

Parameter-Efficient Abstractive Question Answering over Tables or Text

Vaishali Pal¹ Evangelos Kanoulas² Maarten de Rijke²

¹Discovery Lab, University of Amsterdam

²University of Amsterdam

v.pal, e.kanoulas, m.derijke@uva.nl

Abstract

A long-term ambition of information seeking question answering (QA) systems is to reason over multi-modal contexts and generate natural answers to user queries. Today, memory intensive pre-trained language models are adapted to downstream tasks such as QA by fine-tuning the model on QA data in a specific modality like unstructured text or structured tables. To avoid training such memory-hungry models while utilizing a uniform architecture for each modality, parameter-efficient adapters add and train small task-specific bottleneck layers between transformer layers. In this work, we study parameter-efficient abstractive QA in encoder-decoder models over structured tabular data and unstructured textual data using only 1.5% additional parameters for each modality. We also ablate over adapter layers in both encoder and decoder modules to study the efficiency-performance trade-off and demonstrate that reducing additional trainable parameters down to 0.7%–1.0% leads to comparable results. Our models out-perform current state-of-the-art models on tabular QA datasets such as Tablesum and FeTaQA, and achieve comparable performance on a textual QA dataset such as NarrativeQA using significantly less trainable parameters than fine-tuning.

1 Introduction

Information seeking systems over diverse contexts require model capabilities to reason over unstructured and structured data such as free-form text, tables, and images (Agrawal et al., 2016; Vakulenko et al., 2019; Hudson and Manning, 2019; Zhang et al., 2020; Zhu et al., 2021; Deldjoo et al., 2021). Such systems might have the additional requirement of generating natural language responses if deployed as task-oriented conversational agents (Wen et al., 2015; Carnegie and Oh, 2000; Rambow et al., 2001; Ratnaparkhi, 2002). Recent work on open-domain question answering (QA) predominantly addresses these challenges by fine-tuning

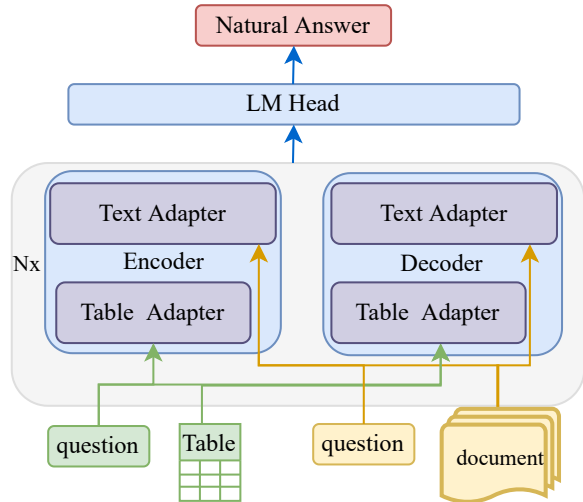


Figure 1: Parameter-efficient transfer learning using modality-specific (table/text) adapters for Abstractive Question Answering

massive pre-trained language models on different modalities such as tables and text (Yin et al., 2020; Herzig et al., 2020, 2021; Katsis et al., 2021; Nan et al., 2021). However, each model trained on a specific input type is incompatible with other modalities and requires modality-specific fine-tuning. For example, in tabular QA (Herzig et al., 2020), the structure of the table is learnt by training additional position embeddings (row and column identifiers) to identify which row and column a table cell belongs to. This renders such modality specific models incompatible with free-form text-based models. Multi-modal models (Zhu et al., 2021) can reason over both tables and text by concatenating the textual context and the flattened table, leading to longer input sequences and limiting the length of the context that can be encoded.

To address these challenges, we study parameter-efficient transfer learning for abstractive QA over tables and over text. We are motivated to use adapter-layers that inject small bottle-neck layers between frozen pre-trained transformer layers as they achieve comparable performance to fine-

tuning on a variety of tasks such as multi-lingual translation (Pfeiffer et al., 2020; Philip et al., 2020; Guo et al., 2020), classification (Houlsby et al., 2019a), text-to-text generation (Lin et al., 2020), domain-adaptation in dialogue state tracking, and response generation (Hung et al., 2021).

Ablation studies on adapter layers (Ruckl e et al., 2020) on masked language models such as BERT-base and RoBERTa over the GLUE benchmark demonstrate that removing beginning adapter layers leads to a minimal drop in performance. Extending adapter layer ablation over separate encoder and decoder modules is non-trivial as the conventional approach of sequential pruning of layers does not extend to consecutive encoder and decoder modules. Our work explores the interaction of adapter layers from both modules in the context of abstractive QA.

Lin et al. (2020) explore the impact of the adapter bottle-neck dimension for various language generation tasks over an auto-regressive model such as GPT-2 (Radford et al., 2019). They do not study tabular data nor ablate adapter layers, which is crucial in understanding impact of individual adapters in sequential transformer module architectures such as encoder-decoder. Our analysis is complementary to (Lin et al., 2020) as we ablate adapter layers to study parameter-performance trade-off whereas they only focus on adapter bottle-neck size. Also, we generalize beyond the text-to-text setting and explore language generation from structured or unstructured input such as tables and text. This introduces domain-shift in both the *task* and *structure* of the downstream data.

We propose a system, named **Parameter, Efficient, Abstractive Question Answering** (PeaQA), shown in Figure 1, which learns to reason over unstructured and structured input using a *shared* pre-trained language model and modality-specific adapter layers. We automatically transform hierarchical tables to regular tables to have a uniform representation without breaking associations between table cells. In addition, we extend the study of ablating adapter layers over both encoder and decoder modules.

Our main contributions are summarized as:

- (1) We perform parameter-efficient abstractive question answering over multi-modal context using only additional 1.5% of trainable parameters for each modality. Our adapter-tuned model outperforms existing work by

a large margin on tabular QA datasets and achieves comparable performance on a textual QA dataset.

- (2) We study tabular QA as a new modality that introduces massive input domain shift to pre-trained language models. We propose a 2-step transformation of hierarchical tables to sequences, which produces a uniform representation to be used by a single, shared pre-trained language model and modality-specific adapter layers. To the best of our knowledge, this is the first work that explores tabular QA question answering in a parameter-efficient manner.
- (3) We ablate adapter layers in both encoder and decoder modules to study their impact and show that beginning layers from both encoder and decoder can be eliminated without significant drop in performance. We also demonstrate that last encoder adapter layers are indispensable and have greater contribution than decoder layers at the same level.

2 Related Work

Tabular question answering. Tabular QA systems aim to answer questions from structured tables, which can be regular or hierarchical. Hierarchical tables can have header cells and body cells spanning across multiple rows and columns (Cheng et al., 2021). In most tabular QA systems (Herzig et al., 2020; Zhu et al., 2021; Katsis et al., 2021), the structure of the table is encoded in the embedding layer of large language models by introducing table specific position information such as row id and column id. Concurrent to our work, abstractive QA over tables (Nan et al., 2021; Cheng et al., 2021) poses additional challenges of generating natural answers by reasoning and aggregating discontinuous facts from the table.

Textual question answering. Question answering over text measures a system’s ability to comprehend free-form text in the user question and context passage(s) and predict an answer. The answer predicted can be extractive in nature, where the system identifies short text spans in the context passage to answer the user query (Lee et al., 2016; Seo et al., 2016; Rajpurkar et al., 2016; Pearce et al., 2021), or it can be abstractive, where it is required to generate a free-form answer (Yin et al., 2016; Mitra, 2017; Bauer et al., 2018; Reddy et al., 2019).

Transfer learning. Transfer learning techniques such as fine-tuning pre-trained models for down-

stream tasks, require a new set of parameters to be learnt for each new task. To avoid such memory intensive transfer learning methods, adapters have been proposed as a parameter-efficient method of adapting to new domains (Houlsby et al., 2019b; Pfeiffer et al., 2020). Adapters have been extended to language generation in a variety of generative tasks such as translation, summarization, multi-turn dialogue, and task-oriented natural language generation (Lin et al., 2020).

Our work combines all the aforementioned aspects to generate abstractive answers from *both* tables and text with only 0.7%–1.0% trainable parameters without compromising performance.

3 Model

We focus on encoder-decoder models for the task of abstractive question answering. We use a BART (Lewis et al., 2019) encoder-decoder architecture which comprises of a bidirectional encoder and an auto-regressive decoder. The input sequence consists of the question, the context title and context sequence preceded with prompts indicating the beginning of the each sub-sequence. Formally, the input sequence is represented as $\langle \text{question} \rangle q_0 q_1 \dots q_m \langle \text{title} \rangle t_1 t_2 \dots t_p \langle \text{context} \rangle c_0 c_1 \dots c_n$, where q_i is the i -th question token, t_j is the j -th title token, and c_k is the k -th context token. The context can either be a text passage or a flattened table. The parameters of the pre-trained BART model are frozen during training. Modality specific adapter layers added to the model are trained on either tabular context or textual context to generate natural answers.

4 Textual Question Answering

To study multi-modal abstractive QA, we first focus on free-form text as context to the system. We train adapter layers for textual context on the NarrativeQA dataset (Kočíský et al., 2018). NarrativeQA is a complex abstractive question answering dataset over stories. The dataset contains 32,747 samples in the training set, 3,461 samples in the validation set, and 10,557 samples in the test set. For our task, we have selected the input context passage to be the human annotated *summary* of each sample which is the Wikipedia page summary of the story and represented as a paragraph. The input to the model is the *question*, *title* and *summary* of each passage and the target is the abstractive answer.

5 Tabular Question Answering

We study tabular QA as a new modality which introduces massive input domain shift to pre-trained language models. Tables enforce structural constraints in their representation which is incompatible with the expected input format of pre-trained language models. To achieve our goal of parameter efficiency by utilizing a uniform pre-trained language model, we only train table specific adapter layers while keeping the pre-trained model frozen. However, this necessitates a uniform input representation for both tables and text. An additional challenge is introduced to maintain uniformity across different table types (regular, hierarchical).

For our task, we explore 2 tabular QA datasets, namely, Tablesum (Zhang et al., 2020) and FeTaQA (Nan et al., 2021). Tablesum consists of 200 unique Wikipedia tables over which questions and abstractive answers are manually annotated; 40% of the samples are questions over hierarchical tables but the tables in their released data are missing information in the hierarchical cells and their work do not handle hierarchies. We address this issue by extracting the wikitable from the respective Wikipedia pages and release a clean version of the dataset.¹

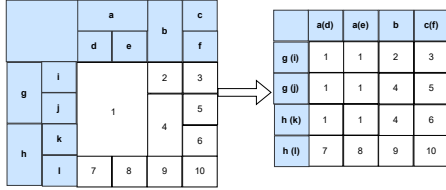
FeTaQA (Nan et al., 2021) is a larger abstractive tabular QA dataset consisting of question and free-form answers over 10,330 regular tables. The dataset consists of 7,326 samples in the training set, 1,001 in the validation set, and 2,003 in the test set. FeTaQA consists of human-annotated answers containing explanations involving entities and relations.

5.1 Table Representation

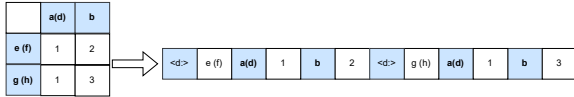
For our work, we choose to represent all tables uniformly in a two-step process: (1) Transformation of a hierarchical table into a regular table; and (2) Linearization of a regular table into a flattened sequence which can be encoded with a language model.

Linearize hierarchical table headers. Hierarchical table headers are linearized into a single row of headers by the following process. A header cell spanning multiple columns is duplicated and split into multiple cells. Next, the cell values over which this header spans are concatenated with the entire split. Repeating this process over all header rows flattens the hierarchical header into a sequential

¹The cleaned data and code can be found at <https://github.com/kolk/Pea-QA>



(a) A multi-span table represented as a regular table.



(b) Linearize regular table to a sequence of *key:value* pairs.

Figure 2: Table representation.

one. We depict this process in Figure 2a, which yields a linear header $a(d)$, $a(d)$, b , $e(f)$.

Linearizing table body. Multi-span table body cells are parsed differently than headers. Each table body cell is replicated with one or multiple header cells depending on its span across columns. Cells that span across multiple rows are replicated with all the spanned rows. This process leads to a regular table. We flatten the regular table in row-major form, concatenating rows sequentially. Each row is a sequence of (*key*, *value*) pairs where a key is a column header and the value is the cell value of that column as depicted in Figure 2b.

6 Experimental Setup

We seek to answer the following research questions with our experiments: (RQ1) How does adapter-tuning perform compared to fine-tuning in the context of multi-modal input? (RQ2) Do all adapter layers across the encoder and decoder contribute equally to performance across tasks/modalities?

6.1 Fine-Tuning

We perform all our experiments on the *large* variant of BART model. We fine-tune the BART-large model over the 3 datasets as the state-of-the-art fine-tuned models utilize different architectures for different datasets making comparison with adapter-tuning difficult. We treat our fine-tuned BART models on the 3 datasets as baselines. We sweep learning rates from $\{8e^{-4}, 6e^{-4}, 3e^{-4}, 1e^{-4}, 5e^{-5}, 4e^{-5}, 3e^{-5}, 2e^{-5}, 1e^{-5}\}$ and select the best performing learning rate for each dataset. We select $4e^{-5}$ for fine-tuning on Tablesum, $8e^{-4}$ on FeTaQA datasets and $2e^{-5}$ to fine-tune NarrativeQA. We use a batch size of 4 and gradient accumulation of 8 to emulate an effective batch size of 32. The maximum target sequence length is set to 200 for tabular QA datasets and to 100 for the textual QA

dataset. On the Tablesum dataset, we follow 5-fold cross validation as described in the original work to evaluate our models. On FeTaQA and NarrativeQA, we utilize the test split for evaluating our models. We train the model on each dataset for 15 epochs and evaluate on Rouge-2, Rouge-L and sacreBLEU metrics.

6.2 Adapter-Tuning

We perform adapter-tuning as a parameter-efficient alternative to adapt BART-large model to the abstractive question answering task across different modalities. We first freeze all layers of the pre-trained BART-large model which was trained on text reconstruction as mentioned in the original BART paper (Lewis et al., 2019). We add bottleneck adapter layers from the Houlsby adapter configuration (Houlsby et al., 2019a) which are trained to adapt to the downstream abstractive question answering task and also to modality specific input context. Each adapter layer has a bottle-neck embedding size of 64. As mentioned in Section 6.1, we sweep learning rates and select the best performing learning rate for each dataset. We select $6e^{-4}$ for the tabular QA datasets Tablesum and FeTaQA, and select $1e^{-1}$ to train the textual QA dataset NarrativeQA. We use the same batch size and maximum target sequence length as fine-tuning for effective comparison. A summary of hyper-parameters are mentioned in Table 1.

Dataset	Params	ATune	FTune
All	scheduler	linear	linear
	batch size	32	32
	seed	6	6
	max epochs	15	15
Tablesum	learning rate	6e-4	4e-5
	input length	200	200
FeTaQA	learning rate	6e-4	8e-4
	input length	100	100
NarrativeQA	learning rate	1e-4	2e-5
	input length	50	50

Table 1: Hyper-parameters for training. **ATune** indicates Adapter-tuning, **FTune** indicates Fine-tuning, **All** indicates all 3 datasets.

6.3 Ablation Study: Adapter Pruning

Adapter-layer pruning has been explored on the GLUE benchmark in (Ruckl  et al., 2020), which

Dataset	Model	Training	Rouge-1	Rouge-2	Rouge-L	BLEU
Tablesum (Zhang et al., 2020)	GPT2	fine-tune	0.272	0.073	0.200	5.35
	T5		0.362	0.143	0.276	10.43
	Ours (Pea-QA)	fine-tune(Baseline)	0.400	0.186	0.316	6.30
		Adapter-tune	0.393	0.186	0.312	6.75
FeTaQA (Nan et al., 2021)	T5-small	fine-tune	0.550	0.330	0.470	21.60
	T5-base		0.610	0.390	0.510	28.14
	T5-large		0.630	0.414	0.530	30.54
	Ours (Pea-QA)	fine-tune(Baseline)	0.632	0.415	0.534	30.81
		Adapter-tune	0.651	0.436	0.553	33.45
NarrativeQA (Kočíský et al., 2018)	Masque (Nishida et al., 2019)	fine-tune	–	–	0.547	–
	Ours (Pea-QA)	fine-tune(Baseline)	0.518	0.268	0.515	21.07
		Adapter-tune	0.510	0.270	0.500	20.08

Table 2: Results: Scores obtained on the Tablesum, FeTaQA and NarrativeQA datasets.

demonstrates that removing adapter layers from the beginning of BERT-base and RoBERTa models leads to minimal performance drop. We extend adapter layer ablation to encoder-decoder architectures and hypothesize that this phenomenon should be observed on both the encoder and decoder modules. However, it is non-trivial how the adapter-layers in the encoder and decoder interact with each other and contribute to performance. Previous studies (Rucklé et al., 2020) on adapter ablation prune consecutive adapter layers in masked language models. This approach does not extend directly to sequential modules of encoder-decoder where intra-module adapters not only contribute to their respective objective of encoding and decoding but also contributes to inter-module interaction and performance. To measure the impact of the adapter layers in different modules, we perform adapter ablation in both the encoder and decoder. First, we uniformly remove adapter layers from both encoder and decoder modules starting from the beginning layers of both modules and finally deleting all layers. This leads to 12 experiments corresponding to eliminating 12 encoder and 12 decoder adapter layers. To study interaction across inter-module adapters at different levels, we conduct 36 experiments of different configurations of adapter elimination from the last 6 levels of encoder and decoder. We analyze the performance by each configuration in Section 7.3.

7 Results

We compare the results of our baseline fine-tuned models with the state-of-the-art fine-tuned mod-

els in Section 7.1. We address (RQ1) “How does adapter-tuning perform compared to fine-tuning in the context of multi-modal input?” in Section 7.2 and (RQ2) “Do all adapter layers across the encoder and decoder contribute equally to performance across tasks/modalities?” in 7.3.

7.1 Fine-Tuned Models

We study the results of our baseline fine-tuned models with the state-of-the-art fine-tuned models for the 3 datasets. The results of the experiments are shown in Table 2. We observe that for the Tablesum dataset, our fine-tuned model outperform the best state-of-art T5 model on Rouge-1 by 3.8%, Rouge-2 by 4.3% and Rouge-L score by 4%. This can be attributed to fine-tuning our model on the clean version of the dataset. Our fine-tuned models perform comparably to the state-of-the-art T5-large on FeTaQA dataset, i.e, 0.2% on Rouge-1, 0.01% higher on Rouge-2, and 0.04% higher on Rouge-L. Our fine-tuning results on NarrativeQA are lower than state-of-the-art models trained with sophisticated reasoning architecture. The focus of this work was primarily on comparing fine-tuning and adapter-tuning and hence we leave explicit reasoning as part of future work.

7.2 Adapter-Tuned Models

We address (RQ1) by comparing the performance of adapter-tuned models to our baseline fine-tuned models. For Tablesum, as observed in Table 2 fine-tuning(baseline) marginally outperforms adapter-tuning with 0.7% higher Rouge-1 and 0.4% higher Rouge-L scores while having the same Rouge-2 score. For FeTaQA, adapter-tune shows a larger

Question: What and when were Akhila Kishore’s first two films?				
Target: akhila kishore made her debut in the kannada film padhe padhe (2013), and appeared in kathai thirai kathai vasanam iyakkam (2014).				
Table:	Year	Film	Role	Language
	2013	Padhe Padhe	Kanchana	Kannada
	2014	Kathai Thirai Kathai Vasanam Iyakkam	Daksha	Tamil
	2015	Inimey Ippadithaan	Akhila	Tamil

Adaper-tune: akhila kishore made her debut in the kannada film padhe padhe (2013) and kathai thirai kathai vasanam iyakkam (2014).				
Fine-tune: kathai thirai kathai vasanam iyakkam (2014) and inimey ippadithaan (2015) were kannada films.				
Question: Who is the starring actor of Aastik?				
Target: aastik is a 1956 hindi film starring shahu modak, paro devi and meenakshi.				
Table:	Title	Director	Cast	
	
	Aastik	S. P. Kalla	Shahu Modak, Paro Devi, Meenakshi, B. M. Vyas, Praveen Paul	
	Alam Ara	Nanubhai Vakil	Daljeet, Chitra, Tiwari, Niranjana Sharma, Minu Mumtaz,...	
	
Adaper-tune: aastik is a 1956 bollywood film starring shahu modak.				
Fine-tune: a directed by s. p. kalla.				
Question: What were the three films directed by Yakub and when were they released?				
Target: yakub directed three films: sagar ka sher in 1937, uski tamanna in 1939, and, in 1949, aiye.				
Table:	Year	Film	Director	
	
	1937	Sagar Ka Sher (Lion of Sagar)	Yakub	
	
	1939	Uski Tamanna (Her Last Desire)	Yakub	
	
1949	Aiye	Yakub		
...		
Adaper-tune: yakub directed three films: sagar ka sher (lion of sagar) in 1937, uski tamanna (her last desire) in 1939 and aiye in 1949.				
Fine-tune: y directed by yakub.				

Table 3: Samples where adapter-tune outperforms fine-tune

performance gain with 1.9% on Rouge-1 and Rouge-L and 2.1% on Rouge-2 compared to fine-tuning. The insignificant gains of fine-tuning over adapter-tuning in tabular QA can be attributed to catastrophic forgetting (French, 1999; Kirkpatrick et al., 2017; Chen et al., 2020) induced by differences in the distribution of downstream tabular data format from the original text data format of pre-training.

To explore this phenomenon further, we analyse examples from FeTaQA dataset in Table 3 where adapter-tuning outperforms fine-tuning. We observe that the fine-tuned model is unable to disambiguate surface-form similarities from the column semantics in the first example. The intended semantics of the named-entity *Akhila Kishore* in the question is *Actor*. While the surface-form is similar to the column value *Akhila*, the intended semantics

is that of the column header *Role*. The fine-tuned model wrongly predicts the second and third row of the tabular context as correct grounding of information while adapter-tuning is able to disambiguate and predicts information from the first 2 rows as answer. We observe that the fine-tuned model also predicts information from the wrong column *Director* instead of *Cast* in the second example. Adapter-tune correctly identifies the column but partially generates the required information in the prediction. The third example depicts both non-factual and non-fluent prediction by the fine-tuned model.

We demonstrate an example of a hierarchical table of Tablesum in Table 4. The question requires aggregation on the table cells and has various summary-like targets associated with it. The hierarchical table mentions *Ed Sheeran* 3 times, but the actual number of occurrence is 10 times, from

Question: how many times was ed sheeran listed as the performer?

Targets:

- Ed Sheeran was listed as a performer twice in the table documenting the top hits of 2014 in Sweden. Other English-Language top performers included Bruce Springsteen, Sam Smith, and Coldplay, implying that English-Language music has significant success in Sweden.
- According to the table, in 2014, Ed Sheeran was only listed as the performer one time. It was for the song that he performed that is called "'I See Fire'", which was out in January and February of 2014

Table:

Week	Date	Song title	Performer	Album title	Performer
1	3 Jan 2014	"Timber"	Pitbull feat. ...	True	Avicii
2	10 Jan 2014				
3	17 Jan 2014	"I See Fire"	Ed Sheeran	High Hopes	Bruce ...
4	24 Jan 2014			True	Avicii
5	31 Jan 2014			Christer Sjögren sjunger Sinatra	Christer Sjögren
6	7 Feb 2014			True	Avicii
7	14 Feb 2014				
8	21 Feb 2014				
9	28 Feb 2014		
31	31 July 2014	Prayer in C...	Lilly Wood &...	Honky Tonk Rebels	Lasse Stefanz
32	7 Aug 2014		
42	16 Oct 2014	"The Days"	Avicii	X	Ed Sheeran
43	23 Oct 2014			Songs for Daddy	Jill Johnson
44	30 Oct 2014				
...

Adapter-tune: I found a table of the year 2014. It shows the week of the week, the song title, the artist, the title of the song, the album title, and the performer. There are 11 times that Ed Sheeran was listed as the performer in the year. The song title is "Timber" and the album is True.

Fine-tune: I found a table of Ed Sheeran's year in 2014. He was listed as the performer 14 times in the year 2014. The first time he was listed was on 3 January 2014 with the song "Timber" and the last time was on 4 April 2014 with "I See Fire".

Table 4: Example from the Tablesun dataset.

Adapter-tune		#Trainable parameters
Encoder adapters removed	Decoder adapters removed	
–	–	6,343,680 (1.56%)
0–2	12–14	4,757,760 (1.17%)
0–4	12–16	3,700,480 (0.91%)
0–6	12–18	2,643,200 (0.65%)
0–8	12–20	1,585,920 (0.39%)
0–10	12–22	528,640 (0.13%)
0–11	12–22	264,320 (0.07%)
fine-tune		406,291,456 (100%)

Table 5: Trainable parameters in the encoder and decoder. Encoder adapter layers are numbered from 0–11 and decoder adapter layers are numbered from 12–22. x – y implies all adapter layers from x to y inclusive.

Week 3 to Week 9, Week 31 and from Week 42 to Week 43. Our table transformation process handles this to produce a regular table with 10 cells containing Ed Sheeran as value. The models can simply aggregate over the mentions. As shown in Table 4,

both models generates long answers summarizing information from the context table. However, as the models do not explicitly handle cell aggregation, we observe factual mistakes in both adapter-tuned and fine-tuned models. The models find Tablesun samples challenging even though the generated language is fluent and readable.

For textual QA, on the NarrativeQA dataset, adapter-tuning performs comparable to fine-tuning with the adapter-tuned model achieving 0.8% lower Rouge-1, 1.8% higher Rouge-2 and 1.5% lower Rouge-L scores than fine-tuning.

We conclude that adapter-tuning performs better than fine-tuning for out-of-domain tabular data and comparable performance on in-domain text.

7.3 Ablation of adapter layers

We study (RQ2) by ablating adapter layers in both the encoder and decoder modules. We uniformly eliminate successive adapter layers from both encoder and decoder starting from the first layer in both modules and finally deleting all layers. This leads to 12 experiments corresponding to 12 en-

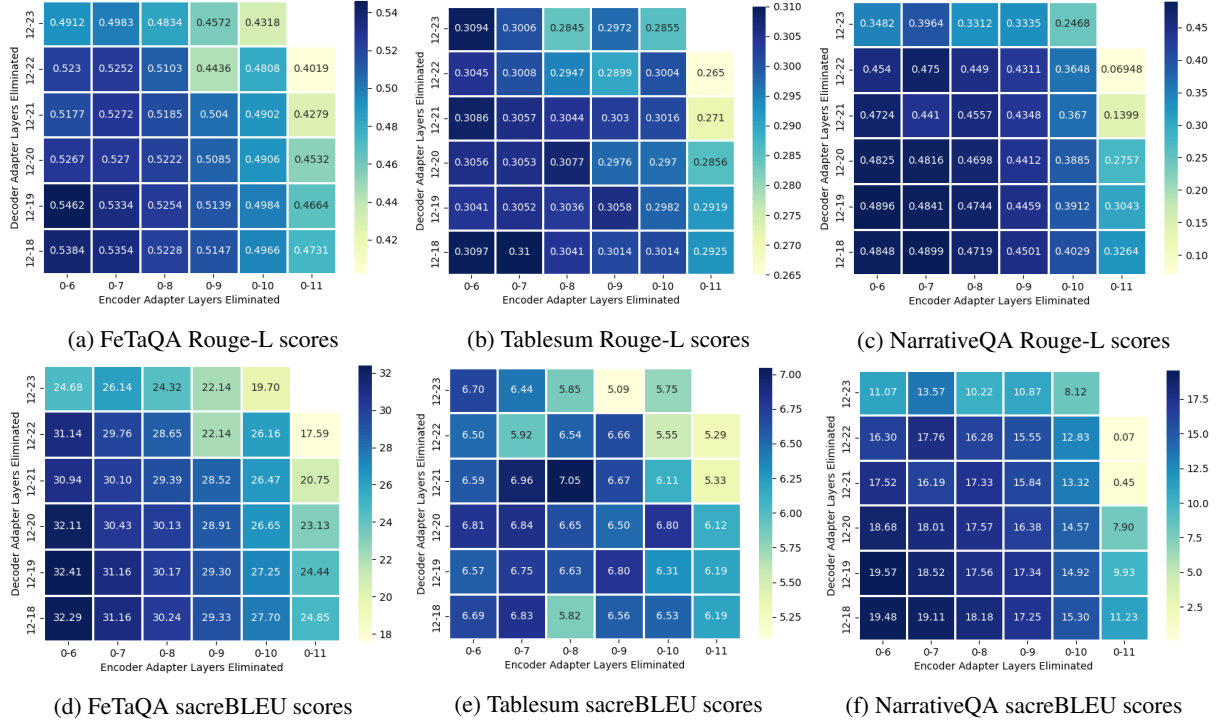


Figure 3: Adapter layer ablation scores. The X-axis represents range of encoder adapter layers deleted, the Y-Axis represents range of decoder adapter layers deleted. $x-y$ implies all adapter layers from x to y inclusive. There are 36 model ablation configurations displayed. The ablation starts from 0 to 6 encoder adapter layers removal and 12 to 18 decoder adapter layer removal represented by the bottom left cell ((0-6), (12-18)) and progressively increases deletion of encoder adapter layers along the X-axis and decoder adapter layers along the Y-axis.

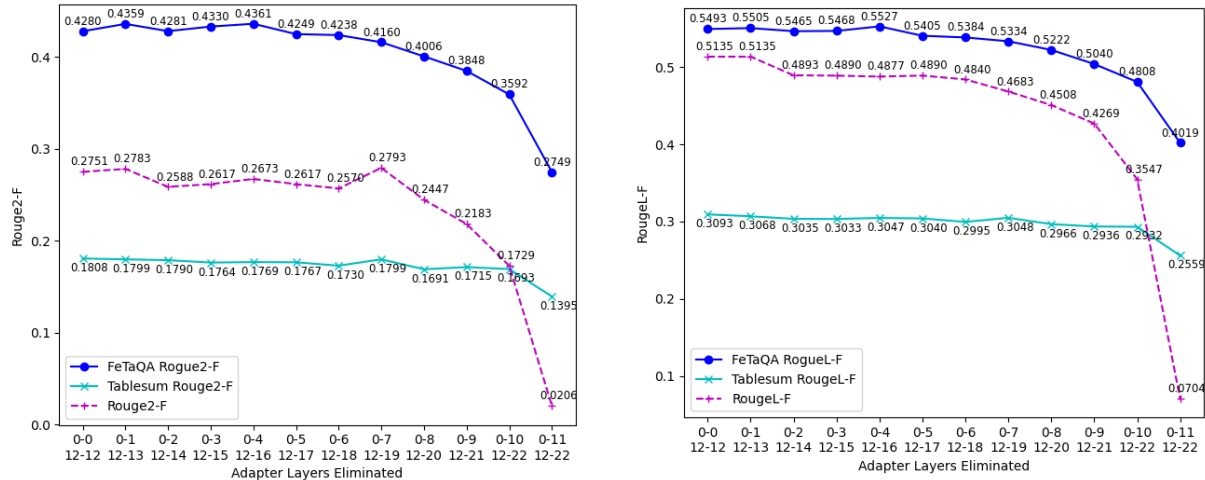


Figure 4: Adapter layer ablation Rouge2 F-scores. The X-axis depicts encoder-adapter layers (0-11) and decoder adapter layers (12-23) deleted progressively. Each $\binom{x-y}{r-s}$ represents F-score with encoder layers p to q deleted and decoder layers r to s deleted.

coder and 12 decoder adapter layers. We number the encoder adapter layers from 0-11 and the decoder adapter layers from 12-23. We measure the performance of the models using Rouge-2, Rouge-L² and sacreBLEU³ scores. The F-scores

²<https://pypi.org/project/rouge-score/>

³<https://github.com/mjpost/sacreBLEU>

Figure 5: Adapter layer ablation Rouge-L scores. The X-axis depicts encoder-adapter layers (0-11) and decoder adapter layers (12-23) deleted progressively. Each $\binom{x-y}{r-s}$ represents F-score with encoder layers p to q deleted and decoder layers r to s deleted.

for each dataset (NarrativeQA, Tablesum, FeTaQA) are shown in Figure 4, 5 and 6, respectively. We observe that as more adapter layers are eliminated, the performance drops across all datasets. However, the performance drop is minimal until the last adapter layers are also deleted. The inflection point varies across dataset but is limited to the last 2 layers of the encoder and decoder. For the Narra-

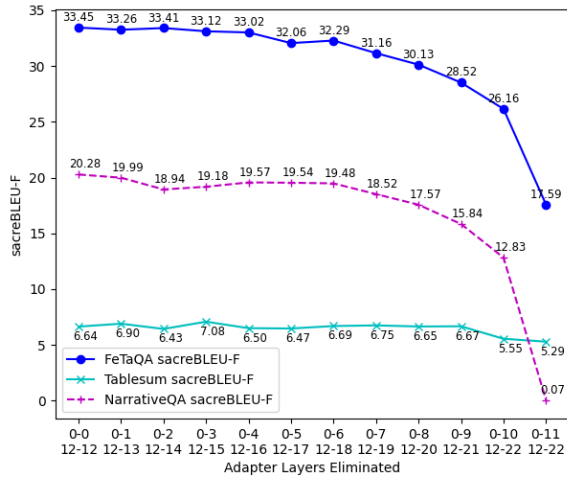


Figure 6: Adapter layer ablation sacreBLEU F-scores. The X-axis depicts encoder-adapter layers (0–11) and decoder adapter layers (12–23) deleted progressively. Each $\binom{x-y}{r-s}$ represents F-score with encoder layers p to q deleted and decoder layers r to s deleted.

tiveQA dataset, this point is when all layers till the second last adapter layer from both the encoder and decoder are deleted. For the FeTaQA and Tablesun datasets, the performance drops sharply only when the last encoder and decoder layers are removed.

To analyze contribution of the i -th adapter layer of encoder and decoder to performance, we perform ablation of adapter layers (0–6), (0–7), ..., (0–11) from encoder and adapter layers (12–18), (12–19), ..., (12–23) from decoder (decoder layers are numbered 12–23). This leads to 36 configurations where a configuration $(p-q, r-s)$ represents removal of all encoder adapters from p -th to q -th layer and all decoder adapters from r -th to s -th. The results are shown in Figure 3. We observe that performance remains comparable as we progressively eliminate adapter layers from encoder and decoder until the last layers. The performance drops steeply when we remove the last encoder and decoder adapter layers depicted towards the top-right corner of RougeL scores in Figures 3a, 3b, and 3c and BLEU scores in Figures 3d, 3e, and 3f. This implies that last adapter layers learns most of the domain information.

We also observe that the last encoder and decoder layers contribute differently to performance. Removing the last encoder layer (column 0–11) leads to substantial score drop across all decoder layers. This indicates that the last encoder layer is indispensable. Keeping only the last decoder adapter (row 12–23) is comparable to keeping last two last encoder layers (column 0–10). We also observe that retaining just the last 50% of adapter

layers from both encoder and decoder increases parameter efficiency by 0.7% parameters as summarized in Table 5 without significant compromise to performance.

8 Conclusion

We are the first to study parameter-efficient transfer learning over tables and text for abstractive question answering using adapters. We demonstrate that parameter efficient adapter-tuning outperforms fine-tuning on out-of-domain tabular data and achieves comparable results on in-domain textual data.

We propose a transformation from hierarchical tables to regular ones and further into a sequential form compatible with pre-trained model. We extend an existing ablation study of adapter layers to encoder-decoder setting and demonstrate that adapter layers from the end of the encoder is indispensable to encoding modality specific information than decoder adapter layers at the same level.

Our results are useful for exploring scalability of QA models in memory constrained situations with comparable performance while scaling across modalities using light-weight adapters.

One of the limitations of our work is that our models do not explicitly reason and aggregate over the table cells. This might lead to fluent but factually incorrect answers on challenging Tablesun dataset. Addressing this limitation is left as future work.

9 Acknowledgements

We would like to thank Elsevier for their support throughout this project and funding this work. This work was also supported by the NWO Innovational Research Incentives Scheme Vidi (016.Vidi.189.039), the NWO Smart Culture - Big Data / Digital Humanities (314-99-301), the H2020-EU.3.4. - SOCIETAL CHALLENGES - Smart, Green And Integrated Transport (814961). All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

References

- Aishwarya Agrawal, Jiasen Lu, Stanislaw Antol, Margaret Mitchell, C. Lawrence Zitnick, Dhruv Batra, and Devi Parikh. 2016. VQA: Visual question answering. *arXiv preprint arXiv:1505.00468*.
- Lisa Bauer, Yicheng Wang, and Mohit Bansal. 2018. Commonsense for generative multi-hop question answering tasks. In *EMNLP*.
- Alice Oh Carnegie and Alice H. Oh. 2000. Stochastic language generation for spoken dialogue systems. In *In Proc. of the ANLP/NAACL 2000 Wrkshp. on Conversational Systems*, pages 27–32.
- Sanyuan Chen, Yutai Hou, Yiming Cui, Wanxiang Che, Ting Liu, and Xiangzhan Yu. 2020. [Recall and learn: Fine-tuning deep pretrained language models with less forgetting](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7870–7881, Online. Association for Computational Linguistics.
- Zhoujun Cheng, Haoyu Dong, Zhiruo Wang, Ran Jia, Jiaqi Guo, Yan Gao, Shi Han, Jian-Guang Lou, and Dongmei Zhang. 2021. Hitab: A hierarchical table dataset for question answering and natural language generation. *arXiv preprint arXiv:2108.06712*.
- Yashar Deldjoo, Johanne R. Trippas, and Hamed Zamani. 2021. [Towards multi-modal conversational information seeking](#). In *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, Virtual Event, Canada, July 11-15, 2021*, pages 1577–1587. ACM.
- Robert M. French. 1999. [Catastrophic forgetting in connectionist networks](#). *Trends in Cognitive Sciences*, 3(4):128–135.
- Junliang Guo, Zhirui Zhang, Linli Xu, Hao-Ran Wei, Boxing Chen, and Enhong Chen. 2020. [Incorporating bert into parallel sequence decoding with adapters](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 10843–10854. Curran Associates, Inc.
- Jonathan Herzig, Thomas Müller, Syrine Krichene, and Julian Eisenschlos. 2021. [Open domain question answering over tables via dense retrieval](#). *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Jonathan Herzig, Pawel Krzysztof Nowak, Thomas Müller, Francesco Piccinno, and Julian Eisenschlos. 2020. [TaPas: Weakly supervised table parsing via pre-training](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4320–4333, Online. Association for Computational Linguistics.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019a. [Parameter-efficient transfer learning for NLP](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2790–2799. PMLR.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019b. [Parameter-efficient transfer learning for NLP](#). *arXiv preprint arXiv:1902.00751*.
- Drew A. Hudson and Christopher D. Manning. 2019. [GQA: A new dataset for real-world visual reasoning and compositional question answering](#). *arXiv preprint arXiv:1902.09506*.
- Chia-Chien Hung, Anne Lauscher, Simone Paolo Ponzetto, and Goran Glavaš. 2021. [DS-TOD: Efficient domain specialization for task oriented dialog](#). *arXiv preprint arXiv:2110.08395*.
- Yannis Katsis, Saneem Chemmengath, Vishwajeet Kumar, Samarth Bharadwaj, Mustafa Canim, Michael Glass, Alfio Gliozzo, Feifei Pan, Jaydeep Sen, Karthik Sankaranarayanan, and Soumen Chakrabarti. 2021. [AIT-QA: Question answering dataset over complex tables in the airline industry](#). *arXiv preprint arXiv:2106.12944*.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. 2017. [Overcoming catastrophic forgetting in neural networks](#). *arXiv preprint arXiv:1612.00796*.
- Tomáš Kočický, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis, and Edward Grefenstette. 2018. [The NarrativeQA reading comprehension challenge](#). *Transactions of the Association for Computational Linguistics*, 6:317–328.
- Kenton Lee, Shimi Salant, Tom Kwiatkowski, Ankur Parikh, Dipanjan Das, and Jonathan Berant. 2016. [Learning recurrent span representations for extractive question answering](#). *arXiv preprint arXiv:1611.01436*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. [Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#).
- Zhaojiang Lin, Andrea Madotto, and Pascale Fung. 2020. [Exploring versatile generative language model via parameter-efficient transfer learning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 441–459, Online. Association for Computational Linguistics.
- Rajarshee Mitra. 2017. [A generative approach to question answering](#). *arXiv preprint arXiv:1711.06238*.

- Linyong Nan, Chiachun Hsieh, Ziming Mao, Xi Victoria Lin, Neha Verma, Rui Zhang, Wojciech Kryściński, Nick Schoelkopf, Riley Kong, Xiangru Tang, Murori Mutuma, Ben Rosand, Isabel Trindade, Renusree Bandaru, Jacob Cunningham, Caiming Xiong, and Dragomir Radev. 2021. Fetaqa: Free-form table question answering. *arXiv preprint arXiv:2104.00369*.
- Kyosuke Nishida, Itsumi Saito, Kosuke Nishida, Kazutoshi Shinoda, Atsushi Otsuka, Hisako Asano, and Junji Tomita. 2019. [Multi-style generative reading comprehension](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2273–2284, Florence, Italy. Association for Computational Linguistics.
- Kate Pearce, Tiffany Zhan, Aneesh Komanduri, and Justin Zhan. 2021. A comparative study of transformer-based language models on extractive question answering. *arXiv preprint arXiv:2110.03142*.
- Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. 2020. [MAD-X: An Adapter-Based Framework for Multi-Task Cross-Lingual Transfer](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7654–7673, Online. Association for Computational Linguistics.
- Jerin Philip, Alexandre Berard, Matthias Gallé, and Laurent Besacier. 2020. [Monolingual adapters for zero-shot neural machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4465–4470, Online. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI Blog*.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [Squad: 100,000+ questions for machine comprehension of text](#). *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*.
- Owen Rambow, Srinivas Bangalore, and Marilyn Walker. 2001. [Natural language generation in dialog systems](#). In *Proceedings of the First International Conference on Human Language Technology Research*.
- Adwait Ratnaparkhi. 2002. [Trainable approaches to surface natural language generation and their application to conversational dialog systems](#). *Computer Speech & Language*, 16(3):435–455. Spoken Language Generation.
- Siva Reddy, Danqi Chen, and Christopher D Manning. 2019. CoQA: A conversational question answering challenge. *Transactions of the Association of Computational Linguistics (TACL)*.
- Andreas Rucklé, Gregor Geigle, Max Glockner, Tilman Beck, Jonas Pfeiffer, Nils Reimers, and Iryna Gurevych. 2020. Adapterdrop: On the efficiency of adapters in transformers. *arXiv preprint arXiv:2010.11918*.
- Minjoon Seo, Aniruddha Kembhavi, Ali Farhadi, and Hannaneh Hajishirzi. 2016. Bidirectional attention flow for machine comprehension. *arXiv preprint arXiv:1611.01603*.
- Svitlana Vakulenko, Kate Revoredo, Claudio Di Ciccio, and Maarten de Rijke. 2019. QRFA: A data-driven model of information-seeking dialogues. In *ECIR 2019: 41st European Conference on Information Retrieval*, pages 541–557. Springer.
- Tsung-Hsien Wen, Milica Gašić, Nikola Mrkšić, Pei-Hao Su, David Vandyke, and Steve Young. 2015. Semantically conditioned LSTM-based natural language generation for spoken dialogue systems. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1711–1721. Association for Computational Linguistics.
- Jun Yin, Xin Jiang, Zhengdong Lu, Lifeng Shang, Hang Li, and Xiaoming Li. 2016. [Neural generative question answering](#). *Proceedings of the Workshop on Human-Computer Question Answering*.
- Pengcheng Yin, Graham Neubig, Wen-tau Yih, and Sebastian Riedel. 2020. [Tabert: Pretraining for joint understanding of textual and tabular data](#). *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Shuo Zhang, Zhuyun Dai, Krisztian Balog, and Jamie Callan. 2020. Summarizing and exploring tabular data in conversational search. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*.
- Fengbin Zhu, Wenqiang Lei, Youcheng Huang, Chao Wang, Shuo Zhang, Jiancheng Lv, Fuli Feng, and Tat-Seng Chua. 2021. TAT-QA: A question answering benchmark on a hybrid of tabular and textual content in finance. *arXiv preprint arXiv:2105.07624*.

APPENDICES

We provide further details on statistics of the datasets used (Appendix A) and on the Rouge-2 scores for an encoder-decoder adapter layer ablation study (Appendix B).

A Dataset Statistics

Statistics of the three datasets, i.e., Tablesum, FeTaQA and NarrativeQA are listed in Table 6. Tablesum has the longest answer length. The answers are summary-like, often, describing aspects of the table contents. The FeTaQA dataset contains answers of mostly single sentences and targeted towards specific facts asked in the question. The NarrativeQA dataset focuses on questions from stories. The answer lengths vary from single words to long sentences. For the tabularQA dataset, Tablesum contains larger tables than the FeTaQA dataset even though it is limited to 200 unique tables over which questions are asked. The FeTaQA dataset’s tables contain more columns on average than Tablesum.

Tablesum	
Domain	Open
Modality	Table
Table-type	Regular
Training samples	798
Validation samples	200
Test samples	–
Max question length	114
Max target length	1,579
Max table row	155
Max table column	8

FeTaQA	
Domain	Open
Modality	Table
Table-type	Hybrid
Training samples	7,326
Validation samples	1,001
Test samples	2,003
Train max question length	165
Train max target length	338
Train max table rows	34
Train max table columns	30
Val max question length	182
Val target length	325
Val max table rows	34
Val max table columns	22
Test max question length	193

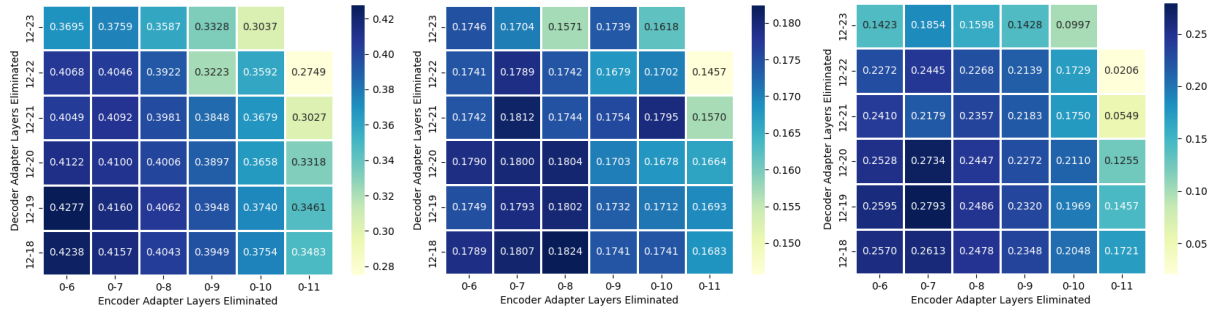
Test max target length	295
Test max table lows	34
Test max table columns	22

NarrativeQA	
Domain	Stories
Modality	Text
Training samples	65,494
Validation samples	6,922
Test samples	21,114
Train max question length	175
Train max target length	171
Train max context length	6,045
Val max question length	158
Val target length	187
Val max context length	6,033
Test max question length	1,220
Test target length	224
Test max context length	6,090

Table 6: Dataset Statistics

B Encoder-Decoder Adapter layer Ablation Rouge-2 Scores

Ablation results (Rouge-2 F-scores) of 36 configurations of adapter layers deleted from the later half of the encoder and decoder. Deleting the last encoder adapter layers leads to massive drop in performance as observed in the last three columns of Figures 7a, 7b and 7c. However, deleting the last decoder adapter layers results in better performance in comparison to the encoder layers at the same level as observed from the top 3 rows.



(a) FeTaQA Rouge-L scores

(b) Tablesum Rouge-L scores

(c) NarrativeQA Rouge-L scores

Figure 7: Adapter layer Rouge-2 ablation scores. The X-axis represents range of encoder adapter layers deleted, the Y-axis represents range of decoder adapter layers deleted. $x-y$ implies all adapter layers from x to y inclusive. There are 36 model ablation configurations displayed. The ablation starts from 0 to 6 encoder adapter layers removal and 12 to 18 decoder adapter layer removal represented by the bottom left cell ((0-6), (12-18)) and progressively increases deletion of encoder adapter layers along the X-axis and decoder adapter layers along the Y-axis.