

QFMTS: Generating Query-Focused Summaries over Multi-Table Inputs

Weijia Zhang^{a,*}, Vaishali Pal^{a,1}, Jia-Hong Huang^a, Evangelos Kanoulas^a and Maarten de Rijke^a

^aInformation Retrieval Lab (IRLab), University of Amsterdam

Abstract. Table summarization is a crucial task aimed at condensing information from tabular data into concise and comprehensible textual summaries. However, existing approaches often fall short of adequately meeting users’ information and quality requirements and tend to overlook the complexities of real-world queries. In this paper, we propose a novel method to address these limitations by introducing query-focused multi-table summarization. Our approach, which comprises a table serialization module, a summarization controller, and a large language model (LLM), utilizes textual queries and multiple tables to generate query-dependent table summaries tailored to users’ information needs. To facilitate research in this area, we present a comprehensive dataset specifically tailored for this task, consisting of 4, 909 query-summary pairs, each associated with multiple tables. Through extensive experiments using our curated dataset, we demonstrate the effectiveness of our proposed method compared to baseline approaches. Our findings offer insights into the challenges of complex table reasoning for precise summarization, contributing to the advancement of research in query-focused multi-table summarization.

1 Introduction

Table summarization involves the distillation of information from tabular data into a succinct format, typically a clear and human-readable description or textual table summary. This process aims to capture the key insights or trends encapsulated within the table, allowing for easier comprehension and interpretation by humans [2, 17, 23, 25, 36]. Traditional methods of table summarization take a single table as input and produce a fixed textual summary [12, 43]. However, the fixed nature of table summaries generated by these traditional approaches often falls short of meeting users’ information and quality requirements adequately. The quality of a summary plays a pivotal role in its utility. For instance, in business contexts, table summaries often play a critical role in shaping future business strategies, with their quality directly influencing the judgments of decision-makers. Poor-quality summaries may fail to capture essential aspects, leading to a misrepresentation of the data and conveying an inaccurate perspective.

To enhance the effectiveness of conventional table summarization, Zhao et al. [51] proposes to leverage textual queries as a starting point to generate query-dependent table summaries aligned with users’ information needs. Their approach accepts a single table and a user-specified textual query as inputs, generating a description or statement tailored to the query’s focus as output. However, in Zhao

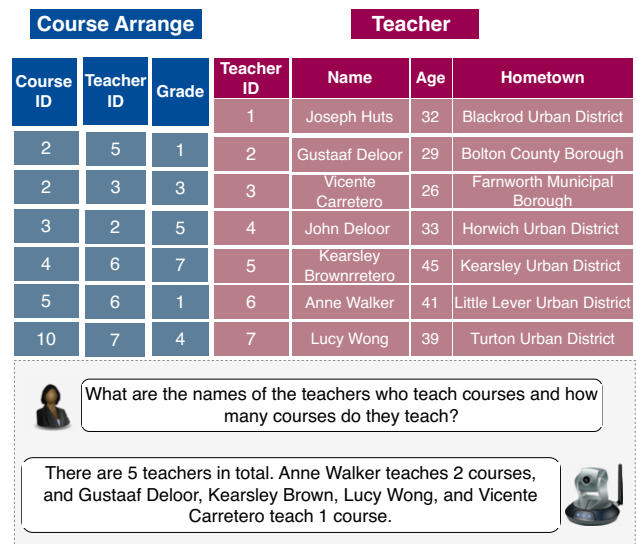


Figure 1: An example of query-focused multi-table summarization. Summarization models should combine the information from the two tables to produce a summary tailored to the query.

et al. [51], the authors presume that fulfilling the information requirements of a given query depends solely on data from a single table. This assumption overlooks the complexity of real-world scenarios, which frequently demand information from multiple data sources. Consequently, the intricate nature of real-world queries highlights the necessity for integrating information from multiple tables to address them effectively.

Let’s consider a practical scenario depicted in Figure 1 to understand how humans address complex real-world queries. For instance, the query “What are the names of the teachers who teach courses and how many courses do they teach?” entails two distinct information requirements – teachers’ names and their course teaching details. While the Teacher table in Figure 1 provides the teachers’ names, relying solely on this table is insufficient to generate a complete query-focused table summary. To completely address the query, we require data from another table, such as the Course Arrange table in Figure 1, which contains information about course arrangements for teachers listed in the Teacher table. By performing table join (i.e., multi-table reasoning) and count (i.e., arithmetic reasoning) operations, we can determine the number of courses taught by each teacher. Thus, addressing such a common real-world complex query involves employing multi-table reasoning and arithmetic reasoning.

* Email: w.zhang2@uva.nl.

¹ Equal contribution.

This practical and challenging scenario remains largely unexplored, underscoring the necessity for further exploration and advancement in this domain.

Motivated by the above observation of how humans handle such a scenario and aiming to bridge this research gap, we propose a new method designed to tackle the aforementioned practical and challenging situation: query-focused multi-table summarization. Our proposed approach involves taking multiple tables and a user-defined query as inputs, aiming to generate a query-relevant textual summary based on the inputs. The proposed approach primarily comprises a table serialization module, a summarization controller, and a large language model (LLM). The flowchart illustrating our proposed method is provided in Figure 2 to aid comprehension. Additionally, in light of the absence of an existing dataset for our introduced task, and to validate the effectiveness of our proposed query-focused multi-table summarization method, we create a query-focused multi-table summarization (QFMTS) dataset. This dataset comprises 4,909 query-summary pairs, each pair associated with multiple tables. For further details about our proposed dataset, please refer to Section 4.

In our comprehensive experiments, we assess the effectiveness of the proposed query-focused multi-table summarization method using our devised QFMTS dataset. The experimental findings demonstrate the superiority of our approach over baseline methods, shedding light on the challenges encountered by existing models in performing complex table reasoning to produce precise table summaries. To the best of our knowledge, we are the first to address the task of query-focused multi-table summarization.

Our primary contributions can be summarized as follows:

- **Introduction of Query-Focused Multi-Table Summarization:** We introduce a novel method tailored for query-focused multi-table summarization, aiming to overcome limitations present in current approaches that predominantly target single-table summarization. Our proposed method comprises a table serialization module, a summarization controller, and a large language model (LLM), offering a structured framework to tackle the complexities inherent in the task of query-focused multi-table summarization.
- **Development of a Comprehensive Dataset:** We develop a QFMTS dataset specifically tailored for the query-focused multi-table summarization task. The QFMTS dataset consists of 4,909 query-summary pairs, each intricately linked with multiple tables. It will serve as a valuable resource for researchers in this field, facilitating the exploration and validation of proposed methods in the future.
- **Extensive Experiments:** We conduct comprehensive experimental validation of our proposed query-focused multi-table summarization method using the introduced QFMTS dataset. Our experimental results demonstrate its effectiveness over baseline methods and provide insights into the challenges encountered in complex table reasoning for precise summarization.

2 Related Work

Table Summarization Table summarization involves generating a concise and informative summary from a given table. To address this task, prior studies [12, 25, 43] have primarily focused on summarizing the entire table without explicitly addressing users' specific information needs. However, in real-world applications, users frequently seek targeted information from table segments, underscoring the importance of generating personalized table summaries. Although Zhao et al. [51] introduced the initial human-annotated dataset for

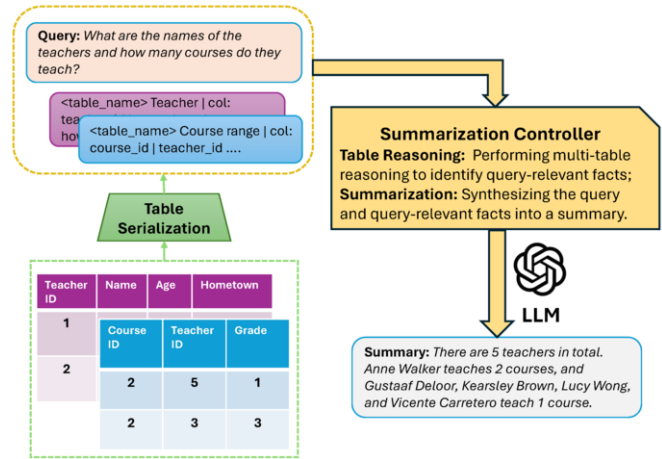


Figure 2: Overview of the proposed approach. Our approach integrates a table serialization module with a summarization controller. Initially, a set of tables is transformed into textual representations through the table serialization module. Subsequently, the summarization controller directs an LLM to perform table reasoning and produce a summary specifically tailored to a given query.

query-focused table summarization, their study restricted to single-table scenarios, lacking consideration of multi-table reasoning such as operations involving table *join* and *union*. In contrast, our work presents two primary distinctions: first, we leverage LLMs to assist in the data annotation; second, our proposed method addresses complex queries that necessitate the integration of information across multiple table contexts.

Query-Focused Text Summarization Query-focused text summarization aims to generate textual summaries based on a specific query and a collection of relevant documents. This research field has been extensively investigated with textual inputs [4, 45, 46, 47, 48, 49]. Traditional studies have faced challenges due to the scarcity of large-scale datasets, often resorting to distant supervision signals from adjacent fields such as generic summarization [39, 40]. Additionally, recent efforts have been directed towards creating synthetic large-scale datasets [15, 16, 26]. Despite these advancements, the application of query-focused summarization to tabular data remains relatively unexplored [51], especially in scenarios where queries span multiple tables. This study seeks to bridge this gap by exploring the efficacy of query-focused summarization in multi-table contexts.

3 Methodology

This section formulates the task of query-focused multi-table summarization and details our proposed approach. As illustrated in Figure 2, our approach comprises two primary components: a table serialization module and a summarization controller.

Task Formulation Query-focused multi-table summarization, denoted as QFMTS, involves generating a coherent and informative summary aimed at addressing a user query across multiple tables. Specifically, given a natural language query q and a set of input tables $\mathcal{T} = t_1, \dots, t_k$, a query-focused multi-table summarization model systematically engages in table-based reasoning across the contents of \mathcal{T} related to q , aiming to produce a textual summary s that effectively resolves the user query while maintaining factual accuracy and comprehensiveness.

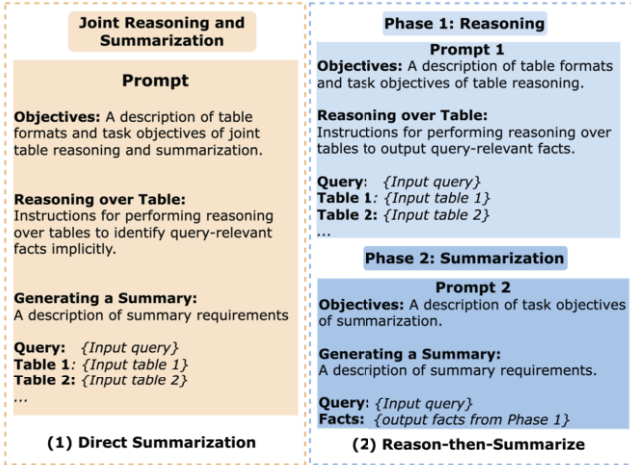


Figure 3: Two distinct approaches of the summarization controller: (1) Direct Summarization (2) Reason-then-Summarize. The former tackles the proposed task in an end-to-end manner. In contrast, the latter tackles the task across two independent phases.

Table Serialization Given that our approach is based on LLMs which process only textual data, it necessitates a table serialization to transform input tabular data into a textual format suitable for processing. In this work, we utilize a technique known as table linearization, which is widely used in table-to-text generation tasks [23, 24, 29, 32]. This technique transforms a table into a textual, sequential format using designated sentinel words. Specifically, a table identified by its name, $name$, consisting of m rows and n columns, is linearized as follows:

<table_name>: $name$ **col:** $h_1 | \dots | h_n$
row 1: $c_{1,1} | \dots | c_{1,n}$ **row m:** $c_{m,1} | \dots | c_{m,n}$.

where h_j and $c_{i,j}$ stand for the j -th column header and i -th row and j -th column cell, respectively.

3.1 Summarization Controller

Given a query and the corresponding linearized tables, our proposed summarization controller efficiently generates a comprehensive summary. This is achieved by integrating an LLM, such as GPT-3.5 [31], with a carefully designed prompting-based method. In this work, we explore two distinct methods: *direct summarization* and *reason-then-summarize*, which are depicted in Figure 3. See Appendix C for detailed prompts in the supplementary material [50].

Direct Summarization Direct summarization (DirectSumm) enables the LLM to jointly perform table reasoning and summarization in an end-to-end manner. We first present the linearized structure of table formats as outlined in Section 3, aiming to improve the LLM’s understanding of tabular data. Given the query and linearized tables, we then prompt the LLM to perform table reasoning across multiple tables to identify query-relevant facts implicitly. i.e., these facts are not explicitly generated by the LLM. It is worth noting that the facts include numerals and entities that address the information needs of the query. Furthermore, we adopt chain-of-thought (CoT) style prompting methods [14, 38] using the directive “Let’s think step by step”. This enhances the LLM’s reasoning ability. Lastly, we instruct the LLM to synthesize the query with the identified query-relevant facts into a comprehensive summary, while adhering to constraints such as summary length.

Reason-then-Summarize Inspired by the established *retrieve-then-generate* paradigm [11, 19], we introduce a novel approach termed *reason-then-summarize* (Reason-then-Summ). This approach tackles the proposed task in sequential phases. The first phase exclusively focuses on enabling the LLM to perform table reasoning across multiple tables based on the query and linearized tables. As a result, the LLM identifies and extracts query-relevant facts from the tables. In the second phase, we prompt the LLM to synthesize a comprehensive summary based on the query and the previously extracted query-relevant facts.

4 Dataset Construction

In order to validate the effectiveness of our proposed method, we create a novel dataset specifically tailored for this query-focused multi-table summarization task. This section details the dataset construction process, including data annotation and quality verification.

Source Data We build our dataset on top of the Spider dataset [41]. The Spider dataset, originally designed for semantic parsing and text-to-SQL tasks, comprises 10,181 natural language queries. These queries are paired with complex SQL queries and one or more tables from various relational databases. In this work, we construct input queries and tables from original textual queries and their corresponding tables to develop our dataset. We utilize SQL queries for data annotation, as detailed in the following subsection. Additionally, our analysis shows that over 50% of queries in the original dataset corresponded to only a single table. This high proportion of single-table inputs raises concerns about the dataset’s efficacy to provide sufficient challenges for multi-table scenarios. To address this, we selectively down-sampled these single-table examples to allow multi-table examples to predominate in the dataset. Since the test set from the original dataset is not publicly available, we randomly allocated 10% of the original training set to serve as our validation set, with the remaining 90% forming the new training set. The original validation set was repurposed as our test set. This results in 3,871 training, 430 validation, and 608 test examples, forming the basis for our dataset.

4.1 Data Annotation

LLMs as Data Annotators The objective of data annotation is to produce high-quality, comprehensive, and accurate summaries tailored to the associated input queries. Prior research in the field of query-focused summarization from single tables, such as QTSUMM [51], has predominantly depended on human experts to annotate summaries based on an input query and an input table. This reliance is primarily due to the complex table reasoning, making summary annotation a challenging task. However, manual annotations are not only time-consuming but also costly. Recent studies [7, 10, 42, 52] have revealed that LLMs can match the annotation quality of crowd-sourced workers while being significantly more cost-effective and efficient. Motivated by these findings, we have employed LLMs as data annotators for our summary annotation, namely LLM_{Anno} . Although LLM_{Anno} does not yet match the table reasoning capabilities of human experts, we have designed a simplified table-to-text annotation task avoiding complex table reasoning. Specifically, for each textual query in our dataset, we follow Pal et al. [33] to first extract the *output table* by executing the corresponding SQL query over the associated relational database. It is worth noting that the execution table contains query-relevant facts/entities required to construct a summary. Subsequently, rather than relying on input tables, we use the execution table as the basis for our annotation. We also have identified that relying solely on

Dataset	Statistics			Reasoning	
	# Example	# Table per Example	# Words in Summary	Numeric	Multi-Table
ROTOWIRE [17]	4,953	1.0	337.1	✗	✗
SciGen [28]	1,338	1.0	116.0	✓	✗
NumericNLG [36]	1,355	1.0	94.2	✓	✗
QTSUMM [51]	7,111	1.0	68.0	✓	✗
QFMTS	4,909	1.83	58.5	✓	✓

Table 1: Comparisons between our dataset and existing table-to-text generation datasets. Our dataset is the only one tailored for query-focused multi-table summarization, supporting both numeric and multi-table reasoning.

the output table frequently results in summaries that lack essential contextual information. Our observations indicate that this missing context can effectively be retrieved directly from the input query. For instance, consider the example output table below:

semester_name	semester_id
summer 2010	2

It lacks contextual information to confirm that the `Summer 2010` semester has the *most registered students* in response to the corresponding query “*What is the semester in which most students registered? Show both the name and the ID.*”. To address this, we have incorporated the given query as a supplementary input for summary annotation. This strategy enhances the overall comprehensiveness and relevance of the annotated summaries. In this work, we employ `gpt-3.5-turbo-0613` as LLM_{Anno} via the public OpenAI API.

Instruction Design The effectiveness of the instruction prompt is crucial in determining the quality of annotated summaries. To this end, we carefully design the instruction to ensure the summary quality, in which the structure of the instruction is shown in Prompt 4.1. The prompt comprises three components: a comprehensive annotation guideline, few-shot demonstrations, and input data. The annotation guideline outlines the expected discourse structure and the summary’s length requirements. We have found that a more precise guideline significantly improves generation quality. To provide further clarity to LLM_{Anno} , we manually write summaries for a few examples as few-shot demonstrations (we use 5-shot in our experiments). Lastly, we include both the input query and the execution table directly in the prompt, specifically requesting LLM_{Anno} to write a summary based on these inputs. Similarly, we leverage table serialization as described in Section 3 to obtain the linearized execution table.

Prompt 4.1: Summary Annotation

Instruction: A comprehensive annotation guideline.

Demonstrations:

Few-shot human-written demonstrations.

Query: {*Input query*}

Table: {*Linearized execution table*}

4.2 Dataset Analysis

Our dataset comprises a total of 4,909 examples, segmented into 3,871 training examples, 430 validation examples, and 608 test examples. The dataset is composed of 32.8% single-table examples, 52.6% double-table examples, and 14.6% examples that incorporate three or more tables. Notably, more than 67% of the examples include at least two tables, highlighting the dataset’s efficacy in facilitating

research in multi-table scenarios. We present a comparative analysis of our dataset against existing table-to-text generation datasets in Table 1. Our dataset averages approximately two input tables per example, in contrast to the prevailing datasets which predominantly focus on single-table scenarios. The summaries from our dataset are sufficiently informative, with an average length of 58.5 words, aligning closely with the norms of existing datasets. Additionally, our dataset is characterized by a rich variety of operations. It includes basic numeric operations such as *sum* and *average*, and extends to more complex multi-table operations like *join* and *union*, which are absent in the QTSUMM dataset, as it is tailored exclusively towards single-table contexts.

4.3 Quality Verification and Control

To assess the quality of annotated summaries effectively, we develop a comprehensive evaluation encompassing both automated and manual evaluations. We define three primary desiderata for quality verification:

- **Faithfulness:** Each statement within the summary must be factually consistent with the facts presented in the execution table.
- **Completeness:** The summary should address all information needs in the user query, representing all facts from the execution table.
- **Fluency:** The summary needs to be articulate, clear, and easily understandable for human readers.

In our experiments, we utilize standard sequence similarity metrics to evaluate completeness. Given the lack of definitive metrics for faithfulness and fluency, we rely on human evaluations to assess these aspects. The results are presented in Table 2.

Split	Automated Evaluation	Human Evaluation [†]	
	Completeness* (%)	Faithfulness	Fluency
Training	91.45	0.98	4.73
Validation	91.75	0.98	4.81
Test	90.75	0.96	4.68
Total	91.48	0.97	4.74

Table 2: Quality evaluation of our dataset. *represents that the completeness is quantified using average ROUGE-L recall scores. †represents that we randomly sample 100 examples from the training, validation, and test set to measure faithfulness (0-1) and fluency (1-5), respectively.

Automated Evaluation We initially assess the completeness of the annotated summary solely based on the execution table. However, as delineated in Subsection 4.1, the execution table, by itself, proves inadequate for a comprehensive evaluation of summary completeness due to its limited contextual information. Consequently, we extend our

evaluation by incorporating the query with the table, thereby enabling a more robust measure of completeness. Specifically, we extract facts, including numerals and entities, from the execution table. These facts are then combined with the query to construct a reference sequence. The completeness of the annotated summary is estimated against this reference sequence. In this work, we quantify the completeness using the lexical similarity metric ROUGE-L [20], a standard metric in table-to-text generation assessments [22, 51]. As our primary focus lies in assessing the presence of information from the query and the execution table within the summary, we focus on the recall scores of ROUGE-L. As shown in Table 2, the ROUGE-L recall scores exceed 90, affirming that the annotated summaries proficiently include not only the facts from the corresponding execution tables but also the contextual information from the queries.

Human Evaluation To assess the faithfulness and fluency of annotated summaries, we randomly select 100 examples from each of the training, validation, and test sets. We engage three annotators, each proficient in SQL and English, to evaluate the summaries in relation to the corresponding SQL queries, input tables, and execution tables. Annotators assign a binary label to assess faithfulness, a common method in table-to-text generation tasks [3, 51]. Summaries that accurately represent the execution tables without any hallucinated content are labeled as 1, while those that do not are labeled as 0. The annotators are also provided with the query and input tables to enhance their understanding and judgment of the summaries. Following the methodology described by Zhao et al. [51], we measured fluency using a 5-point Likert scale, ranging from 1 (least fluent) to 5 (most fluent). The average score from the three annotators determined the faithfulness and fluency rating for each summary. Table 2 presents the results: the summaries achieved an average faithfulness score of 0.97 and a fluency score of 4.74, indicating that over 97% of the summaries are faithful to the corresponding execution tables and are deemed sufficiently fluent by the annotators. To measure the inter-annotator agreement, we employed the Fleiss Kappa scores [8], achieving Kappa scores of 0.97 for faithfulness and 0.80 for fluency. These scores indicate almost perfect agreement and substantial agreement, respectively.

Human Post-Correction Despite our quality verification confirming the high quality of the dataset, we acknowledge the need for further correction of the validation and test sets to ensure their accuracy, as they play a critical role in selecting optimal model checkpoints and measuring model performance, respectively. Additionally, biases may arise from using output summaries produced by LLM_{Anno} to construct these sets. This risk is particularly pronounced if the same LLM_{Anno} is employed as the baseline model, potentially leading to artificially enhanced performance results. To address these concerns, we have implemented a rigorous post-correction process on the annotated summaries within both sets. This involves a detailed manual review to identify and rectify any missing information or hallucinated content in the summaries, based on the corresponding query and execution table. Furthermore, we have undertaken to rephrase each summary in a manner that more closely resembles human expression, thereby reducing potential biases. See Appendix A for more details in the supplementary material [50].

5 Experiment

In this section, we outline baseline models selected for performance comparison. We then provide the implementation details of the baseline models and our proposed method. Lastly, we describe the evaluation protocols employed in our experiments.

5.1 Baseline Models

In this work, we conduct experiments to evaluate two distinct neural network architectures: encoder-decoder models and decoder-only LLMs. Given that encoder-decoder models typically have significantly fewer parameters than LLMs, we fine-tune these models using our training dataset to facilitate a fair comparison. In contrast, we utilize the LLMs as backbone models for our proposed methods without updating their parameters in the experiments. We benchmarked against the following state-of-the-art models:

BART [18] represents a pre-trained encoder-decoder architecture known for its efficacy in text generation tasks. We have fine-tuned two variants of BART, namely `bart-base` with 139 million parameters and `bart-large` with 406 million parameters [18], which are referred to as BART-base-FT and BART-large-FT, respectively.

TAPEX [24] is a table-to-text generation model, trained using a large-scale synthetic dataset that includes executable SQL queries and their corresponding outputs. In our experiments, we employ the version that utilizes `bart-large` as the backbone.

OmniTab [13] is based on the same architecture as TAPEX but has been additionally trained on a synthetic dataset designed for table question answering tasks. Like TAPEX, our implementation leverages `bart-large` as its backbone.

MultiTab [33] is a table-to-text generation model that has been additionally trained on a synthetic, multi-table question answering dataset. In our experiments, we utilize `bart-base` as its backbone, as the released version exclusively supports this backbone.

Llama-2 [37] includes a set of open-source LLMs that have been pre-trained across vast datasets. We explore two adaptations of Llama-2. First, we fine-tune it on our dataset, namely Llama-2-FT. Second, we utilize llama-2 as the backbone for our approach.

GPT [30, 31] comprises a family of LLMs developed by OpenAI, demonstrating their remarkable text generation capabilities across numerous tasks. In our experiments, we deploy versions `gpt-3.5-turbo-0613` and `gpt-4-0613`, applying these models as the backbones for our approach.

5.2 Implementation Details

Fine-Tuning We fine-tune encoder-decoder models on the training set using the AdamW optimizer [27]. This is conducted over 32 epochs with a learning rate of $1e^{-4}$, batch size of 256, and the maximum sequence length of 1024. For fine-tuning Llama-2 models, we adopt the QLoRA algorithm [5] to fine-tune `llama-2-chat-7B` due to the computation restrictions. The maximum sequence length is 4096. Input sequences for the fine-tuned models were composed by concatenating the query with all linearized input tables. For instance, the final input sequence for an example with k tables is represented as `query [table1] . . . [tablek]`, where `[tablei]` is the linearized representation of the i -th input table. Model performance is evaluated by selecting the best checkpoints based on the loss from the validation set. All experiments are conducted on a single A6000 GPU.

Few-Shot Prompting We prepend 3-shot demonstrations into the prompts to facilitate in-context learning [1] for our method. We set the temperature, top-p, and maximum output tokens to 0.1, 0.95 and 400, respectively. Due to budget constraints, we only report the results of the Reason-then-Summ method with the backbone of GPT-4.

Model	Backbone	Text-Based Metric			Table-Based Metric	
		SacreBLEU	ROUGE-L	BERTScore	STR-EM	PARENT
<i>Fine-tuned</i>						
BART-base-FT	BART-base	39.74	63.14	62.38	24.63	13.51
BART-large-FT	BART-large	43.40	64.84	66.06	33.54	18.66
TAPEX-FT	BART-large	43.99	65.12	66.43	38.78	22.16
MultiTab-FT*	BART-base	44.41	65.68	67.13	42.70	24.83
OmniTab-FT	BART-large	45.58	67.19	68.76	44.60	26.46
Llama-2-FT	Llama-2-7B	54.06	71.82	73.66	61.71	28.69
<i>Our Prompting-based</i>						
DirectSumm	Llama-2-7B	12.49	32.63	21.85	45.32	7.45
DirectSumm	GPT-3.5	33.58	57.02	60.18	53.93	22.21
Reason-then-Summ	Llama-2-7B	16.45	37.13	25.13	48.16	12.17
Reason-then-Summ	GPT-3.5	40.84	62.68	64.98	56.24	24.36
Reason-then-Summ	GPT-4	42.32	64.36	67.36	66.83	32.37

Table 3: Summarization performance of our approaches with various backbones and fine-tuned models on the test set of our dataset. “FT” stands for the fine-tuned version of the corresponding model. The best results are highlighted in **bold**. *represents that the released version only supports the backbone of BART-base.

Model	Backbone	Faithfulness	Fluency
OmniTab-FT	BART-large	0.19	4.79
Reason-then-Summ	GPT-3.5	0.28	4.72
Reason-then-Summ	GPT-4	0.56	4.84

Table 4: Human evaluations of representative models on the test set. Three expert annotators are recruited to evaluate 100 random examples for each model. The best results are in **bold**.

Model	Number of Tables					
	Single-Table			Multi-Table		
	R-L	BSc	PA	R-L	BSc	PA
OmniTab-FT (BART-large)	70.73	72.11	37.73	65.67	67.20	21.56
Reason-then-Summ (GPT-3.5)	64.95	67.04	28.92	61.61	64.18	19.29
Reason-then-Summ (GPT-4)	66.03	69.23	38.11	62.58	64.98	29.87

Table 5: Comparisons between single-table and multi-table examples on the test set. R-L, BSc, and PA stand for ROUGE-L, BERTScore, and PARENT, respectively.

5.3 QFMTS Evaluation

To assess model performance, we employ automated and human evaluations. For automated evaluation, we measure the quality of the generated summary based on the corresponding reference summary and the execution table, as detailed in Subsection 4.1. For human evaluation, we focus on evaluating two key aspects: *faithfulness* and *fluency*.

Text-Based Automated Evaluation We first evaluate the quality of a generated summary w.r.t. the corresponding reference textual summary by estimating the similarity between them in general aspects, such as fluency and accuracy. Following Zhao et al. [51], we adopt two lexical-based metrics, SacreBLEU [34] and ROUGE-L (longest common sub-sequences) [21], along with a semantic-based metric, BERTScore [44]. We report the F1 versions for both ROUGE-L and BERTScore. We use `deberta-xl-large-mnli` [9] as the backbone for BERTScore.

Table-Based Automated Evaluation In contrast to text-based evaluations, table-based evaluations focus more on specific aspects, such

as completeness and faithfulness of the generated summary. To assess these aspects, we employ two metrics: String Exact Match (STR-EM) [35] and PARENT [6]. STR-EM quantifies the proportion of facts or entities from the execution table that are accurately represented in the generated summary. PARENT evaluates summary completeness by integrating both the reference summary and the execution table. PARENT has demonstrated a significant correlation with human judgments.

Human Evaluation In addition to the automated evaluations, we conduct human evaluations, specifically targeting *faithfulness* and *fluency*. These evaluations follow the detailed annotation guidelines specified in Subsection 4.3. For each model, 100 generated summaries were randomly sampled from the test set, and their quality was assessed by three expert annotators.

6 Results and Analysis

Main Results We present the summarization performance of various models in Table 3. We observe that our Reason-then-Summ method markedly outperforms the DirectSumm approach in both text-based and table-based evaluations. For instance, when we use GPT-3.5 as the backbone, Reason-then-Summ surpasses DirectSumm by about 5 points regarding BERTScore (64.98 vs. 60.18). This is because the Reason-then-Summ method tackles the task into sequential phases of table reasoning and summarization. This specialized focus on multi-table reasoning in the initial phase enables the generation of more query-relevant facts and entities, thus enhancing the LLM’s reasoning ability. Consequently, the subsequent summarization task benefits from generating more relevant facts, yielding superior performance compared to the DirectSumm approach, which jointly addresses table reasoning and summarization tasks. Additionally, our results indicate that baseline models fine-tuned on our training dataset largely surpass the DirectSumm methods employing GPT-3.5 or Llama-2 as the backbone in text-based evaluations. In comparison, our Reason-then-Summ approach using GPT-3.5 or GPT-4 as the backbone demonstrates competitive performance alongside the BART-large-FT and MultiTab-FT. This also indicates the effectiveness of Reason-then-Summ. Furthermore, the Llama-2-FT model achieves the highest performance. Particularly in text-based evaluation, when compared to the Reason-then-Summ with the backbone

Llama-2, the fine-tuned Llama-2-FT shows significant improvements, underscoring our dataset’s efficacy as a robust training resource for query-focused multi-table summarization scenarios. Conversely, in table-based evaluations, however, the trend reverses. For instance, our Reason-then-Summ method employing GPT-4 as the backbone, largely outperforms the leading Llama-2-FT model. This discrepancy indicates that while smaller, fine-tuned models may produce plausible summaries, they lack proficiency in table reasoning across multiple tables to gather query-relevant facts, thereby leading to inferior performance on table-based metrics. In comparison, our Reason-then-Summ method produces more facts relevant to the queries, demonstrating superior performance in terms of table reasoning ability.

Human Evaluation Table 4 illustrates the results of sampled human evaluation on the test set. Despite lower scores on automated text-based evaluation metrics such as SacreBLEU and ROUGE scores, our Reason-then-Summ method, which integrates GPT-4, significantly outperforms the fine-tuned OmniTab-FT in human evaluations, particularly in terms of faithfulness (0.56 vs. 0.19). This disparity highlights the superior reasoning ability of our methods over fine-tuned baseline models. Furthermore, our findings reveal a mismatch between text-based automated metrics and human evaluations, aligning with the observations made by Zhao et al. [51]. In contrast, our table-based metrics demonstrate a strong correlation with human judgments concerning faithfulness. This indicates that table-based evaluations are complementary to text-based evaluations, enabling a more comprehensive evaluation for system performance comparison.

Single- vs. Multi-Table To enhance our understanding of the challenges presented in multi-table scenarios, we conducted a performance comparison between single-table and multi-table inputs from our test set, as illustrated in Table 5. It is worth noting that approximately 30% of the examples in our dataset are characterized by single-table inputs. Our analysis reveals that the presence of multiple input tables significantly deteriorates the performance across all evaluated metrics for every model tested. This decline in performance is particularly significant for the smaller OmniTab-FT, whereas it is least noticeable for our method utilizing GPT-4 as the backbone. For instance, considering PARENT scores, the decrease observed with OmniTab-FT is approximately 16 points, moving from 37.73 to 21.56. In contrast, our Reason-then-Summ shows a more modest decrease of about 8 points, dropping from 38.11 to 29.87. This reduction is nearly half that observed with OmniTab-FT. These findings suggest that while multi-table reasoning poses greater challenges compared to single-table scenarios, increased model capacity can effectively narrow this performance gap.

Qualitative Analysis To enhance our understanding of the strengths of our approach and challenges within the task, we conduct a manual analysis of the summaries generated by Reason-then-Summ with the backbone of GPT-3.5 on the test set, including success and failure cases. We observe that our method successfully performs arithmetic and multi-table operations in some cases. A success case illustrates the strengths of our approach. For the query “Which employee received the most awards in evaluations? Give me the employee name.” over two input tables:

Employee			Evaluation		
ID	Name	Age	ID	Year _awarded	Bonus
1	George Chuter	23	1	2011	3000
2	Lee Mears	29	2	2015	3200
...	1	2016	2900
...

With the reference summary “The recipient of the most awards in evaluations is *George Chuter*.”, our method reasons over the 2 tables, performing complex table operations, such as *count* and *join*. Specifically, the method finds two records of awards of *George Chuter* in the table *Evaluation* and aggregates the total number of awards. After joining the two tables, the method accurately identifies *George Chuter* as the person with the most awards, generating “The employee who received the most awards in evaluations is *George Chuter*.”

A failure case illustrates the challenges of multi-table scenarios. Consider the query “What are the names of all European countries with at least 3 manufacturers?” over three input tables:

Continents		Countries		
Cont Id	Continent	Country Id	Country Name	Cont- inent
1	America	2	Germany	2
2	Europe	3	France	2
3	Asia	1	USA	1
4	Africa	8	Korea	3
5	Australia

Car Makers

Id	Maker	Full Name	Country
2	Volkswagen	Volkswagen	2
3	bmw	BMW	2
...
7	citroen	Citroen	3
...
14	opel	Opel	2
15	peugeot	Peugeot	3
16	renault	Renault	3
...
22	kia	Kia Motors	8

With the reference summary “There are 2 European countries with at least 3 manufacturers. The names of these countries are *France* and *Germany*.”, in which the correct country names are marked in blue. The method incorrectly generates “There are 2 European countries which have at least 3 manufacturers. Their names are *France* and *Korea*.”, in which the incorrect country name is marked in red. Even though this generated summary exhibits a high degree of *fluency*, it is only partially *faithful* and *complete* due to the incorrect inclusion of *Korea*, a country which is not located in Europe. This case exemplifies the complexity and challenges of multi-table operations since the proposed approach struggles to combine information from the three tables based on the corresponding column headers, *Country ID* and *Cont Id*.

7 Conclusion

In conclusion, this work addresses the shortcomings of current table summarization techniques through the introduction of an innovative method for query-focused multi-table summarization. Our proposed method leverages user queries and analyzes multiple tables to generate summaries that directly cater to users’ information needs. Additionally, we make a significant contribution to the field by presenting a comprehensive dataset tailored explicitly for this query-focused multi-table summarization task, thereby enabling further research. Extensive evaluations conducted demonstrate the superior performance of our method compared to existing baselines, underscoring the complexities associated with accurate summarization in the context of intricate table reasoning. Overall, our work not only propels advancements in query-focused multi-table summarization but also offers valuable insights to guide future exploration and development in this field.

Acknowledgments

This research was supported by the China Scholarship Council (CSC) under grant number 202008440470.

References

- [1] T. B. Brown, B. Mann, and et al. Language Models are Few-Shot Learners. In *NeurIPS*, 2020.
- [2] W. Chen, J. Chen, Y. Su, Z. Chen, and W. Y. Wang. Logical natural language generation from open-domain tables. In *ACL*, pages 7929–7942, 2020.
- [3] W. Chen, H. Wang, and et al. TabFact: A large-scale dataset for table-based fact verification. In *ICLR*, 2020.
- [4] H. T. Dang. DUC 2005: Evaluation of question-focused summarization systems. In *Proceedings of the Workshop on Task-Focused Summarization and Question Answering*, pages 48–55, 2006.
- [5] T. Dettmers, A. Pagnoni, A. Holtzman, and L. Zettlemoyer. QLoRA: Efficient finetuning of quantized LLMs. *CoRR*, abs/2305.14314, 2023. doi: 10.48550/arXiv.2305.14314.
- [6] B. Dhingra, M. Faruqui, A. Parikh, M.-W. Chang, D. Das, and W. Cohen. Handling divergent reference texts when evaluating table-to-text generation. In *ACL*, pages 4884–4895, 2019.
- [7] B. Ding, C. Qin, L. Liu, Y. K. Chia, B. Li, S. Joty, and L. Bing. Is GPT-3 a good data annotator? In *ACL*, pages 11173–11195, 2023.
- [8] J. L. Fleiss. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378, 1971.
- [9] P. He, X. Liu, J. Gao, and W. Chen. Deberta: Decoding-Enhanced Bert with Disentangled Attention. In *ICLR*, 2021.
- [10] X. He, Z. Lin, Y. Gong, A.-L. Jin, H. Zhang, C. Lin, J. Jiao, S. M. Yiu, N. Duan, and W. Chen. AnnoLLM: Making large language models to be better crowdsourced annotators. *CoRR*, abs/2303.16854, 2023.
- [11] G. Izacard and E. Grave. Leveraging passage retrieval with generative models for open domain question answering. In *EACL*, pages 874–880, 2021.
- [12] P. Jain, A. Laha, K. Sankaranarayanan, P. Nema, M. M. Khapra, and S. Shetty. A mixed hierarchical attention based encoder-decoder approach for standard table summarization. In *NAACL-HLT*, pages 622–627, 2018.
- [13] Z. Jiang, Y. Mao, P. He, G. Neubig, and W. Chen. OmniTab: Pretraining with natural and synthetic data for few-shot table-based question answering. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 932–942. Association for Computational Linguistics, 2022.
- [14] T. Kojima, S. S. Gu, M. Reid, Y. Matsuo, and Y. Iwasawa. Large language models are zero-shot reasoners. In *NeurIPS*, 2022.
- [15] S. Kulkarni, S. Chammas, W. Zhu, F. Sha, and E. Ie. AQuaMuSe: Automatically generating datasets for query-based multi-document summarization. *CoRR*, abs/2010.12694, 2020.
- [16] M. T. R. Laskar, M. Rahman, I. Jahan, E. Hoque, and J. X. Huang. CQ-SumDP: A ChatGPT-Annotated resource for query-focused abstractive summarization based on debaterpedia. *CoRR*, abs/2305.06147, 2023.
- [17] R. Lebre, D. Grangier, and M. Auli. Neural text generation from structured data with application to the biography domain. In *EMNLP*, pages 1203–1213, 2016.
- [18] M. Lewis, Y. Liu, N. Goyal, M. Ghazvininejad, A. Mohamed, O. Levy, V. Stoyanov, and L. Zettlemoyer. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *ACL*, pages 7871–7880, 2020.
- [19] P. S. H. Lewis, E. Perez, and et al. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *NeurIPS*, 2020.
- [20] C.-Y. Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81. Association for Computational Linguistics, 2004.
- [21] C.-Y. Lin and E. Hovy. Automatic evaluation of summaries using N-gram co-occurrence statistics. In *NAACL*, pages 150–157, 2003.
- [22] Y. Lin, T. Ruan, J. Liu, and H. Wang. A survey on neural data-to-text generation. *IEEE Transactions on Knowledge and Data Engineering*, 2023.
- [23] A. Liu, H. Dong, N. Okazaki, S. Han, and D. Zhang. PLOG: Table-to-logic pretraining for logical table-to-text generation. In *EMNLP*, 2022.
- [24] Q. Liu, B. Chen, J. Guo, M. Ziyadi, Z. Lin, W. Chen, and J.-G. Lou. TAPEX: Table pre-training via learning a neural SQL executor. In *ICLR*, 2022.
- [25] S. Liu, J. Cao, R. Yang, and Z. Wen. Long text and multi-table summarization: Dataset and method. In *EMNLP Findings*, pages 1995–2010, 2022.
- [26] Y. Liu, Z. Wang, and R. Yuan. QuerySum: A multi-document query-focused summarization dataset augmented with similar query clusters. In *AAAI*, pages 18725–18732, 2024.
- [27] I. Loshchilov and F. Hutter. Decoupled weight decay regularization. In *ICLR*. OpenReview.net, 2019.
- [28] N. S. Moosavi, A. Rücklé, D. Roth, and I. Gurevych. SciGen: A dataset for reasoning-aware text generation from scientific tables. In *NeurIPS*, 2021.
- [29] L. Nan, C. Hsieh, and et al. FeTaQA: Free-form table question answering. *Transactions of the Association for Computational Linguistics*, 10:35–49, 2022.
- [30] OpenAI. GPT-4 technical report. *CoRR*, abs/2303.08774, 2023.
- [31] L. Ouyang, J. Wu, and et al. Training language models to follow instructions with human feedback. In *NeurIPS*, 2022.
- [32] V. Pal, E. Kanoulas, and M. Rijke. Parameter-efficient abstractive question answering over tables or text. In *Proceedings of the Second DialDoc Workshop on Document-Grounded Dialogue and Conversational Question Answering*, pages 41–53, 2022.
- [33] V. Pal, A. Yates, E. Kanoulas, and M. de Rijke. MultiTabQA: Generating tabular answers for multi-table question answering. In *ACL*, pages 6322–6334, 2023.
- [34] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu. Bleu: A method for automatic evaluation of machine translation. In *ACL*, pages 311–318, 2002.
- [35] I. Stelmakh, Y. Luan, B. Dhingra, and M.-W. Chang. ASQA: Factoid questions meet long-form answers. In *EMNLP*, pages 8273–8288, 2022.
- [36] L. H. Suadaa, H. Kamigaito, K. Funakoshi, M. Okumura, and H. Takamura. Towards table-to-text generation with numerical reasoning. In *ACL*, pages 1451–1465, 2021.
- [37] H. Touvron, L. Martin, and et al. Llama 2: Open foundation and fine-tuned chat models. *CoRR*, abs/2307.09288, 2023. doi: 10.48550/arXiv.2307.09288.
- [38] J. Wei, X. Wang, and et al. Chain-of-thought prompting elicits reasoning in large language models. In *NeurIPS*, 2022.
- [39] Y. Xu and M. Lapata. Coarse-to-fine query focused multi-document summarization. In *EMNLP*, pages 3632–3645, Nov. 2020.
- [40] Y. Xu and M. Lapata. Generating query focused summaries from query-free resources. In *ACL*, pages 6096–6109, 2021.
- [41] T. Yu, R. Zhang, and et al. Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-SQL task. In *EMNLP*, pages 3911–3921, 2018.
- [42] Y. Yu, Y. Zhuang, and et al. Large language model as attributed training data generator: A tale of diversity and bias. *CoRR*, abs/2306.15895, 2023.
- [43] S. Zhang, Z. Dai, K. Balog, and J. Callan. Summarizing and exploring tabular data in conversational search. In *SIGIR*, pages 1537–1540, 2020.
- [44] T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, and Y. Artzi. BERTScore: Evaluating Text Generation with BERT. In *ICLR*. OpenReview.net, 2020.
- [45] W. Zhang, S. Vakulenko, T. Rajapakse, and E. Kanoulas. Scaling up query-focused summarization to meet open-domain question answering. *ArXiv preprint*, abs/2112.07536, 2021.
- [46] W. Zhang, S. Vakulenko, T. Rajapakse, Y. Xu, and E. Kanoulas. Tackling query-focused summarization as a knowledge-intensive task: A pilot study. *Gen-IR workshop at SIGIR*, 2023.
- [47] W. Zhang, M. Aliannejadi, J. Pei, Y. Yuan, J.-H. Huang, and E. Kanoulas. A comparative analysis of faithfulness metrics and humans in citation evaluation. *LLM4Eval workshop at SIGIR*, 2024.
- [48] W. Zhang, M. Aliannejadi, Y. Yuan, J. Pei, J.-H. Huang, and E. Kanoulas. Towards fine-grained citation evaluation in generated text: A comparative analysis of faithfulness metrics. In *INLG*, 2024.
- [49] W. Zhang, J.-H. Huang, S. Vakulenko, Y. Xu, T. Rajapakse, and E. Kanoulas. Beyond relevant documents: A knowledge-intensive approach for query-focused summarization using large language models. In *ICPR*, 2024.
- [50] W. Zhang, V. Pal, J.-H. Huang, E. Kanoulas, and M. de Rijke. Qfmts: Generating query-focused summaries over multi-table inputs. *arXiv preprint arXiv:2405.05109*, 2024.
- [51] Y. Zhao, Z. Qi, and et al. QTSumm: Query-focused summarization over tabular data. In *EMNLP*, pages 1157–1172, 2023.
- [52] Y. Zhao, H. Zhang, and et al. Investigating table-to-text generation capabilities of large language models in real-world information seeking scenarios. In *EMNLP*, pages 160–175, 2023.