaluating algorithms for table understanding in PDF documents. In *Proceedings of the 2012 ACM symposium on Document engineering*, pages 45–48, 2012.

[13] Liangcai Gao, Yilun Huang, Hervé Déjean, Jean-Luc Meunier, Qinqin Yan, Yu Fang, Florian Kleber, and Eva Lang. ICDAR 2019 competition on table detection and recognition (cTDaR). In *2019 International Conference on Document Analysis and Recognition (ICDAR)*, pages 1510–1515. IEEE, 2019.

[14] Xinxin Wang. Tabular abstraction, editing, and formatting, 1996.

[15] Max Göbel, Tamir Hassan, Ermelinda Oro, and Giorgio Orsi. ICDAR 2013 table competition. In *2013 12th International Conference on Document Analysis and Recognition*, pages 1449–1453. IEEE, 2013.

[16] Zewen Chi, Heyan Huang, Heng-Da Xu, Houjin Yu, Wanxuan Yin, and Xian-Ling Mao. Complicated table structure recognition. *arXiv preprint arXiv:1908.04729*, 2019.

[17] Noah Siegel, Nicholas Lourie, Russell Power, and Waleed Ammar. Extracting scientific figures with distantly supervised neural networks. In *Proceedings of the 18th ACM/IEEE on joint conference on digital libraries*, pages 223–232, 2018.

[18] Xu Zhong, Jianbin Tang, and Antonio Jimeno Yepes. Publaynet: largest dataset ever for document layout analysis. In *2019 International Conference on Document Analysis and Recognition (ICDAR)*, pages 1015–1022. IEEE, 2019.

[19] Saul B Needleman and Christian D Wunsch. A general method applicable to the search for similarities in the amino acid sequence of two proteins. *Journal of molecular biology*, 48(3):443–453, 1970.

[20] Sharad Seth, Ramana Jandhyala, Mukkai Krishnamoorthy, and George Nagy. Analysis and taxonomy of column header categories for web tables. In *Proceedings of the 9th IAPR International Workshop on Document Analysis Systems*, pages 81–88, 2010.

[21] Jianying Hu, Ramanujan S Kashi, Daniel P Lopresti, and Gordon Wilfong. Table structure recognition and its evaluation. In *Document Recognition and Retrieval VIII*, volume 4307, pages 44–55. International Society for Optics and Photonics, 2000.

[22] David Pinto, Andrew McCallum, Xing Wei, and W Bruce Croft. Table extraction using conditional random fields. In *Proceedings of the 26th annual international ACM SIGIR conference on Research and development in informaion retrieval*, pages 235–242, 2003.

[23] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. *arXiv preprint arXiv:1506.01497*, 2015.

[24] Amihood Amir, Tzvika Hartman, Oren Kapah, B Riva Shalom, and Dekel Tsur. Generalized LCS. *Theoretical computer science*, 409(3):438–449, 2008.

10 Appendix

Here we provide additional details about the dataset creation, training, and evaluation procedures.

10.1 Dataset

Header correction The canonicalization procedure operates on cells in the row and column headers. The source XML annotations, however, do not label row headers, and we found that they sometimes contain incomplete annotations of the column headers, as well. Before canonicalization, we use the assumption that the logical structure of the headers in their abstract representations is a tree to identify missing row header and incomplete column header annotations. Accurately labeling the full row header of a table for functional analysis is considered outside the scope of this paper. However, the high accuracy of our row header identification method is useful to correct oversegmented cells in the first column, leading to a significant net improvement in segmentation correctness for these cells. The method to add missing rows to the column header extends the header to more rows for 56,495 tables (6.0%).

Projected row headers Here we briefly call attention to a cell annotation label that is added to the PubTables-1M ground truth as a result of the canonicalization procedure. This type of cell is referred to as a *projected multi-level row header* [21], *section header* [22], or by us as simply a *projected row header*. An example of a projected row header can be seen in row 3 of the table in Figure 2. Projected row headers label logically distinct sections (or rows) of a table. These are a common source of ambiguity, as annotators differ on how to segment such rows into cells. As each projected row header corresponds to one node in the tree representation of the row header, during canonicalization we consolidate the entire row into a single spanning cell. For the tables in PMCOA, we consider this annotation of the spanning cell as part of the row header accurate enough to include as part of the canonicalized ground truth.

Algorithm 1 PubTables-1M Annotation Correction and Canonicalization Algorithm

PRE-PROCESS
  Split every blank (having no text content) spanning cell into blank grid cells
CORRECT/INFER THE COLUMN HEADER ANNOTATION
  if the first row starts with a blank cell **then** label the first row as being part of the column header
  if there is at least one row labeled as part of the column header **then**
    for each column **do** determine the first row in that column that has a cell that does not span multiple columns
    Determine the last row of the rows identified in the previous step
    Label every row in the table up to the row determined in the previous step to be a part of the column header
  end if
CANONICALIZE THE COLUMN HEADER CELLS
  for each cell in the column header **do** merge the cell with any neighboring cells (above or below) in the column header that span the exact same columns
  for each cell in the column header **do** merge the cell with blank cells below it in the column header: expand the rowspan of the cell to as many succeeding rows as possible such that it is only merged with blank cells that are also in the column header
  for each cell in the column header **do** merge the cell with blank cells above it: expand the rowspan of the cell to as many preceding rows as possible such that it is only merged with blank cells
INFER (PORTIONS OF) THE ROW HEADER
  for each row **do: if** the row is outside of the column header and has exactly one non-blank cell **then** label the row as a projected row header
  if any cell in the first column below the column header is a spanning cell or blank **then** label the first column as part of the row header
CANONICALIZE THE ROW HEADER CELLS
  for each projected row header **do** merge all of the cells in the row into a single cell
  for each cell in the row header (first column only) **do** merge the cell with blank cells below it: expand the rowspan of the cell to as many succeeding rows as possible such that it is only merged with blank cells

Canonicalization algorithm for PubTables-1M We describe the algorithm for correcting and canonicalizing the annotations in the PubTables-1M dataset in Algorithm 1. The goal of canonicalization is to ensure unique, unambiguous ground truth structure annotations for the table structure recognition task. The

principal action for accomplishing this is to eliminate oversegmentation of cells in the column and row headers. Oversegmentation occurs when a node in the header's logical tree representation corresponds to more than one cell in the grid representation of the table. This is undesirable if permitted in a table's grid structure ground truth labels because it introduces ambiguity in how to interpret and parse a table's grid structure from its presentation.

Even when we ensure there is exactly one cell annotation for every tree node, there is still the potential for ambiguity in where exactly the boundaries between these cells lie. This is because sometimes sections of whitespace that occupy entire grid cells in the header can be considered a part of a spanning cell on either side of the whitespace. How we associate such whitespace with surrounding cells can affect which grid locations the surrounding cells are considered to occupy, which is central to the task of table structure recognition. The canonicalization procedure, therefore, not only eliminates oversegmentation but also provides a consistent whitespace association scheme for unambiguously interpreting a table's grid structure. We choose a particular association scheme from among multiple possible schemes, which provides consistency in the annotations and helps ensure the table structure recognition task is well-posed.

While the intended outcome of canonicalization can be applied to any table structure annotation, we note that the canonicalization steps in Algorithm 1 are designed to achieve this specifically for the kinds of source annotations found in the PMCOA dataset. For example, included in the overall procedure but distinct from canonicalization itself, are correction steps to infer missing portions of the column and row headers. These steps are needed because in the source annotations from PMCOA the column header is not always fully annotated, and the row header is not annotated at all. Additionally, the step taken at Line 19 is motivated from the observation in the PMCOA dataset that oversegmentation of the row header occurs when cells have text with top vertical alignment but not middle vertical alignment or bottom vertical alignment. Therefore this step associates the whitespace with the nearest cell above the whitespace rather than below.

One final thing worth noting is that Algorithm 1, while necessary to ensure unambiguous ground truth, is limited in the kind of corrections it attempts to make. Because a table's presentation form is not changed by the canonicalization procedure, the procedure cannot fix other possible aspects of a table's underlying design unrelated to oversegmentation that could prevent it from having an unambiguous interpretation. Similarly, it does not attempt to detect or correct all possible annotation mistakes. Instead, we use a quality control procedure to detect and remove samples whose annotations are unreliable. While the concept of canonicalization is applicable to tables generally, additional work may be required to apply the particular procedure described here to tables from other datasets or to distill the rules of the procedure into a specification for designing unambiguous tables.

10.2 Training

For both DETR models, we use a ResNet-18 backbone, six layers in the encoder, and six layers in the decoder. For TD, we use 15 object queries, and for TSR and FA we use 125 object queries, each chosen to be slightly more than the maximum number of objects in each set's training samples. Besides this, we use the same default architecture settings for each model.

To create training data for the TD model, we render the PDF pages to images with a maximum length of 1000 pixels and appropriately scale the bounding boxes for the objects to image coordinates. For TSR and FA, we first render the page containing the table as an image with a maximum length of 1000 pixels, scale and pad the table bounding box with an additional 30 pixels on all sides (or fewer on a side if there are less than 30 pixels available on that side), and crop the image to this bounding box. The padding enables more variation in training through cropping augmentations.

We use no custom components, losses, or procedures for training the model, other than standard data augmentations, such as random cropping and resizing. We only add to the models a simple *conflict resolution* step used strictly at inference time, followed by a conversion step from the set of objects to a logical table. The conflict resolution step only involves removing objects or adjusting their bounding boxes to eliminate overlap between objects of the same class. For the sake of evaluation, we also align the bounding boxes to the text extracted from the document, though this action is taken after text extraction and has no effect on the outcome.

All of the experiments are performed using a single NVidia Tesla V100 GPU. We initialize the models with weights pre-trained on ImageNet and train each model for 20 epochs using all default hyperparameters and training settings except for those we note here. For both models, we use a learning rate drop of 1 and gamma of 0.9. For the TSR and FA model, we also use an initial learning rate of 0.00005 and a no-object class weight of 0.4. We limited hyperparameter tuning to one short experiment to determine the initial learning rate. We ran training experiments with three different initial learning rates of 0.0002, 0.0001, 0.00005 and chose to use the learning rate for each model that had the best performance on the validation set after one epoch of training.

10.3 Evaluation

$\mathbf{GriTS_{RawLoc}}$ To assess how well the DETR TSR model performs with no post-processing, we define a fourth metric, $\text{GriTS}_{\text{RawLoc}}$. $\text{GriTS}_{\text{RawLoc}}$ uses the same similarity function as $\text{GriTS}_{\text{Loc}}$ but the matrix of predicted cell

bounding boxes are the raw output of the model, which we compare to the true dilated bounding boxes. The difference between $\text{GriTS}_{\text{Loc}}$ and $\text{GriTS}_{\text{RawLoc}}$ mostly measures the impact of the conflict resolution stage on performance.

Note that for evaluating TSR models according to cell location recognition, we report the cell locations after the conflict resolution stage that, in addition to removing overlap between objects of the same class, also adjusts the row and column bounding boxes to tightly surround the bounding boxes for the words they contain.