Tables Decoded: DELTA for Structure, TARQA for Understanding
Supplementary Material

1. Symbols and abbreviations

Table (1) presents the terminology used in the proposed approaches, along with their full forms and the corresponding tasks in which they are applied. Table (2) lists all the abbreviations used throughout the paper, providing readers with an easy reference to understand the terminology better.

2. Motivation for Doubly Decoupled Approach

Table Reconstruction approaches can be divided into:

- **End-to-end VLMs** In conventional table understanding pipelines, physical structure, logical structure, and content recognition (OCR) are often bundled together into a monolithic framework. We refer to this as the VLMs paradigm, where all three aspects: cell boundaries, spanning relations, and textual extraction are jointly modeled. While such end-to-end systems simplify design, they are typically hard to interpret, debug, and adapt across domains and languages.
- **Conventional Decoupling:** Recent approaches such as CascadeTabNet [6], TATR [9], GTE [10], TableFormer [5], and LGPMA [7] attempt to decouple structure recognition from content recognition. However, within the structure recognition stage, the physical structure (rows, columns, cell boundaries) and the logical structure (spanning cells, header associations, merged cells) remain entangled. This coupling often limits flexibility and makes it challenging to localise errors, as mispredictions in geometry and semantics influence each other.
- **Our Doubly-Decoupled Framework: (DELTA)** We introduce a finer decomposition by separating physical and logical structure recognition into independent stages. In the structure recognition stage, the physical structure (cell boundaries) and the logical structure (cell arrangements) should be loosely coupled. The physical structure is inherently tied to image coordinates, whereas the logical structure only requires predicting the arrangement and relationships among cells. By decoupling the two, we can keep the logical layer language-agnostic, while isolating coordinate-dependent errors in the physical layer. This separation improves flexibility, simplifies debugging, and allows each component to be strengthened independently. Specifically, TATR models the physical layout, while SPRINT captures logical relations. Their outputs are then combined into a complete table structure, which is subsequently passed to OCR for content extraction. This double decoupling offers greater control and modularity.

3. DELTA in detail

Our proposed framework, DELTA, tackles the long-standing challenge of disentangling physical and logical structures in table structure recognition. Unlike traditional approaches, where geometric layout (rows, columns, cell boundaries) and logical layout (cell arrangement, spans, header associations) are tightly coupled, DELTA predicts them independently and then combines their outputs in a principled manner. Specifically, we employ SPRINT for logical structure prediction, which generates compact OTSL/HTML sequences and offers an ideal balance of speed, accuracy, and language independence, making it robust across multilingual and noisy documents. For the physical structure, we use TATR, a DETR-based model pre-trained on large-scale datasets, that leverages only row and column predictions to ensure clean grid alignment. These two components are integrated to reconstruct complete HTML tables with explicit bounding boxes, row spans, and column spans, making the system modular, interpretable, and extensible. Beyond structure prediction, we also examine the role of OCR for content recognition and conduct comprehensive ablations to validate and justify our design choices.

3.1. SPRINT for Logical Structure

SPRINT is an image-to-sequence model that employs a Global Context Attention (GCA)-based encoder and a transformer-based decoder to generate compact sequence representations (OTSL/HTML) of logical table structures. We adopt SPRINT for logical structure prediction because it achieves an ideal balance of speed, accuracy, and language independence. Unlike methods that rely heavily on OCR or language-specific cues, SPRINT focuses solely on the structural layout, making it inherently robust across multilingual and noisy document settings. This design aligns perfectly with our objective of decoupling physical and logical structure, as SPRINT cleanly predicts the cell arrangement without being confounded by text semantics. Moreover, since it has already demonstrated state-of-the-art performance on table structure recognition benchmarks, it provides a reliable backbone for our framework. We report ablations reported on SPRINT to demonstrate these strengths in Table (3). For further implementation details, we refer the reader to the original SPRINT [3] paper.

3.2. TATR for Physical Structure

To extract the physical structure, we use the TATR [9] V1.1 model pre-trained on FinTabNet, PubTabNet, and

Abbreviation	Description	Task Used For
DELTA	**D**oubly d**E**coupled tab**L**e recons**T**ruction **A**pproach (a proposed approach)	Table Reconstruction, measured by TEDS-Structure and TEDS Scores
TarQA	**TA**ble structu**R**e-aware **Q**uestion **A**nswering (a fine-tuned LLM on OTSL sequences)	TabQA Table-based Question Answering
TORQUE	**T**able **O**riented **R**econstruction and **Q**uestion-answering **U**pon d**E**vanagari (curated benchmark)	Hindi Table Reconstruction Hindi TabVQA
DELTA + TarQA	DELTA converts Table Image to OTSL TarQA answers Questions with OTSL	Decoupled VQA task
TarQA-OTSL	TarQA Fine-tuned on OTSL sequences.	TabVQA
TarQA-HTML	TarQA Fine-tuned on HTML sequences.	TabVQA
GT OTSL + TarQA-OTSL	Ground Truth OTSL given to TarQA-OTSL	TabVQA
GT HTML + TarQA-HTML	Ground Truth OTSL given to TarQA-OTSL	TabVQA

Table 1. Abbreviations and Their Descriptions

Abbreviation	Description
ANLS	Average Normalized Levenshtein Similarity
EM	Exact Match
GT	Ground Truth
LLMs	Large Language Models
OCR	Optical Character Recognition
OTSL	Optimized Table Structure Language
p.p.	percentage point
SPRINT	Script-agnostic Structure Recognition in Tables
TabQA	Table Question Answering
TabVQA	Table Visual Question Answering
TATR	Table Transformer
TEDS	Tree Edit Distance -based Similarity
TSR	Table Structure Recoginition
VLMs	Vision Language Models
WTQ	WikiTableQuestions

Table 2. General abbreviations used in the paper.

PubTables-1M. TATR, built on DETR [1], predicts six classes, of which we only leverage `table-row` and `table-column` to estimate the rows and columns. For inference, we set the detection threshold to 0.25 and apply non-maximum suppression (NMS) with an IoU threshold of 0.25 on `table-row` predictions to minimize overlap and improve consistency. The resulting values are then aligned with the output sequence predicted by SPRINT, ensuring coherence between physical and logical structures.

3.3. TATR and SPRINT for Complete TSR

This step is responsible for integrating the logical structure (tag sequence predicted by SPRINT) with the physical structure (list of bounding boxes corresponding to detected rows and columns) to produce a final HTML representation of the table. Each `<td>` element in the output is annotated with its bounding box coordinates, as well as rowspan and colspan attributes whenever merged cells are detected. While row and column bounding boxes intersect to form candidate cells, the crucial constraint is that there exists a one-to-one mapping between each logical cell predicted by SPRINT and its corresponding physical bounding box. The algorithm enforces this alignment to guarantee that both spatial positioning and spanning attributes are preserved.

Formally, as seen in Algorithm (1), it takes as input an OTSL matrix M (generated by SPRINT, note that it can be converted to HTML in a lossless manner, but since SPRINT directly gives an OTSL string, we leverage that initially) and a set of bounding boxes $Cells$ (derived from TATR). The OTSL matrix encodes the table layout, where each entry specifies whether a position corresponds to a cell (`C`), a new row marker (`N`), or is empty. The process begins by initializing an empty HTML string H. For each entry (i, j) in M:

- If it corresponds to a cell, the bounding box $cell$ is re-

Test	Training	SPRINT Config	TEDS-S Simple	TEDS-S Complex	TEDS-S Overall
PubTabNet	PubTabNet	*Layers: 3, Shape: 32*128	97.91	91.17	94.61
	PubTabNet	Layers: 4, Shape: 32*128	98.12	92.84	95.53
	All	Layers: 6, Shape: 32*128	**98.11**	92.98	95.60
	All	Layers: 6, Shape: 128*128	98.00	**93.32**	**95.71**
FinTabNet	FinTabNet	Layers: 6, Shape: 32*32	**98.39**	94.57	96.41
	FinTabNet	Layers: 6, Shape: 32*128	98.30	97.46	97.88
	All	Layers: 6, Shape: 128*128	98.31	97.73	98.01
	FinTabNet	Layers: 6, Shape: 128*128	98.35	**97.74**	**98.03**
PubTables-1M	PubTables-1M	Layers: 6, Shape: 32*128	98.19	92.69	95.50
	All	Layers: 8, Shape: 32*128	98.88	93.34	96.00
	All	Layers: 6, Shape: 32*128	98.87	94.80	96.75
	All	Layers: 6, Shape: 128*128	**98.92**	**96.54**	**97.68**

Table 3. Results on different test sets for the SPRINT component of DELTA trained on various datasets. The training set of 'All' refers to the combined training dataset of all three datasets. The config is dictated by two parameters, mainly the number of decoder layers and the shape (dimensions) of the input image, which is resized in the preprocessing stage. * indicates that the maximum permissible length of prediction was set to 192 for that experiment, and the length was set to 224 otherwise. All the results are reported on the canonical validation set of PubTabNet [4] and canonical test sets of FinTabNet [4] and PubTables-1M [4]

trieved from $Cells$. The function `get_cell_spans` computes the extent of row and column spans by checking consecutive overlaps in the SPRINT output.

- If spans are present, the bounding box is extended accordingly, and an HTML td tag with the appropriate rowspan and/or colspan attributes is generated and appended to H.
- If no spans are present, a simple td tag with its bounding box is appended. In both cases, the processed bounding box is stored in the $bbox$ attribute.

Whenever an entry is marked as `N`, the algorithm closes the current row and begins a new one. After all entries are processed, the final HTML string is completed.

The output consists of the structured HTML string H, which encodes the table with explicit row and column spans that maintain the bounding boxes associated with each logical cell. This design ensures that even complex tables with merged rows/columns are reconstructed faithfully, while retaining both logical order (from SPRINT) and spatial grounding (from TATR).

3.4. OCR Ablations

The TSR module of DELTA achieves highly reliable structure predictions, with TEDS-Structure scores exceeding 95% on standard datasets. However, the overall TEDS score is hampered by OCR quality, making OCR the primary performance bottleneck. OCR errors arise from noise, complex layouts, slanted text, and special symbols. Since OCR follows TSR in our pipeline, its choice is critical. To assess this impact, we compare two widely used engines: Tesseract [8] and EasyOCR [2]. Tesseract provides broad multilingual support, modularity, and ease of integration but suffers from lower accuracy on noisy data, lacks GPU acceleration, and is slow at inference. In contrast, EasyOCR is GPU-compatible, nearly three times faster, and consistently more accurate. It integrates seamlessly with detection and structure recognition modules, offers greater control over outputs, and is modular enough to be replaced with CRNN-based or fine-tuned models. Empirically, EasyOCR achieves consistently stronger results (Table 4), improving TEDS scores of almost all the datasets. The weighted average increases by 12.45 p.p. across FinTabNet, PubTabNet, FinTabNetQA, and TORQUE, which also translates into higher TabVQA accuracy. The slight drop in performance on PubTabNet stems from its relatively clean images with few OCR errors, rather than a limitation of EasyOCR. In this setting, Tesseract and EasyOCR perform nearly identically, with negligible differences in TEDS scores. In contrast, PubTables-1M poses a greater challenge: its large scale makes Tesseract impractically slow for experiments, whereas EasyOCR strikes a balance between speed and accuracy, achieving a reasonable TEDS score of 54.8. Table (5) highlights the latency comparison, underscoring EasyOCR's superior efficiency. Accordingly, EasyOCR is used as the default OCR module in our DELTA pipeline, though it can be readily replaced with a stronger alternative.

To isolate OCR as the primary performance bottleneck, we prepared Ground Truth (GT) mapped predictions, where the content of each predicted HTML cell was replaced with the corresponding ground-truth content while preserving the predicted structure. Care was taken to ensure accurate row-wise fidelity and cell mapping. This setup is equivalent to DELTA predictions followed by perfect OCR. As shown by the resulting TEDS scores (Table 4) above 90% in most cases. This shows that the content errors are almost entirely

Algorithm 1 Converting OTSL Matrix (logical structure) and cell boxes (physical structure) into HTML Sequence.

```
Input: OTSL matrix M of size R × C, list of cell bounding boxes Cells
Output: HTML table string H, list of structured cells S
H ← "<table><tbody>"
for i = 1 → R do
    Append "<tr>" to H
    for j = 1 → C do
        if M[i, j] = C then
            cell ← Cells[i, j]
            (rs, cs) ← get_cell_spans(M, i, j)
            if rs > 0 ∧ cs > 0 then
                Extend td to cover row and column spans
                Append <td rowspan=rs+1 colspan=cs+1 bbox=cell> to H
            else if rs > 0 then
                Extend td vertically
                Append <td rowspan=rs+1 bbox=cell> to H
            else if cs > 0 then
                Extend td horizontally
                Append <td colspan=cs+1 bbox=cell> to H
            else
                Append <td bbox=cell> to H
            end if
        else if M[i, j] = N then
            Append "</tr>" to H
        end if
    end for
end for
Append "</tbody></table>" to H
return H
```

attributable to OCR. This study clearly indicates that integrating a stronger OCR module can substantially boost the table reconstruction quality of DELTA, effectively narrowing the problem.

4. Evaluation Metrics

Table (6) lists the metrics used for evaluation. TEDS and TEDS-S are applied to the table reconstruction task on PubTabNet, PubTables-1M, FinTabNet, and TORQUE; TEDS-S captures structural fidelity, while TEDS score takes into account both structure and content, making them well-suited for measuring both layout accuracy and content alignment. For the TabQA and TabVQA tasks on WTQ, FinTabNetQA, and TORQUE, we report ANLS, EM, and Relieved Accuracy, all of which range from 0 to 100. These metrics are standard in QA benchmarks, reflecting exact correctness (EM), tolerance to minor variations (ANLS), and robustness to different semantic answer formats (Relieved Accuracy). Together, they ensure fair and comprehensive comparisons across diverse methods.

Dataset	**Tesseract**	**Easy OCR**	**GT Mapped**
FinTabNet	41.5	**55.9**	91.2
PubTabNet	**54.4**	53.0	87.8
FinTabNetQA	70.0	**84.0**	92.5
PubTables-1M	-	54.8	86.8
TORQUE	40.8	**63.6**	83.5

Table 4. TEDS scores across different OCR modules and ground truth content mapped to the structure of DELTA. Results show consistent improvements on FinTabNet, FinTabNetQA, and TORQUE datasets.

OCR Engine	**Average Latency per image (in secs)**
Tesseract	13.9
EasyOCR	**4.3**

Table 5. Average Latency comparison of Tesseract and EasyOCR per table image in seconds calculated on the FinTabNet test set.

5. OTSL Ablation

Figure (1) presents qualitative examples illustrating the compactness of the OTSL format relative to HTML. We include two scenarios: one featuring a large table and another involving a large table with a more complex structure. These examples demonstrate that OTSL not only preserves structural fidelity but also provides a significantly more compact representation than HTML. The compactness of OTSL further benefits LLMs by reducing context length, enabling more efficient and accurate table understanding. In our analysis, several input tables that were incorrectly processed in HTML format were correctly interpreted when encoded in OTSL, highlighting its effectiveness. This reduced representation allows larger tables to be processed without exceeding context limits, ultimately contributing to improved overall accuracy. Finally, the handling of mathematical expressions depends on the OCR module applied after structure recognition; therefore, OTSL itself does not impose any inherent limitations on the extraction of mathematical equations.

6. Qualitative Results

In this section, we present qualitative results for table reconstruction on FinTabNet and TORQUE datasets, as well as for TabQA and TabVQA tasks.

Metric	Definition
Tree Edit Distance-based Similarity (TEDS)	Measures the similarity between predicted and ground truth HTML table structures using tree edit distance. It evaluates both the structure and content of tables, which we use to evaluate DELTA pipeline.
TEDS-Structure (TEDS-S)	A TEDS variant that evaluates only the HTML tag sequence, ignoring content. We use it to assess DELTA's table structure recognition.
Average Normalized Levenshtein Similarity (ANLS)	Measures textual similarity between predicted and ground truth answers using the normalised Levenshtein distance, providing partial credit for near matches. We use this metric to evaluate TARQA on the WTQ dataset.
Exact Match (EM)	A strict binary metric that returns 1 if the predicted answer matches the ground truth exactly, and 0 otherwise.
Relieved-Accuracy	This metric considers predictions correct if any normalised form matches the ground truth, ignoring units and formatting. It emphasizes semantic equivalence while evaluating DELTA + TARQA on the FinTabNetQA dataset for the TabVQA task.

Table 6. Overview of metrics used for evaluating DELTA and TARQA.

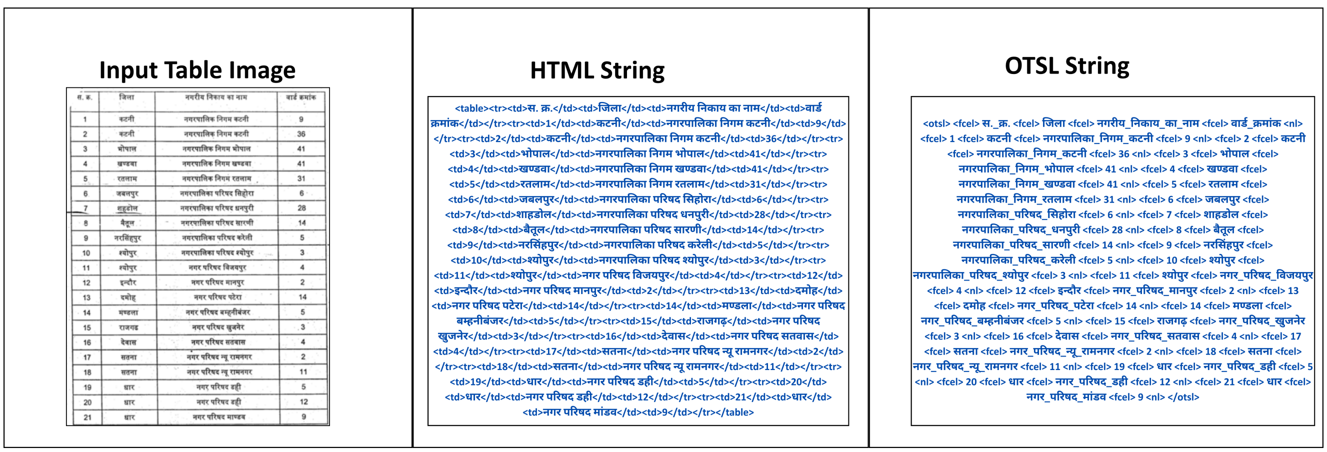

(a) HTML Character Count: 438, OTSL Character Count: 409

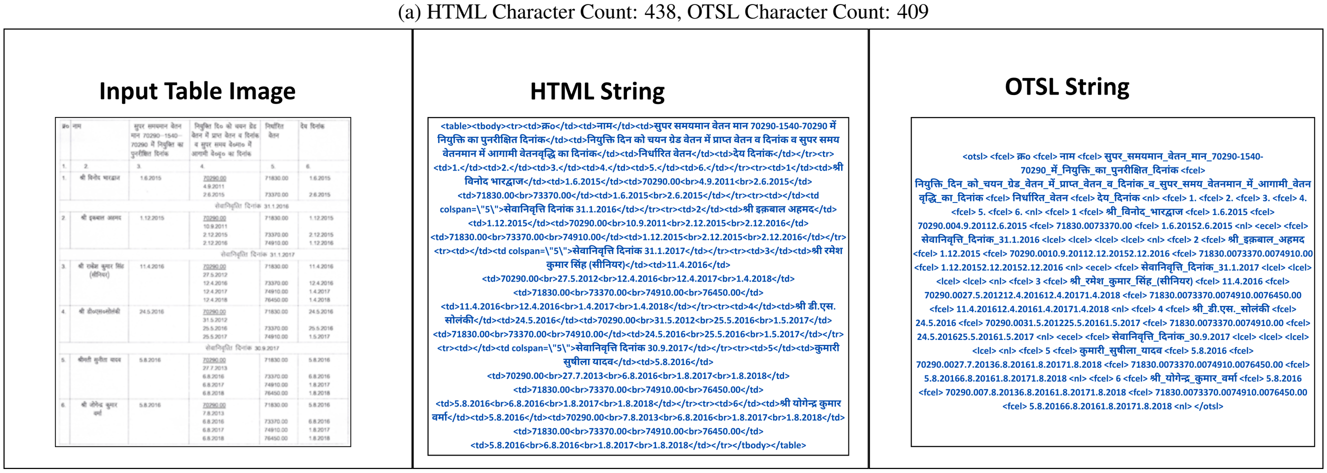

(b) HTML Character Count: 1607, OTSL Character Count: 1430

Figure 1. Qualitative examples from the TORQUE dataset illustrating the compactness of the OTSL format in terms of Character Counts from the input image.

6.1. Table Reconstruction

Figure (2a) demonstrates a successful case of DELTA, where the table image has a clean layout with multiple columns and well-separated numeric content. This results in perfect structural accuracy (TEDS-S = 100) and a high overall score (TEDS = 80.98). Similarly, Figure (2b) shows another positive example, with a clear two-column structure and well-defined row–column divisions, yielding TEDS-S = 100 and TEDS = 88.11. In contrast, Figure (3a) highlights a failure case arising from OCR extraction errors: while the structural score remains high (TEDS-S = 95.15), the overall content fidelity is poor (TEDS = 19.35). Figure (3b) presents another challenging example, where low image resolution and a borderless table with bullet points hinder accurate reconstruction, leading to low scores (TEDS-S = 26.09, TEDS = 19.95).

Similarly, Figure (4) presents qualitative examples of table reconstruction results on the TORQUE dataset. Figure (4a) demonstrates a successful reconstruction, reflected in a high TEDS-S and TEDS score of 97.22 and 96.06, respectively. In contrast, Figure (4b) highlights a failure case where slanted text and low-resolution input hinder accurate prediction, resulting in a significantly lower TEDS-S and TEDS score of 44.44 and 14.60, respectively.

6.2. Table Question Answering

Figure (5a) and Figure (5b) present qualitative examples from the WTQ dataset for TabQA tasks, where the answers predicted by the proposed TarQA framework are consistent with the ground truth. In contrast, Figure (6a) and Figure (6b) illustrate cases where the predictions deviate from the ground truth. For Figure (6a), the model outputs a value occurring immediately after the correct answer, indicating a limitation in its contextual understanding. For Figure (6b), the prediction is incorrect because the question is of a comparative type, for which the model has not been explicitly fine-tuned.

6.3. Table Visual Question Answering

Figure (7a) and Figure (7b) show illustrative TabVQA examples from the FintabnetQA dataset using the DELTA+TarQA pipeline. These examples highlight that even with difficult, borderless tables, the framework can accurately predict both numerical and textual answers, aligning with the ground truth. On the other hand, Figure (8a) and Figure (8b) showcase instances of divergence. In Figure (8a), the model becomes confused due to the complexity of the question type, while in Figure (8b), poor image quality leads to incorrect answer. These cases highlight both the strengths and current limitations of the approach, while also indicating clear directions, improved OCR, and enhanced reasoning for future enhancements.

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

[2] JaidedAI. Easyocr: Ready-to-use ocr with 80+ supported languages and all popular writing scripts. `https://github.com/JaidedAI/EasyOCR`, 2020. Accessed: 2025-04-20. 3

[3] Dhruv Kudale, Badri Vishal Kasuba, Venkatapathy Subramanian, Parag Chaudhuri, and Ganesh Ramakrishnan. Sprint: Script-agnostic structure recognition in tables. *arXiv preprint arXiv:2503.11932*, 2025. 1

[4] Maksym Lysak, Ahmed Nassar, Nikolaos Livathinos, Christoph Auer, and Peter Staar. Optimized table tokenization for table structure recognition. In *Document Analysis and Recognition - ICDAR 2023*, pages 37–50, Cham, 2023. Springer Nature Switzerland. 3

[5] Ahmed Nassar, Nikolaos Livathinos, Maksym Lysak, and Peter Staar. Tableformer: Table structure understanding with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4614–4623, 2022. 1

[6] 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, 2020. 1

[7] Liang Qiao, Zaisheng Li, Zhanzhan Cheng, Peng Zhang, Shiliang Pu, Yi Niu, Wenqi Ren, Wenming Tan, and Fei Wu. Lgpma: Complicated table structure recognition with local and global pyramid mask alignment. In *International conference on document analysis and recognition*, pages 99–114. Springer, 2021. 1

[8] R. Smith. An overview of the tesseract ocr engine. In *Ninth International Conference on Document Analysis and Recognition (ICDAR 2007)*, pages 629–633, 2007. 3

[9] Brandon Smock, Rohith Pesala, and Robin Abraham. Pubtables-1m: Towards comprehensive table extraction from unstructured documents. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4634–4642, 2022. 1

[10] Xinyi Zheng, Doug 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, 2020. 1

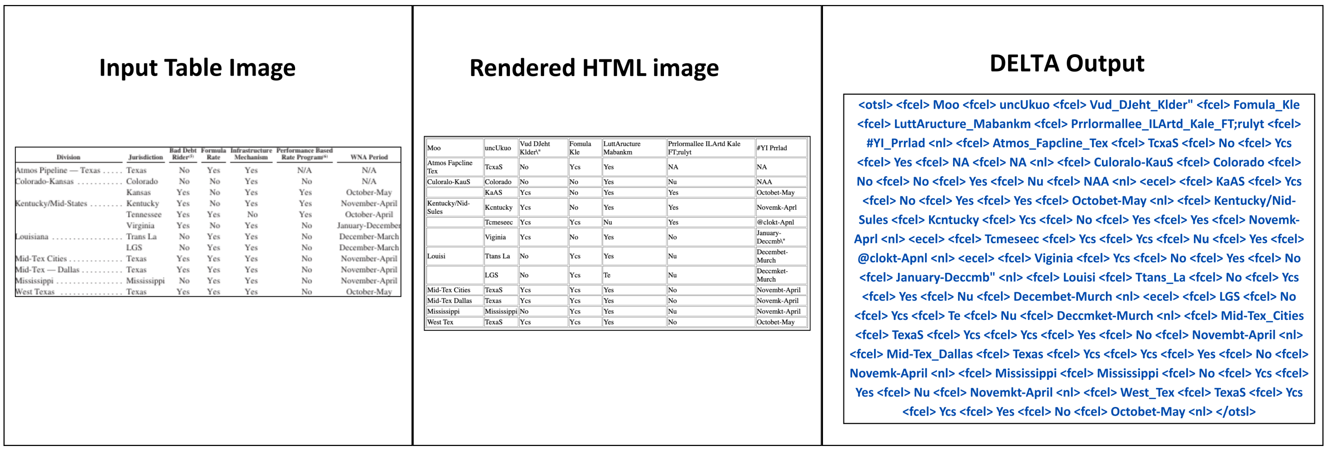

(a) TEDS-S = 100 and TEDS = 80.98

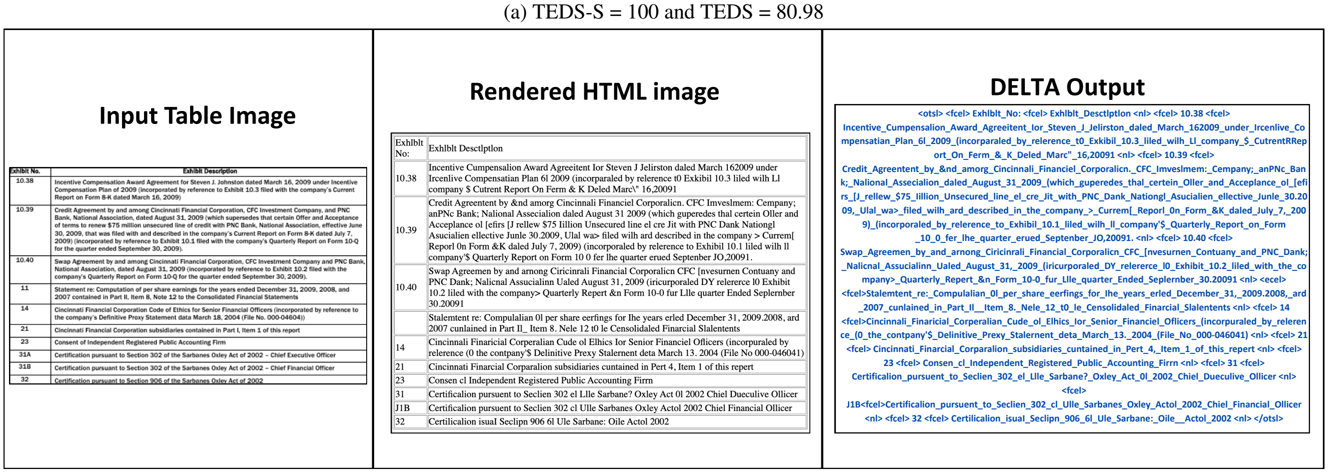

(b) TEDS-S = 100 and TEDS = 88.11.

Figure 2. Qualitative examples from the FinTabNet dataset illustrating table reconstruction outputs produced by DELTA framework.

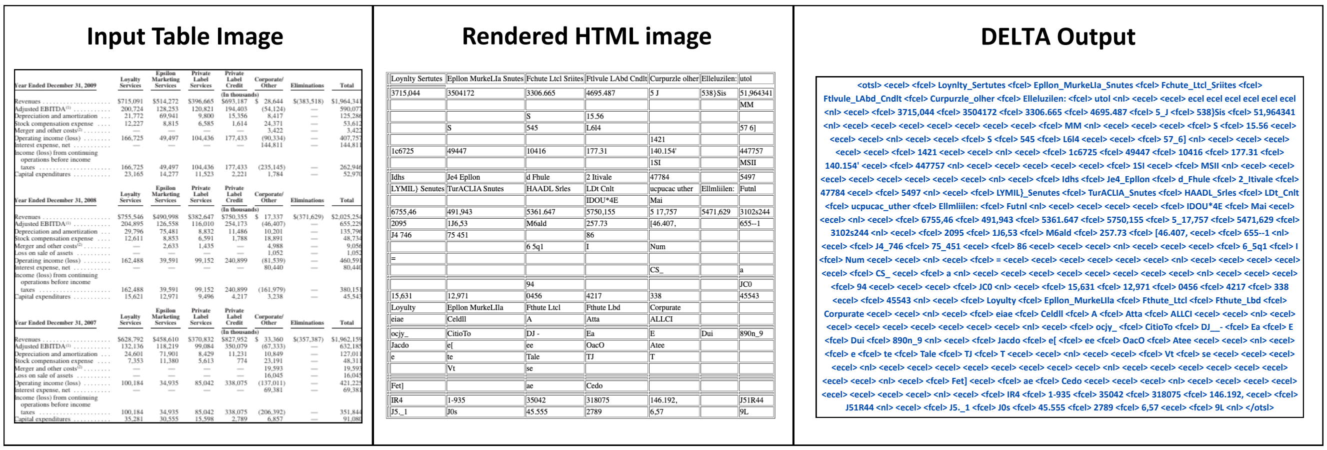

(a) TEDS-S = 95.15 and TEDS = 19.35

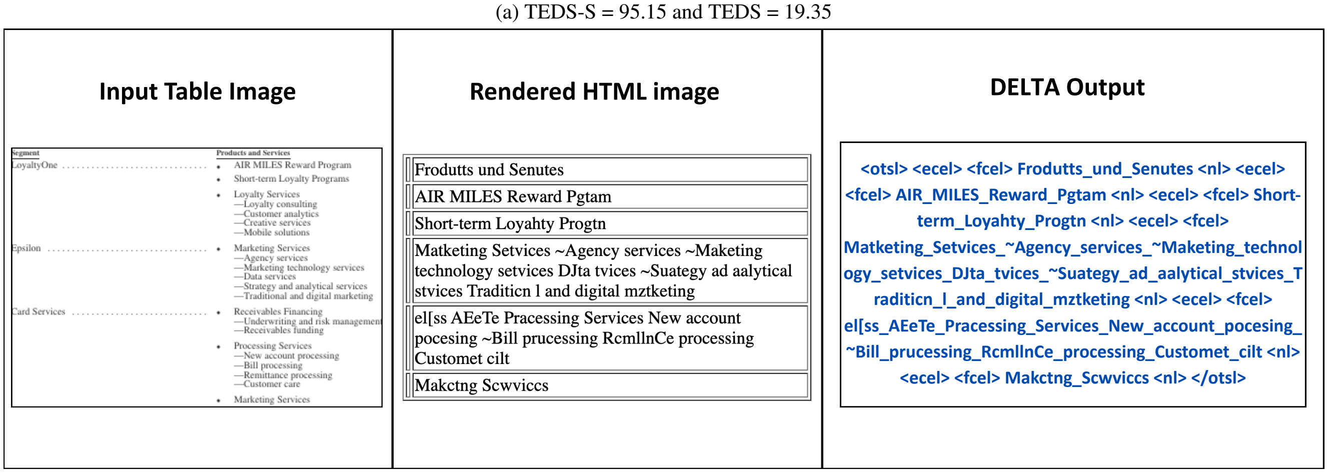

(b) TEDS-S = 26.09 and TEDS = 19.95

Figure 3. Inaccurate qualitative examples from the FinTabNet dataset illustrating table reconstruction outputs generated by DELTA.

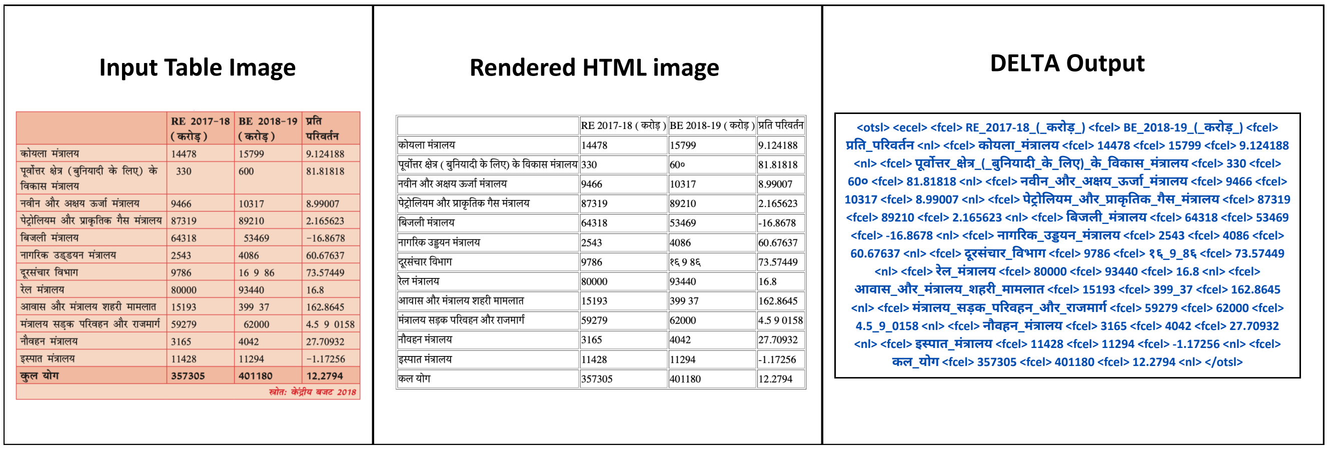

	RE 2017-18 (करोड़)	BE 2018-19 (करोड़)	प्रति परिवर्तन
कोयला मंत्रालय	14478	15799	9.124188
पूर्वोत्तर क्षेत्र (बुनियादी के लिए) के विकास मंत्रालय	330	600	81.81818
नवीन और अक्षय ऊर्जा मंत्रालय	9466	10317	8.99007
पेट्रोलियम और प्राकृतिक गैस मंत्रालय	87319	89210	2.165623
बिजली मंत्रालय	64318	53469	−16.8678
नागरिक उड्डयन मंत्रालय	2543	4086	60.67637
दूरसंचार विभाग	9786	16 9 86	73.57449
रेल मंत्रालय	80000	93440	16.8
आवास और मंत्रालय शहरी मामलात	15193	399 37	162.8645
मंत्रालय सड़क परिवहन और राजमार्ग	59279	62000	4.5 9 0158
नौवहन मंत्रालय	3165	4042	27.70932
इस्पात मंत्रालय	11428	11294	−1.17256
कुल योग	**357305**	**401180**	**12.2794**
			स्रोत: केंद्रीय बजट 2018

	RE 2017-18 (करोड़)	BE 2018-19 (करोड़)	प्रति परिवर्तन
कोयला मंत्रालय	14478	15799	9.124188
पूर्वोत्तर क्षेत्र (बुनियादी के लिए) के विकास मंत्रालय	330	60०	81.81818
नवीन और अक्षय ऊर्जा मंत्रालय	9466	10317	8.99007
पेट्रोलियम और प्राकृतिक गैस मंत्रालय	87319	89210	2.165623
बिजली मंत्रालय	64318	53469	-16.8678
नागरिक उड्डयन मंत्रालय	2543	4086	60.67637
दूरसंचार विभाग	9786	१६ 9 8६	73.57449
रेल मंत्रालय	80000	93440	16.8
आवास और मंत्रालय शहरी मामलात	15193	399 37	162.8645
मंत्रालय सड़क परिवहन और राजमार्ग	59279	62000	4.5 9 0158
नौवहन मंत्रालय	3165	4042	27.70932
इस्पात मंत्रालय	11428	11294	-1.17256
कल योग	357305	401180	12.2794

(a) Shows a prediction that closely aligns with the input table, achieving a high TEDS-S and TEDS score of 97.22 and 96.06, respectively.

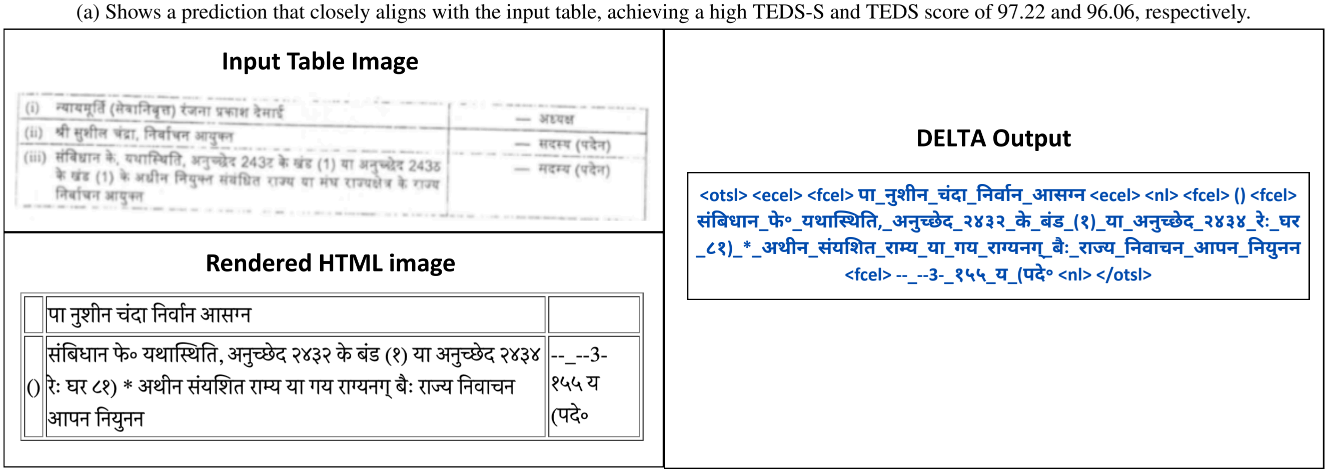

(i) न्यायमूर्ति (सेवानिवृत्त) रंजना प्रकाश देसाई	— अध्यक्ष
(ii) श्री सुशील चंद्रा, निर्वाचन आयुक्त	— सदस्य (पदेन)
(iii) संविधान के, यथास्थिति, अनुच्छेद 243ट के खंड (1) या अनुच्छेद 243ठ के खंड (1) के अधीन नियुक्त संबंधित राज्य या संघ राज्यक्षेत्र के राज्य निर्वाचन आयुक्त	— सदस्य (पदेन)

	पा नुशीन चंदा निर्वान आसग्न	
()	संबिधान फे० यथास्थिति, अनुच्छेद २४३२ के बंड (१) या अनुच्छेद २४३४ रे: घर ८१) * अथीन संयशित राम्य या गय राग्यनग् बै: राज्य निवाचन आपन नियुनन	--_--3-१५५ य (पदे०

(b) Depicts a challenging case with slanted text and low-resolution input, leading to notable differences between the predicted and input tables and a lower TEDS-S and TEDS score of 44.44 and 14.60, respectively.

Figure 4. Qualitative results of the proposed DELTA framework on a TORQUE sample for the table reconstruction task.

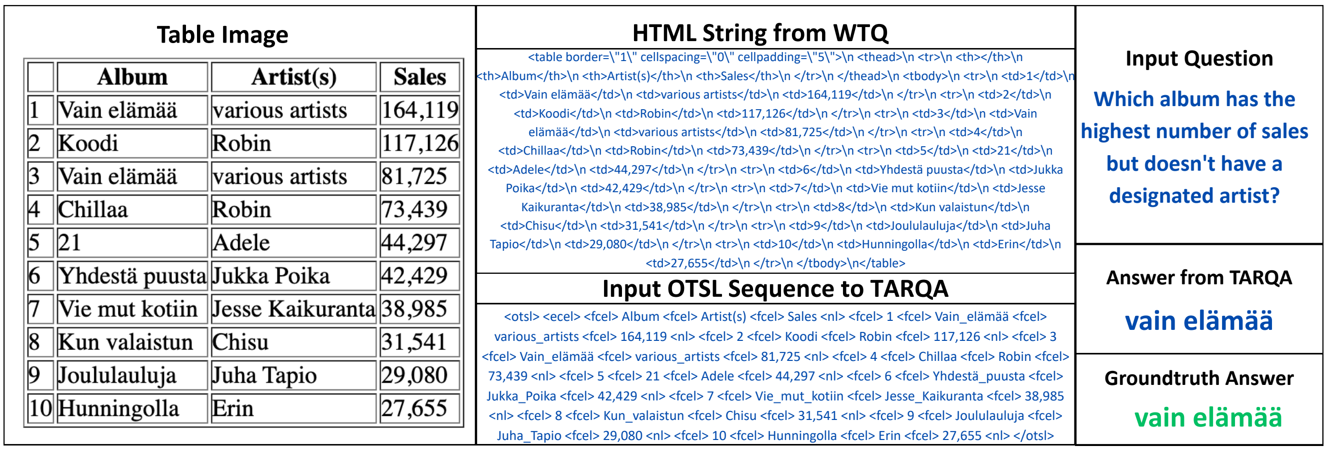

(a) The predicted answers align with the ground truth

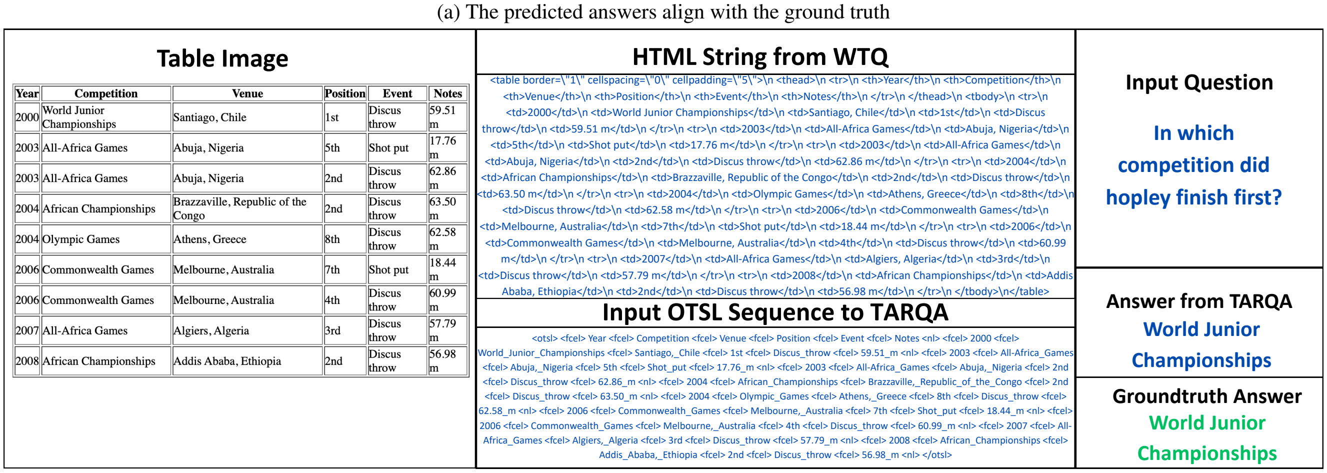

(b) The predicted answers align with the ground truth

Figure 5. Qualitative examples from the WTQ dataset for TabQA tasks, illustrating answers generated TarQA-OTSL.

Table Image

Pos	No	Driver	Team	Laps	Time/Retired	Grid	Points
1	1	Sébastien Bourdais	Newman/Haas Racing	66	1:51:31.146	2	34
2	9	Justin Wilson	RuSPORT	66	+3.528 secs	1	29
3	5	Will Power	Team Australia	66	+46.536 secs	4	26
4	2	Bruno Junqueira	Newman/Haas Racing	66	+1:04.023	3	23
5	15	Alex Tagliani	Team Australia	66	+1:18.033	8	22
6	6	Oriol Servià	PKV Racing	66	+1:28.745	7	19
7	4	Nelson Philippe	CTE Racing - HVM	66	+1:29.997	10	17
8	27	Andrew Ranger	Mi-Jack Conquest Racing	65	+ 1 Lap	16	16
9	3	David Martínez	Forsythe Racing	65	+ 1 Lap	9	13
10	7	Buddy Rice	Forsythe Racing	65	+ 1 Lap	14	11
11	34	Charles Zwolsman	Mi-Jack Conquest Racing	65	+ 1 Lap	12	10
12	18	Antônio Pizzonia	Rocketsports Racing	65	+ 1 Lap	15	9
13	11	Jan Heylen	Dale Coyne Racing	65	+ 1 Lap	17	8
14	10	Ryan Briscoe	RuSPORT	64	+ 2 Laps	5	7
15	19	Andreas Wirth	Dale Coyne Racing	64	+ 2 Laps	18	6
16	20	Katherine Legge	PKV Racing	63	+ 3 Laps	13	5
17	8	Mario Domínguez	Rocketsports Racing	59	Retired	11	4
18	14	Dan Clarke	CTE Racing - HVM	7	Differential	6	3

HTML String from WTQ

<table border=\"1\" cellspacing=\"0\" cellpadding=\"5\">\n <thead>\n <tr>\n <th>Pos</th>\n <th>No</th>\n <th>Driver</th>\n <th>Team</th>\n <th>Laps</th>\n <th>Time/Retired</th>\n <th>Grid</th>\n <th>Points</th>\n </tr>\n </thead>\n <tbody>\n <tr>\n <td>1</td>\n <td>1</td>\n <td>Sébastien Bourdais</td>\n <td>Newman/Haas Racing</td>\n <td>66</td>\n <td>1:51:31.146</td>\n <td>2</td>\n <td>34</td>\n </tr>\n <tr>\n <td>2</td>\n <td>9</td>\n <td>Justin Wilson</td>\n <td>RuSPORT</td>\n <td>66</td>\n <td>+3.528 secs</td>\n <td>1</td>\n <td>29</td>\n </tr>\n <tr>\n <td>3</td>\n <td>5</td>\n <td>Will Power</td>\n <td>Team Australia</td>\n <td>66</td>\n <td>+46.536 secs</td>\n <td>4</td>\n <td>26</td>\n </tr>\n <tr>\n <td>4</td>\n <td>2</td>\n <td>Bruno Junqueira</td>\n <td>Newman/Haas Racing</td>\n <td>66</td>\n <td>+1:04.023</td>\n <td>3</td>\n <td>23</td>\n </tr>\n <tr>\n <td>5</td>\n <td>15</td>\n <td>Alex Tagliani</td>\n <td>Team Australia</td>\n <td>66</td>\n <td>+1:18.033</td>\n <td>8</td>\n <td>22</td>\n </tr>\n <tr>\n <td>6</td>\n <td>6</td>\n <td>Oriol Servià</td>\n <td>PKV Racing</td>\n <td>66</td>\n <td>+1:28.745</td>\n <td>7</td>\n <td>19</td>\n </tr>\n <tr>\n <td>7</td>\n <td>4</td>\n <td>Nelson Philippe</td>\n <td>CTE Racing - HVM</td>\n <td>66</td>\n <td>+1:29.997</td>\n <td>10</td>\n <td>17</td>\n </tr>\n <tr>\n <td>8</td>\n <td>27</td>\n <td>Andrew Ranger</td>\n <td>Mi-Jack Conquest Racing</td>\n <td>65</td>\n <td>+ 1 Lap</td>\n <td>16</td>\n <td>16</td>\n </tr>\n <tr>\n <td>9</td>\n <td>3</td>\n <td>David Martínez</td>\n <td>Forsythe Racing</td>\n <td>65</td>\n <td>+ 1 Lap</td>\n <td>9</td>\n <td>13</td>\n </tr>\n <tr>\n <td>10</td>\n <td>7</td>\n <td>Buddy Rice</td>\n <td>Forsythe Racing</td>\n <td>65</td>\n <td>+ 1 Lap</td>\n <td>14</td>\n <td>11</td>\n </tr>\n <tr>\n <td>11</td>\n <td>34</td>\n <td>Charles Zwolsman</td>\n <td>Mi-Jack Conquest Racing</td>\n <td>65</td>\n <td>+ 1 Lap</td>\n <td>12</td>\n <td>10</td>\n </tr>\n <tr>\n <td>12</td>\n <td>18</td>\n <td>Antônio Pizzonia</td>\n <td>Rocketsports Racing</td>\n <td>65</td>\n <td>+ 1 Lap</td>\n <td>15</td>\n <td>9</td>\n </tr>\n <tr>\n <td>13</td>\n <td>11</td>\n <td>Jan Heylen</td>\n <td>Dale Coyne Racing</td>\n <td>65</td>\n <td>+ 1 Lap</td>\n <td>17</td>\n <td>8</td>\n </tr>\n <tr>\n <td>14</td>\n <td>10</td>\n <td>Ryan Briscoe</td>\n <td>RuSPORT</td>\n <td>64</td>\n <td>+ 2 Laps</td>\n <td>5</td>\n <td>7</td>\n </tr>\n <tr>\n <td>15</td>\n <td>19</td>\n <td>Andreas Wirth</td>\n <td>Dale Coyne Racing</td>\n <td>64</td>\n <td>+ 2 Laps</td>\n <td>18</td>\n <td>6</td>\n </tr>\n <tr>\n <td>16</td>\n <td>20</td>\n <td>Katherine Legge</td>\n <td>PKV Racing</td>\n <td>63</td>\n <td>+ 3 Laps</td>\n <td>13</td>\n <td>5</td>\n </tr>\n <tr>\n <td>17</td>\n <td>8</td>\n <td>Mario Domínguez</td>\n <td>Rocketsports Racing</td>\n <td>59</td>\n <td>Retired</td>\n <td>11</td>\n <td>4</td>\n </tr>\n <tr>\n <td>18</td>\n <td>14</td>\n <td>Dan Clarke</td>\n <td>CTE Racing - HVM</td>\n <td>7</td>\n <td>Differential</td>\n <td>6</td>\n <td>3</td>\n </tr>\n </tbody>\n</table>

Input OTSL Sequence to TARQA

<otsl> <fcel> Pos <fcel> No <fcel> Driver <fcel> Team <fcel> Laps <fcel> Time/Retired <fcel> Grid <fcel> Points <nl> <fcel> 1 <fcel> 1 <fcel> Sébastien_Bourdais <fcel> Newman/Haas_Racing <fcel> 66 <fcel> 1:51:31.146 <fcel> 2 <fcel> 34 <nl> <fcel> 2 <fcel> 9 <fcel> Justin_Wilson <fcel> RuSPORT <fcel> 66 <fcel> +3.528_secs <fcel> 1 <fcel> 29 <nl> <fcel> 3 <fcel> 5 <fcel> Will_Power <fcel> Team_Australia <fcel> 66 <fcel> +46.536_secs <fcel> 4 <fcel> 26 <nl> <fcel> 4 <fcel> 2 <fcel> Bruno_Junqueira <fcel> Newman/Haas_Racing <fcel> 66 <fcel> +1:04.023 <fcel> 3 <fcel> 23 <nl> <fcel> 5 <fcel> 15 <fcel> Alex_Tagliani <fcel> Team_Australia <fcel> 66 <fcel> +1:18.033 <fcel> 8 <fcel> 22 <nl> <fcel> 6 <fcel> 6 <fcel> Oriol_Servià <fcel> PKV_Racing <fcel> 66 <fcel> +1:28.745 <fcel> 7 <fcel> 19 <nl> <fcel> 7 <fcel> 4 <fcel> Nelson_Philippe <fcel> CTE_Racing_-_HVM <fcel> 66 <fcel> +1:29.997 <fcel> 10 <fcel> 17 <nl> <fcel> 8 <fcel> 27 <fcel> Andrew_Ranger <fcel> Mi-Jack_Conquest_Racing <fcel> 65 <fcel> +_1_Lap <fcel> 16 <fcel> 16 <nl> <fcel> 9 <fcel> 3 <fcel> David_Martínez <fcel> Forsythe_Racing <fcel> 65 <fcel> +_1_Lap <fcel> 9 <fcel> 13 <nl> <fcel> 10 <fcel> 7 <fcel> Buddy_Rice <fcel> Forsythe_Racing <fcel> 65 <fcel> +_1_Lap <fcel> 14 <fcel> 11 <nl> <fcel> 11 <fcel> 34 <fcel> Charles_Zwolsman <fcel> Mi-Jack_Conquest_Racing <fcel> 65 <fcel> +_1_Lap <fcel> 12 <fcel> 10 <nl> <fcel> 12 <fcel> 18 <fcel> Antônio_Pizzonia <fcel> Rocketsports_Racing <fcel> 65 <fcel> +_1_Lap <fcel> 15 <fcel> 9 <nl> <fcel> 13 <fcel> 11 <fcel> Jan_Heylen <fcel> Dale_Coyne_Racing <fcel> 65 <fcel> +_1_Lap <fcel> 17 <fcel> 8 <nl> <fcel> 14 <fcel> 10 <fcel> Ryan_Briscoe <fcel> RuSPORT <fcel> 64 <fcel> +_2_Laps <fcel> 5 <fcel> 7 <nl> <fcel> 15 <fcel> 19 <fcel> Andreas_Wirth <fcel> Dale_Coyne_Racing <fcel> 64 <fcel> +_2_Laps <fcel> 18 <fcel> 6 <nl> <fcel> 16 <fcel> 20 <fcel> Katherine_Legge <fcel> PKV_Racing <fcel> 63 <fcel> +_3_Laps <fcel> 13 <fcel> 5 <nl> <fcel> 17 <fcel> 8 <fcel> Mario_Domínguez <fcel> Rocketsports_Racing <fcel> 59 <fcel> Retired <fcel> 11 <fcel> 4 <nl> <fcel> 18 <fcel> 14 <fcel> Dan_Clarke <fcel> CTE_Racing_-_HVM <fcel> 7 <fcel> Differential <fcel> 6 <fcel> 3 <nl> </otsl>

Input Question

Who finished directly after the driver who finished in 1:28.745?

Answer from TARQA

Andrew Ranger

Groundtruth Answer

Nelson Philippe

(a) The predictions differ from the ground truth

Table Image

#	Player	Goals	Caps	Career
1	Landon Donovan	57	155	2000–present
2	Clint Dempsey	36	103	2004–present
3	Eric Wynalda	34	106	1990–2000
4	Brian McBride	30	95	1993–2006
5	Joe-Max Moore	24	100	1992–2002
6T	Jozy Altidore	21	67	2007–present
6T	Bruce Murray	21	86	1985–1993
8	Eddie Johnson	19	62	2004–present
9T	Earnie Stewart	17	101	1990–2004
9T	DaMarcus Beasley	17	114	2001–present

HTML String from WTQ

<table border=\"1\" cellspacing=\"0\" cellpadding=\"5\">\n <thead>\n <tr>\n <th>#</th>\n <th>Player</th>\n <th>Goals</th>\n <th>Caps</th>\n <th>Career</th>\n </tr>\n </thead>\n <tbody>\n <tr>\n <td>1</td>\n <td>Landon Donovan</td>\n <td>57</td>\n <td>155</td>\n <td>2000–present</td>\n </tr>\n <tr>\n <td>2</td>\n <td>Clint Dempsey</td>\n <td>36</td>\n <td>103</td>\n <td>2004–present</td>\n </tr>\n <tr>\n <td>3</td>\n <td>Eric Wynalda</td>\n <td>34</td>\n <td>106</td>\n <td>1990–2000</td>\n </tr>\n <tr>\n <td>4</td>\n <td>Brian McBride</td>\n <td>30</td>\n <td>95</td>\n <td>1993–2006</td>\n </tr>\n <tr>\n <td>5</td>\n <td>Joe-Max Moore</td>\n <td>24</td>\n <td>100</td>\n <td>1992–2002</td>\n </tr>\n <tr>\n <td>6T</td>\n <td>Jozy Altidore</td>\n <td>21</td>\n <td>67</td>\n <td>2007–present</td>\n </tr>\n <tr>\n <td>6T</td>\n <td>Bruce Murray</td>\n <td>21</td>\n <td>86</td>\n <td>1985–1993</td>\n </tr>\n <tr>\n <td>8</td>\n <td>Eddie Johnson</td>\n <td>19</td>\n <td>62</td>\n <td>2004–present</td>\n </tr>\n <tr>\n <td>9T</td>\n <td>Earnie Stewart</td>\n <td>17</td>\n <td>101</td>\n <td>1990–2004</td>\n </tr>\n <tr>\n <td>9T</td>\n <td>DaMarcus Beasley</td>\n <td>17</td>\n <td>114</td>\n <td>2001–present</td>\n </tr>\n </tbody>\n</table>

Input OTSL Sequence to TARQA

<otsl> <fcel> # <fcel> Player <fcel> Goals <fcel> Caps <fcel> Career <nl> <fcel> 1 <fcel> Landon_Donovan <fcel> 57 <fcel> 155 <fcel> 2000–present <nl> <fcel> 2 <fcel> Clint_Dempsey <fcel> 36 <fcel> 103 <fcel> 2004–present <nl> <fcel> 3 <fcel> Eric_Wynalda <fcel> 34 <fcel> 106 <fcel> 1990–2000 <nl> <fcel> 4 <fcel> Brian_McBride <fcel> 30 <fcel> 95 <fcel> 1993–2006 <nl> <fcel> 5 <fcel> Joe-Max_Moore <fcel> 24 <fcel> 100 <fcel> 1992–2002 <nl> <fcel> 6T <fcel> Jozy_Altidore <fcel> 21 <fcel> 67 <fcel> 2007–present <nl> <fcel> 6T <fcel> Bruce_Murray <fcel> 21 <fcel> 86 <fcel> 1985–1993 <nl> <fcel> 8 <fcel> Eddie_Johnson <fcel> 19 <fcel> 62 <fcel> 2004–present <nl> <fcel> 9T <fcel> Earnie_Stewart <fcel> 17 <fcel> 101 <fcel> 1990–2004 <nl> <fcel> 9T <fcel> DaMarcus_Beasley <fcel> 17 <fcel> 114 <fcel> 2001–present <nl> </otsl>

Input Question

Who scored more goals: clint dempsey or eric wynalda?

Answer from TARQA

eric wynalda

Groundtruth Answer

Clint Dempsey

(b) The predictions differ from the ground truth

Figure 6. Qualitative examples from the WTQ dataset for TabQA tasks, illustrating answers generated by TARQA-OTSL.

Input Table Image

Period	Total Number of Shares Purchased[(a)]	Average Price Paid per Share[(b)]	Total Number of Shares Purchased as Part of Publicly Announced Plans or Programs	Maximum Dollar Value of Shares that May Yet Be Purchased Under the Plans or Programs[(c)]
10/01/18-10/31/18	36,701	$ 82.02	-	$ 5,579,603,383
11/01/18-11/30/18	3,145,000	63.75	3,138,171	5,379,603,637
12/01/18-12/31/18	7,812,656	60.86	7,804,590	4,904,604,184
Total	10,994,357	61.76	10,942,761	

Input Question

What was the total number of shares purchased in the period from 10/01/18 to 10/31/18?

Answer from TARQA

36,701

Groundtruth Answer

36,701

(a) The predicted answers coincide with the ground truth

Input Table Image

December 31, 2015 *dollars in millions*	Fair Value of Level 3 Assets	Valuation Technique	Significant Unobservable Input	Range (Weighted-Average)
Recurring				
Other investments — principal investments — direct:	$ 50	Individual analysis of the condition of each investment		
Debt instruments			EBITDA multiple	N/A (5.40)
Equity instruments of private companies			EBITDA multiple (where applicable)	5.40 - 6.70 (6.60)
Nonrecurring				
Impaired loans [(a)]	—	Fair value of underlying collateral	Discount	00.00 - 34.00% (15.00%)
Goodwill	1,060	Discounted cash flow and market data	Earnings multiple of peers	10.30 - 17.80 (12.79)
			Equity multiple of peers	1.25 - 1.56 (1.43)
			Control premium	10.00 - 30.00% (19.18%)
			Weighted-average cost of capital	12.00 - 13.00% (12.54%)

Input Question

What is the valuation technique used for Impaired loans?

Answer from TARQA

Fair value of underlying collateral

Groundtruth Answer

Fair value of underlying collateral

(b) The predicted answers coincide with the ground truth

Figure 7. Illustrative TabVQA examples from the FintabnetQA dataset showcasing the performance of DELTA+TARQA-OTSL.

Input Table Image

	Year ended December 31,		
	2014	2013	2012
Computed tax at statutory tax rate	$297	$212	$ 31
State income taxes, net of federal tax benefit (1)	22	15	5
Non-deductible expenses and other (2)	8	8	(8)
Foreign taxes	(17)	(17)	(15)
Total	$310	$218	$ 13

Input Question

What was the total for the year ended December 31, 2013?

Answer from TARQA

$1,080,544.00

Groundtruth Answer

$218

(a) The ground truth and predicted answers are diverging from each other.

Input Table Image

(In millions)	2008	2007	2006 (a)
Stock Option Awards	$35.9	$35.2	$61.9
Restricted Share/Unit Awards	21.2	15.9	7.5
Total Stock-based Compensation Expense	$57.1	$51.1	$69.4

(a) Includes $33.8 million and $2.9 million of stock option and restricted share expense, respectively, resulting from the accelerated vesting upon the change of control that occurred as a result of the Fisher merger.

Input Question

How much was the Restricted Share/Unit Awards in 2006?

Answer from TARQA

fas

Groundtruth Answer

7.5 million

(b) The ground truth and predicted answers are diverging from each other.

Figure 8. Illustrative TabVQA examples from the FintabnetQA dataset showcasing the performance of DELTA+TARQA-OTSL.