



Replacing Multi-Step Assembly of Data Preparation Pipelines with One-Step LLM Pipeline Generation for Table QA

Fengyu Li
Zhejiang University
Hangzhou, China
fengyuli@zju.edu.cn

Junhao Zhu
Zhejiang University
Hangzhou, China
zhujunhao@zju.edu.cn

Kaishi Song
Zhejiang University
Hangzhou, China
sks@zju.edu.cn

Lu Chen
Zhejiang University
Hangzhou, China
luchen@zju.edu.cn

Zhongming Yao
Northeastern University
Shenyang, China
yaozming@stumail.neu.edu.cn

Tianyi Li
Aalborg University
Aalborg, Denmark
tianyi@cs.aau.dk

Christian S. Jensen
Aalborg University
Aalborg, Denmark
csj@cs.aau.dk

ABSTRACT

Table Question Answering (TQA) aims to answer natural language questions over structured tables. Large Language Models (LLMs) enable promising solutions to this problem, with operator-centric solutions that generate table manipulation pipelines in a multi-step manner offering state-of-the-art performance. However, these solutions rely on multiple LLM calls, resulting in prohibitive latencies and computational costs.

We propose Operation-R1, the first framework that trains light-weight LLMs (e.g., Qwen-4B/1.7B) via a novel variant of reinforcement learning with verifiable rewards to produce high-quality data-preparation pipelines for TQA in a single inference step. To train such an LLM, we first introduce a self-supervised rewarding mechanism to automatically obtain fine-grained pipeline-wise supervision signals for LLM training. We also propose variance-aware group resampling to mitigate training instability. To further enhance robustness of pipeline generation, we develop two complementary mechanisms: operation merge, which filters spurious operations through multi-candidate consensus, and adaptive rollback, which offers runtime protection against information loss in data transformation. Experiments on two benchmark datasets show that, with the same LLM backbone, Operation-R1 achieves average absolute accuracy gains of 8.83 and 4.44 percentage points over multi-step preparation baselines, with 79% table compression and a 2.2 $\times$ reduction in monetary cost.

PVLDB Reference Format:

Fengyu Li, Junhao Zhu, Kaishi Song, Lu Chen, Zhongming Yao, Tianyi Li, and Christian S. Jensen. Replacing Multi-Step Assembly of Data Preparation Pipelines with One-Step LLM Pipeline Generation for Table QA. PVLDB, 19(9): 2372 - 2384, 2026.
doi:10.14778/3819518.3819557

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/ZJU-DAILY/Operation-R1.git>.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.
Proceedings of the VLDB Endowment, Vol. 19, No. 9 ISSN 2150-8097.
doi:10.14778/3819518.3819557

1 INTRODUCTION

Table Question Answering (TQA) aims to answer natural language (NL) queries over (semi-)structured tables, spanning spreadsheets, databases, web tables. It empowers non-technical users to explore and extract insights without writing formal query languages such as SQL, enhances analytical productivity, and enables applications in finance [42], healthcare [1], and scientific research [23]. For example, TQA for business intelligence enables users to query sales, customer, and financial tables using natural language (e.g., “Which product had the highest growth in Q3?”), lowering the technical barrier to data analytics.

Given their semantic understanding and instruction-following capabilities Large Language Models (LLMs) have advanced TQA substantially. As illustrated in Fig.1, one line of LLM-based TQA methods employ *direct text-based reasoning* [34, 38]. These methods serialize tables into plain text and answer NL queries in a single step by reasoning directly over the text. However, they scale poorly to large tables because LLMs struggle to reliably identify relevant information within long contexts. Studies show that text-based reasoning deteriorates sharply as table size grows [8].

Another line of methods rely on *Text-to-SQL/code generation* [7, 24, 38, 40]. These methods translate NL queries into compositional executable programs (e.g., SQL or Python) and derive answers through program execution. Although often effective, they are fragile: LLMs easily introduce subtle errors into long, compound programs, causing executions to fail or produce corrupted results.

Motivated by limitations of above reasoning approaches, recent methods adopt *operator-centric* approaches [8, 32], where an LLM assembles a sequence of table manipulation operators (e.g., select, sort_by) to restructure input tables into more compact, question-relevant tables that facilitate the TQA the most. To discover effective table manipulation operations, these methods employ multi-step reasoning, where table manipulation operations are generated step-by-step in the context of input questions, tables, and intermediate orchestration results. This operator-centric approach facilitates more transparent, interpretable reasoning and enables methods that achieve state-of-the-art performance.

Despite their accuracy and interpretability advantages, operator-centric pipelines incur substantially higher latency and computational cost due to repeated LLM invocations. This overhead, which




Figure 1: (a) Direct text-based reasoning, (b) Program-based reasoning, (c) Operation-based multi-agent framework, and (d) Pipeline-specific model.

is avoided by one-step approaches, is the primary barrier to practical deployment of these otherwise desirable operator-specific methods.

To enable the practical deployment of these methods, we introduce Operation-R1, a framework that trains a lightweight LLM to generate a high-quality data preparation pipeline for TQA *in a single inference step*. Inspired by the success of DeepSeek-R1 [10] and TableR1 [36], we employ Reinforcement Learning with Verifiable Rewards (RLVR) to teach compact models (e.g., Qwen3-4B, Qwen-1.7B) to orchestrate well-structured data preparation tailored to each question–table pair.

However, building such a one-step orchestrator introduces several unique challenges not addressed by prior work. Below, we outline these challenges and how they motivate the key design components of Operation-R1.

Challenge I: Designing a data-preparation-aware reward mechanism for pipeline generation. Training a lightweight LLM to generate effective data preparation pipelines requires reliable supervision signals that reflect pipeline quality. However, obtaining such signals is difficult because no ground-truth—or “golden” preprocessing pipelines”—exists for TQA. The only available supervision in typical TQA datasets is final answers to each question, which is insufficient for directly supervising pipeline construction.

A natural surrogate is to evaluate a generated pipeline using an LLM-as-Judge [26]: execute the pipeline, present the transformed table to an LLM, and check whether it can infer the correct answer. Yet this approach has three critical drawbacks: (i) Low accuracy. As shown in our experiments (Table 1, Section 4), even strong LLMs (e.g., Qwen3-8B) achieve under 80% accuracy on WikiTQ when

directly reasoning over tables, which means it may introduce noise into supervision signals. (ii) Sparse feedback. The evaluation yields only a binary reward per pipeline, providing little insight into which operations are helpful or harmful. (iii) High computational cost. Invoking an LLM-as-Judge for every pipeline drastically increases training cost. Thus, a verifiable, fine-grained, and low-cost reward mechanism is essential.

To address this challenge, we introduce a self-supervised reward mechanism tailored to data preparation. We focus on *Cell-Focused QA*, a setting where the answer appears verbatim as a table cell. This property allows us to verify each operation by checking whether it removes or preserves the answer-critical cell. Pipeline quality is then evaluated *operation by operation*, enabling detailed attribution of failures and successes. To complement semantic correctness, we further incorporate an efficiency metric that measures how effectively the pipeline reduces table size. The combined reward provides fine-grained, stepwise feedback that directly reflects each operation’s utility, yielding informative and stable supervision for training the pipeline generator.

Challenge II: Maintaining training stability under fine-grained reward supervision. Although our approach introduces finer-grained supervision, integrating it with existing RLVR algorithms leads to training instability. In standard RLVR workflows, models are trained not with raw reward values but with *normalized relative advantages*, which are computed by comparing a group of model responses for the same input. This normalization, while effective for coarse rewards, inevitably distorts the underlying learning signal. With our more granular pipeline rewards, this distortion

becomes even more pronounced, ultimately misleading model optimization. Our early analysis reveals two sources of distortion: (a) Intra-group low-variance amplification. When reward variance within a sampled group is small, even noise-level differences are artificially amplified by the advantage normalization, causing the model to overfit to random fluctuations rather than meaningful reward differences. (b) Inter-group reward distortion. Conversely, large and informative reward gaps across different groups are flattened after normalization, suppressing important learning signals that should guide the policy toward better pipelines.

To alleviate these issues, we propose a *variance-aware group resampling* strategy that replaces the static grouping used in standard RLVR with a dynamic, variance-sensitive sampling process. The method imposes two constraints—each derived from the two distortion phenomena above—on sampled response groups. If any constraint is violated, the group is discarded and resampled. This ensures that each accepted group provides diverse, high-fidelity supervision signals, thereby preventing noise amplification and stabilizing model training.

Challenge III: Ensuring robustness during preprocessing. Even with a carefully finetuned LLM designed to generate high-quality pipelines, failures are inevitable when the model encounters unusual or previously unseen NL questions. A single faulty operation—such as overly aggressive filtering or an incorrect aggregation—can permanently remove question-critical information and make downstream reasoning impossible. This reveals a core tension between *compression efficiency* and *information preservation*: while aggressive table compression helps reduce reasoning complexity, any loss of essential data invalidates the entire preprocessing pipeline. The challenge, therefore, is to design a robustness mechanism that mitigates generation errors without compromising too much on efficiency and interpretability.

To address this, Operation-R1 incorporates two complementary mechanisms: *Operation Merge* and *Adaptive Rollback*. Operation Merge tackles the nondeterminism inherent in pipeline generation by producing multiple candidate pipelines in parallel and constructing a final, more reliable pipeline through a consensus-based voting algorithm over an operation trie. This filters out spurious or unstable operations that may appear in single-pass outputs. Adaptive Rollback provides runtime safeguards against information loss. When the downstream QA model signals that the processed table lacks sufficient information, the system incrementally reverts to earlier states—i.e., tables resulting from fewer preprocessing operations—ultimately reverting to the original table if necessary.

Together, these mechanisms strike a balance between the efficiency of single-pass preprocessing and the robustness required for real-world TQA. They enable our lightweight framework to reliably handle diverse questions and edge cases without relying on multi-agent coordination or iterative refinement.

The key contributions are summarized as below:

- To our knowledge, we present the first RLVR-based LLM post-training framework that enables lightweight LLMs to generate data preparation pipelines for TQA in a *single inference step*.
- We introduce a self-supervised pipeline reward mechanism based on Cell-Focused QA that provides fine-grained, stepwise feedback for RLVR training. We further propose a variance-aware

group resampling strategy to ensure stable optimization under such fine-grained rewards.

- We develop two complementary mechanisms to balance compression efficiency with information preservation: *Operation Merge*, which filters spurious operations through multi-candidate consensus, and *Adaptive Rollback*, which offers runtime protection against information loss.
- Extensive experiments on WikiTQ and TabFact show that Operation-R1 achieves absolute accuracy improvements of 9.62% and 6.05% over non-preprocessed baselines, respectively, while reducing table size by 79% on average and requiring only a single lightweight model inference.

2 PRELIMINARIES

2.1 Task Formulation

Tabular Question Answering (TQA). Let q be an NL question and T a table with columns (attributes) $C = \{C_1, C_2, C_3, \dots\}$ and rows $\mathcal{R} = \{R_1, R_2, R_3, \dots\}$, where the cell at row i and column j is denoted by $c_{i,j}$. The goal of TQA is to generate a correct answer a to question q based on the content of table T .

Depending on the purpose of the question, TQA typically includes two problem types: (i) *Table-based Question Answering*: Given q and T , the task is to generate a natural language answer a or extract relevant cell values from T . (ii) *Table-based Fact Verification*: Given a statement q and a table T , the task is to determine whether the statement is *supported*, *refuted*, or *unverifiable* given the table.

Question-Aware Data Preparation. Answering NL questions directly over an uncurated table is often ineffective and inefficient. The reasons are twofold: (i) many table entries are irrelevant to the question, forcing the model to sift through large amounts of unnecessary information—the classic “needle in a haystack” problem; and (ii) answers frequently require transformations or computations over the table rather than direct lookup, necessitating intermediate data preparation steps.

Thus, before generating answers, it is beneficial to apply a series of table transformation operations (data preparation operations, or simply *operations*). Formally, let op be an operation that transforms a table T into a table T' , i.e., $T' = op(T)$. Given a pipeline $P = [op_1, op_2, \dots, op_k]$, a curated table is produced as follows:

$$T' = op_k(op_{k-1}(\dots(op_2(op_1(T)))\dots))$$

We consider five data preparation operators. Based on prior studies [8] and our analysis of benchmark and real-world tables, these operators cover most table manipulation needs in TQA. Below, we describe each operator and its parameters.

Select. The select operator removes columns from an input table. Formally, given an input table T and a set of column names C , $select_C(T) = T[C]$ denotes a processed table from T by retaining the only columns in C . This operator is useful when many columns are irrelevant to a question and may distract answering.

Filter. The filter operator selects rows that satisfy a given predicate. Given an input table T , a column name c , a comparison operator $op \in \{=, >, <, \dots\}$, and a threshold value v , The result is denoted $filter_{c,op,v}(T)$. For example, $filter_{Country',=,'USA'}(T)$ retains

rows where the country column cell equals ‘USA’. This operation enables eliminating irrelevant tuples and retaining only those related to a question.

Sort_by. The `sort_by` operator sorts rows based on a specific column and optionally retains the top- k rows. Given an input table T , a column name c , an order $o \in \{\text{asc}, \text{desc}\}$, and an optional parameter k , we denote the result by `sort_byc,o,k( $T$ )`. By default, when k is omitted, all rows are returned, which we denote by `sort_byc,o( $T$ )`. This operation supports reasoning about ranking, extrema, and comparisons.

Group_by. The `group_by` operator groups rows by a specific column and counts the occurrences of each distinct value. Given an input table T and a column name c , we denote the result by `group_byc( $T$ )`, which produces a table with distinct values in column c and their corresponding counts. This operation is especially helpful for counting tasks that LLMs often struggle with.

Add_column. The `add_column` operator adds a new column inferred from existing columns according to a natural language description. Given an input table T , a new column name c' and a description d , we denote the result by `add_columnc',d( $T$ )`. For example, `add_column'Gender','infer genders from the column Name'( $T$ )` adds a Gender column based on column Name. This operator relies on semantic understanding and is implemented by invoking an LLM.

Operator `add_column` differs fundamentally from the other operators. We thus categorize the data preparation operators into two groups: structured operators, whose logic can be implemented by built-in or user-defined functions; and semantic operators that require reasoning over the data.

2.2 RL with Verifiable Rewards

Reinforcement Learning with Verifiable Rewards (RLVR). RLVR [14] is a family of LLM post-training methods in which an LLM is optimized through RL using automatically verifiable feedback as the reward signal. Unlike preference-model-based or human-rated rewards, RLVR rewards are computed by a verifier—an automated, deterministic, and programmatically checkable function that evaluates whether an output satisfies a well-defined success criterion. The verifier may check logical or syntactic constraints, unit-test outcomes, exact or semantic equivalence to reference solutions, or it may use other rule-based correctness conditions. Concretely, given a prompt or input x and a policy model (i.e., an LLM) π_θ producing an output o (which may include a reasoning trace, or a chain-of-thought, and a final answer), a verifier computes a reward $r(o; x)$ (that is either a binary value $\{0, 1\}$ or a continuous value in range $[0, 1]$). RLVR uses an RL algorithm to update the policy model π_θ to maximize expected rewards. A typical optimization objective is formulated as follows:

$$\max_{\pi_\theta} \mathbb{E}_{x \sim \mathcal{D}, o \sim \pi_\theta(\cdot | x)} [r(o; x) - \beta \text{KL}[\pi_\theta(\cdot | x) \| \pi_{\text{ref}}(\cdot | x)]], \quad (1)$$

where π_{ref} is a reference policy (e.g. a pre-trained LLM) and $\beta > 0$ controls the strength of the regularization. RLVR is applicable most straightforwardly in domains where outputs can be verified unambiguously, such as arithmetic reasoning, code synthesis, QA with known gold answers [6, 10, 11].

Group Relative Policy Optimization (GRPO). A member of the RLVR family, GRPO [10, 25] has gained in prominence following the

emergence of DeepSeek-R1. GRPO is a policy gradient algorithm designed to train an LLM in settings, where one can sample multiple candidate outputs per input and assign verifiable rewards to each candidate. Instead of relying on a learned value function or critic, GRPO uses within-group relative comparisons across the sampled outputs for the same input to compute advantages and update the policy. The use of the contrastive reward signal helps stabilize training when faced with sparse or binary rewards.

Formally, given an input x , GRPO samples a group G of candidate outputs from the current (or “old”) policy model π_{old} , i.e., $G = \{o_i\}_{i=1}^{|G|}$. Each output o_i is then scored by a verifier to produce a reward $R_i = r(o_i; x)$. Each output’s advantage is defined as the normalized difference from the group’s mean reward:

$$A_i = \frac{R_i - \mu_G}{\sigma_G}, \quad (2)$$

where $\mu_G = \frac{1}{|G|} \sum_{i=1}^{|G|} R_i$ is the group mean reward, and $\sigma_G = \sqrt{\frac{1}{|G|} \sum_{i=1}^{|G|} (R_i - \mu_G)^2}$ is the group standard deviation.

GRPO optimizes the policy model π_θ (i.e., an LLM) by maximizing the following token-level objective:

$$\begin{aligned} \max_{\pi_\theta} \mathbb{E}_{x \sim \mathcal{D}, \{o_i\}_{i=1}^{|G|} \sim \pi_\theta(\cdot | x)} & \left[\frac{1}{|G|} \sum \frac{1}{|o_i|} \sum_{o_i} \left\{ \right. \right. \\ & \left. \left. \min[\rho_{i,t}^\theta A_{i,t}, \text{clip}(\rho_{i,t}^\theta, 1 - \epsilon, 1 + \epsilon) A_{i,t}] - \beta \text{KL}(\pi_\theta \| \pi_{\text{ref}}) \right\} \right], \quad (3) \end{aligned}$$

where $\rho_{i,t}^\theta = \frac{\pi_\theta(o_{i,t} | x, o_{i,<t})}{\pi_{\text{old}}(o_{i,t} | x, o_{i,<t})}$ is the probability ratio at token position t , and ϵ is the clipping parameter that constrains policy updates. Although RLVR methods are prevalent for LLM post-training, they face problems, including a lack of verifiable rewards and unstable training, as discussed in the introduction, when they are applied to generate a high-quality data preparation pipeline, motivating the Operation-R1 approach.

3 OPERATION-R1

3.1 Framework Overview

To bridge the gap between TQA and data preparation, we propose Operation-R1, a question-oriented data-preparation orchestration framework for high-quality tabular question answering (Figure 2). Given a question and a table, Operation-R1 automatically composes and executes a tailored transformation sequence that reshapes the table to better expose the answer. Unlike prior operator-centric methods that assemble pipelines through repeated LLM invocations, Operation-R1 generates the entire pipeline in a single forward pass, structurally eliminating multi-step orchestration overhead.

As shown in Fig. 2, the offline stage of Operation-R1 aims to produce a TQA-oriented data preparation pipeline generator, built using open-sourced LLMs. To finetune an LLM for the task of QA-oriented data preparation, we propose ORPO (Operation-wise Group Relative Policy Optimization), a novel RVRL algorithm. ORPO features two designs that aim to make the post-training process more effective, efficient, and robust: (i) a self-supervised pipeline reward mechanism, providing finer-grained supervision signals tailored to TQA for model training; and (ii) a variance-aware group




Figure 2: Overview of Operation-R1, including ORPO-based offline training, parallel operator generation, operation-tree merging, and adaptive rollback for reliable tabular QA.

resampling strategy, enhancing the stability of the LLM training in case of reward collapse.

During the online stage, once receiving a natural language question q over a table T , Operation-R1 leverages the data preparation pipeline generator derived from the offline stage to produce a sequence of operators. Each operator is organized as a JSON object containing the operator name, parameters, and an explanation (the usages of the operation). Due to hallucinations and nondeterminism in LLM inference, the Operation-R1 builds N data preparation pipelines concurrently and then produces a consistent pipeline using an operation merge algorithm.

The JSON-represented operators are interpreted as physical operators and executed in the underlying environment (e.g., Python). Structured operators (e.g., select, filter) are compiled into code snippets and executed directly; semantic operators (e.g., add_column) are executed via LLM calls based on NL description. This hybrid execution extends the capabilities of data preparation, expanding the scope of questions that TQA can address.

Eventually, the processed table T' together with natural language question q are fed into a QA model to obtain the desired answer. To avoid the “over-processing” problem where key information is lost in the processing of data preparation, Operation-R1 also supports an adaptive rollback mechanism: When the QA model cannot identify relevant information from the table, the system will revert to the earlier table that contains more information.

In the following, we detail the post-training algorithm of ORPO and then the online question answering process.

3.2 ORPO: Operation-wise Group Relative Policy Optimization

As discussed earlier, there are numerous benefits to using a customized LLM to generate a QA-oriented data preparation pipeline, and GRPO performs well at customizing an LLM for a specific task [10, 39]. The challenge however, is that as for data preparation operation generation requires assessing both final answers correctness and the quality of intermediate operations. We thus propose a low-cost surrogate operation reward mechanism to quantify the quality of generated operations in terms of both correctness and efficiency, forming the foundation for our ORPO training strategy.

3.2.1 Self-supervised Pipeline Reward Mechanism. We first introduce the reward designs for operations in a generated data preparation pipeline in terms of correctness and efficiency.

Operation correctness rewarding. Ideally, a generated operation would be evaluated by comparing its output to a ground-truth preprocessed table; however, such a table cannot be obtained in advance; however, such a table cannot be obtained in advance. We thus instead propose a surrogate operation rewarding mechanism that performs intermediate operation evaluations in a self-supervised manner on a “cell-focused QA task”.

DEFINITION 1 (CELL-FOCUSED QA TASK). Given a table T containing a set of cell values $\{c_{ij}\}$, a question q over T is a cell-focused QA task if every value of the dataset-provided gold answer set $\mathcal{A}(q)$ of question answers is required to appear verbatim as a cell value in table T . Thus, the answer set $\mathcal{A}(q)$ formally satisfies $\forall a_k \in \mathcal{A}(q) (\exists c_{i,j} \in T (a_k = c_{i,j}))$, where “=” denotes exact string match.

Example. The question in Figure 1 (“How long did it take for the New York Americans to win the national cup after 1936?”) is *non-cell-focused* as its answer “17 years” does not appear verbatim in the table. A cell-focused question over the same table would be “What was the Playoffs result in the 1937/38 season?”, whose answer “1st Round” appears as a cell value.

On cell-focused QA tasks, we regard the operation evaluation as a judgment of whether the produced intermediate table contains the answers, which is a proxy for the comparison between the intermediate result and the virtual ground truth.

Given a cell-focused QA task q over an initial table T and its answer set $\mathcal{A}(q)$, an LLM generates a sequence of operations $P = [op_1, op_2, \dots, op_n]$. The produced table of an intermediate operation $op_k \in P$ is denoted by T_k . The correctness of the operation op_k is defined as follows.

$$cr(op_k, T_k; \mathcal{A}(q)) = \mathbb{1}(a = c, \forall a \in \mathcal{A}(q), \exists c \in T_k), \quad (4)$$

where $\mathbb{1}$ is an indicator function. For simplicity, we omit $\mathcal{A}(q)$ from the notation of the operation correctness.

Based on the definition of correctness of a single operation, we propose a cumulative reward function to rate the correctness of a generated pipeline. Given a pipeline with n data preparation operators, the accumulative accuracy reward assigned to the first

k ($k \leq n$) steps in the pipeline is defined as follows.

$$r_{\text{acc}}(k) = \frac{\sum_{i \leq k} cr(op_i, T_i)}{n} \quad (5)$$

This accumulative accuracy reward function assigns higher scores to pipelines where every data preparation operation can retain answers to the question. While it encourages LLMs to generate pipelines, lengthy pipelines are also encouraged, which is not what users want.

Operation efficiency rewarding. To avoid generating excessive data preparation pipelines, in addition to correctness, efficiency is considered as a reward to punish LLM generations that include redundant, useless operations in a pipeline.

Since one of the goals of data preparation is to eliminate irrelevant, noisy data from a given table, we define operator efficiency in terms of a compression ratio that measures reductions in the column and row sizes of the given table by an operation.

Without loss of generality, given an operator op_k from a generated pipeline and its input table T_{k-1} with columns C_{k-1} and rows $\mathcal{R}_{k-1}$, the table produced by op_k is denoted by T_k and has columns C_k and rows $\mathcal{R}_k$. The row-level and column-level compression ratio of operator op_k is defined as $\frac{|\mathcal{R}_k|}{|\mathcal{R}_{k-1}|}$ and $\frac{|C_k|}{|C_{k-1}|}$, respectively.

Similar to the accumulative operator accuracy reward function, we define an accumulative operator efficiency reward function based on the compression ratio. As for an initial table T with columns C and rows $\mathcal{R}$, the accumulative efficiency reward assigned to the first k ($k \leq n$) operations in the pipeline is defined:

$$r_{\text{compress}}(k) = 0.5 \cdot \left(1 - \frac{|\mathcal{R}_k|}{|\mathcal{R}|}\right) + 0.5 \cdot \left(1 - \frac{|C_k|}{|C|}\right) \quad (6)$$

Overall reward function. In addition to the operation correctness and operation efficiency rewards, we adopt a length reward from DAPO [39] to constrain LLMs from generating overlong responses, alleviating “over thinking” [28] as well as reducing latency and cost.

Formally, let o denote an LLM response that contain a data preparation pipeline, and let $|o|$ be its token length. The length reward function on o is defined as follows.

$$r_{\text{length}}(o) = \begin{cases} 0 & |o| \leq L_{\text{max}} - L_{\text{cache}} \\ \frac{(L_{\text{max}} - L_{\text{cache}}) - |o|}{L_{\text{cache}}} & L_{\text{max}} - L_{\text{cache}} < |o| \leq L_{\text{max}} \\ -1 & |o| > L_{\text{max}}, \end{cases} \quad (7)$$

where L_{max} is set according to system constraints such as memory and rollout budget, and L_{cache} is a slack token budget that defines the width of the soft tolerance region before strong penalties apply, thereby discouraging excessively long responses while avoiding abrupt penalties on near-limit samples.

The overall reward function integrates the previous components with tunable weights:

$$r(o, n) = r_{\text{acc}}(n) + \lambda_1 \cdot r_{\text{compress}}(n) + \lambda_2 \cdot r_{\text{length}}(o), \quad (8)$$

where the weights (λ_1 and λ_2) balance operation correctness, operation and output efficiency, each of which are tuned according to performance on each datasets. Here we set $\lambda_1 = \lambda_2 = 0.5$.

3.2.2 LLM Post-training by ORPO. We integrate the self-supervised pipeline rewarding mechanism into the LLM post-training GRPO method to achieve the operation-wise group relative policy optimization method, ORPO. Given a question q and a table T from the

training dataset, ORPO samples a group G of LLM outputs from the current LLM π_{old} , i.e., $G = \{o_i\}_{i=1}^{|G|}$. Each output o_i contains a data preparation pipeline $P_i = [op_1, op_2, \dots, op_{n_i}]$ and each output is then scored by the pipeline reward function $R_i = r(o_i, n_i)$. ORPO optimizes the LLM by maximizing the following token-level objective:

$$\max_{\pi_{\theta}} \mathbb{E}_{(q, T) \sim \mathcal{D}, \{o_i\}_{i=1}^{|G|} \sim \pi_{\theta}(\cdot | q, T)} \left[\frac{1}{|G|} \sum_G \frac{1}{|o_i|} \sum_{o_i} \left\{ \min \left[\rho_{i,t}^{\theta} A_{i,t}^{\text{ORPO}}, \text{clip}(\rho_{i,t}^{\theta}, 1 - \epsilon, 1 + \epsilon) A_{i,t}^{\text{ORPO}} \right] - \beta \text{KL}(\pi_{\theta} \| \pi_{\text{ref}}) \right\} \right], \quad (9)$$

where $\rho_{i,t}^{\theta} = \frac{\pi_{\theta}(o_{i,t} | q, T, o_{i,<t})}{\pi_{\theta_{\text{old}}}(o_{i,t} | q, T, o_{i,<t})}$, and $A_{i,t}^{\text{ORPO}}$ represents a normalized difference from the group mean operation reward, defined as:

$$A_{i,t}^{\text{GRPO}} = \frac{R_i - \mu_r}{\sigma_r + \epsilon}, \quad (10)$$

where μ_r and σ_r are the mean and standard deviation of rewards among the sampled completions o_i in the same batch, and ϵ is a small constant that ensures numerical stability.

Variance-aware group resampling strategy (VGR strategy). Although ORPO introduces finer-grained supervision signals, it also causes more instability in the LLM training, which manifests itself in two forms. The first is *low-variance amplification*: when rewards assigned to different LLM responses in a group are very similar (i.e., the reward variance is close to zero), the group-wise normalization in ORPO amplifies tiny, often noise-level differences into large positive and negative advantages, causing the optimizer to overreact to spurious contrasts between almost indistinguishable trajectories and leading to unstable gradient updates. Second, we observe a *reward distortion problem*: in ORPO, the raw rewards $\{R_i\}$ are transformed into group-wise normalized advantages A_i^{ORPO} (Eq. 10), and the policy update is driven purely by such relative differences instead of by absolute reward values. This means that the model training does not receive real supervision signals. For instance, two vastly different groups of absolute reward scores $[0.9, 0.5, 0.5]$ and $[0.3, 0.1, 0.1]$ will be normalized to the nearly identical advantages $[1.41, -0.70, -0.70]$ and $[1.43, -0.72, -0.72]$. Here training samples with a high reward of 0.9 and with a low reward of 0.3 across two groups will have the same effect (both are regarded as positive signals), consequently misleading the LLM training. This creates a critical vulnerability as highlighted in recent relative optimization studies [17, 30]: the optimization process inadvertently reinforces the “least bad” sample in a sub-optimal group (e.g., the 0.3 reward case) simply because it outperforms its peers locally, while failing to distinguish between relative superiority and absolute correctness.

These two issues are largely absent in traditional GRPO applications such as mathematical reasoning or code synthesis, where binary correctness rewards naturally induce high intra-group variance and provide absolute oracle signals robust to cross-group normalization noise (e.g., “least bad”) [18]. In our setting, data preparation considers both compression efficiency and information preservation, inducing dense, continuous rewarding. Our task further demands precise operator composition—a single erroneous operator causes catastrophic downstream failure—and our light-weight models (1.7B/4B) trained on limited examples ($\sim 5K$) could

be more sensitive to noisy supervision than larger models, making robust supervision construction necessary.

To address these challenges, we propose a *variance-aware group resampling strategy*, which yields dynamic groups that provide diverse, informative training samples based on two group constraints. Consider a group G of LLM outputs containing various data preparation pipelines $G = \{o_i\}_{i=1}^{|G|}$. Each LLM output is assigned with a reward by the reward function, forming a group of rewards $\mathcal{R} = \{R_1, R_2, \dots, R_{|G|}\}$. The group mean reward μ_G is computed by $\mu_G = \frac{1}{|G|} \sum_{i=1}^{|G|} R_i$, and the reward variance σ_G is calculated by $\sigma_G = \frac{1}{|G|} \sum_{i=1}^{|G|} (R_i - \mu_G)^2$. To overcome the two dilemmas discussed above, we have following two constraints to the group G :

- *Variance constraint*: the reward variance of the group has to satisfy $\sigma_G \geq \epsilon$. This constraint avoids the *signal collapse effect*.
- *Quality constraint*: the maximum reward value $R_{\max}$ within the group has to satisfy $R_{\max} \geq \tau$, which guarantees that at least one output provides a high-quality data preparation pipeline. This constraint overcomes the *reward distortion problem*.

Once a group violates any of the constraints, it triggers group resampling *within the same model update step*: the current group is discarded and a new group is generated for the same input before any gradient update is performed. This resampling repeats until a qualifying group is found, after which gradients are computed and the model is updated. Our variance-aware group resampling contributes to collecting a more diverse and higher-quality data preparation pipeline reasoning trajectories as the training data, and our empirical studies in Section 4.4 validates that the resampling strategy enhances stability of LLM training under fine-grained reward regimes and mitigates early-stage reward collapse.

3.3 Operation Merge

In the online phase, the pipeline generator trained by ORPO is used to generate a sequence of data preparation operators tailored to the specific question and given table. To enhance the robustness and consistency of pipeline generation, inspired by [2], Operation-R1 asks the generator producing N pipeline candidates and then returns one of them with the most “votes” as the final data preparation pipeline. The vote for an operator represents how many pipeline candidates include that operator. The more votes a pipeline candidate has, the more reliable operators it contains.

We propose an operation merge algorithm to count the votes and find the most trustworthy pipeline among all candidates efficiently. The core data structure in operation merge is an operation trie built upon all pipeline candidates. In the operation trie, each branch represents a candidate pipeline, and each node in a branch represents an operator in the pipeline. If two candidate pipelines share the same operator sub-sequence, they share the same sub-branch in the operation trie. Each node in the trie has a weight recording how many candidate pipelines contain the operator (*i.e.*, the number of votes). With an operation trie, the problem of searching a pipeline that most LLM orchestrations advocate can be solved by identifying a branch of a trie with the maximum sum of node weights. In case of multiple branches having same node weights, a longer branch is preferred as the final pipeline.

Algorithm 1: Operation Merge

Input : Candidate pipelines $\mathcal{L} = \{L_1, L_2, \dots, L_N\}$
Output : The merged pipeline L^*
 $L^* \leftarrow \emptyset$
 // Step 1: Aggregate all select operations
 $C_{\text{sel}} \leftarrow \bigcup_{L_i \in \mathcal{L}} \{\text{columns in select}(\cdot) \text{ operators of } L_i\}$
 append select(C_{sel}) to L^*
 // Step 2: Build an operation trie over the remaining operators
 $\mathcal{L}' \leftarrow \{L_i \setminus \{\text{select}(\cdot)\} \mid L_i \in \mathcal{L}\}$
 initialize an empty trie $\mathcal{T}$ with root node r
 initialize the node weight $w(v) \leftarrow 0$ for all nodes $v \in V(\mathcal{T})$
foreach pipeline $L'_i \in \mathcal{L}'$ **do**
 insert L'_i into $\mathcal{T}$ as a branch starting from r
 foreach node v along the insertion path **do**
 $w(v) \leftarrow w(v) + 1$
 // Return a path with maximum node weights
 $\mathcal{P} \leftarrow \{P \mid P \text{ is a root-to-leaf path in } \mathcal{T}\}$
 $P^* \leftarrow \arg \max_{P \in \mathcal{P}} [\sum_{v \in P} w(v)]$
 append all operators from P^* to L^*
return L^*

The complete merging procedure is in Algorithm 1, which strikes a balance between inclusivity and precision: it captures all potentially informative schema modifications (via union and retention) while enforcing consensus for other transformations. Empirically, this approach has been shown to improve the consistency of downstream QA performance by mitigating the effects of noisy or unstable operator predictions from individual generations.

3.4 Adaptive Rollback

Although we design lots of strategies and techniques to avoid generating incorrect data preparation pipelines, it still remains risks of over-processing and information loss. *I.e.*, essential information may be inadvertently removed from the table. To address this challenge, we design an adaptive rollback mechanism that enables self-correction when information loss is detected.

The rollback process is triggered when the QA model detects that the processed table T' loses the necessary information to answer the question and it outputs “No data available”. This output serves as a self-diagnostic signal indicating the model’s uncertainty under missing data conditions.

To balance effectiveness and efficiency, we propose binary rollback, which identifies the first state that supports correct reasoning through binary search. Given a table T and a produced data preparation pipeline $P = [op_1, op_2, \dots, op_n]$, we define a series of rollback states (the produced tables in each states) as follows:

$$T_{\text{rb}}^{(k)} = \begin{cases} op_{l_k}(\dots op_1(T) \dots) & l_k \geq 1 \\ T & l_k = 0 \end{cases}, \quad l_k = \left\lfloor \frac{n}{2^k} \right\rfloor \quad (11)$$

Initially, the full pipeline is applied and obtains an initial state. If the system emits “No data available”, we recover the state at the middle position $\lfloor l_k/2 \rfloor$ of the current pipeline, and check whether

the LLM can reason out a correct answer based on the produced table. If not, the search continues for the first half sub-pipeline.

4 EXPERIMENTS

4.1 Experiments Setup

Datasets. We evaluate Operation-R1 using 4 datasets, WikiTQ [22], TabFact [5], FeTaQA [20], and TableBench [34].

- **WikiTQ** is a dataset for table question answering, consisting of complex natural language questions over diverse Wikipedia tables, with 17,689 training and 4,344, test pairs covering wide variations in table structures and reasoning types.
- **TabFact** is a table-based fact verification dataset built from Wikipedia, where each statement-table pair is labeled as *True* or *False*, emphasizing logical reasoning and consistency checking over tabular evidence.
- **FeTaQA** is a free-form table question answering dataset. Built from Wikipedia tables and derived from selected ToTTo examples, it emphasizes multi-cell evidence aggregation, reasoning, and long-form answer generation.
- **TableBench** is a comprehensive TQA benchmark covering four reasoning categories, including Fact Checking (FC), Numerical Reasoning (NR), Trend Forecasting, and Chart Generation.

Baselines. We evaluate Operation-R1 against 10 baselines.

- **No-Prep** performs direct table QA without preprocessing, using the prompting template from Table-R1 (Yang et al.) [36].
- **DATER** [38] is an LLM-driven decomposer that splits a large evidence table into salient sub-tables and factorizes complex questions into executable sub-questions via a parsing-execution-filling pipeline.
- **Binder** [7] is a program-centric approach that binds language models to symbolic execution (e.g., SQL/code), compiling model outputs into interpretable, executable programs.
- **ReAcTable** [40] is a ReAct-style TQA system that incrementally constructs intermediate table representations and invokes external SQL/Python executors for progressive data transformation and answer derivation.
- **CoTable (Chain-of-Table)** [32] is a tabular reasoning system that treats an evolving table as the chain of thought, iteratively generating operations to update the table and expose verifiable intermediate results before producing a final answer.
- **AutoPrep** [8] is a question-aware data preparation multi-agent framework that decomposes complex table reasoning tasks into structured data preparation subtasks through collaborative agents, including a planner, a preprocessor, and an executor.
- **Table-R1 (Yang et al.)** [36] introduces two variants for end2end table reasoning. *Table-R1-SFT* is a supervised fine-tuning model trained on $\sim 48.5k$ carefully filtered reasoning traces generated by DeepSeek-R1 across three task types (TQA, TFV, FF-TQA), employing full-parameter fine-tuning. *Table-R1-Zero*, instead of being initialized from the SFT checkpoint, is trained directly from the same base model using $\sim 69k$ verifiable samples with the GRPO algorithm. To enable a controlled comparison under identical settings, we further introduce two additional variants, *Table-R1-GRPO* and *Table-R1-DAPO*, which apply GRPO and DAPO respectively to train an LLM to directly predict final answers.

Both share the same backbone (Qwen3-4B, thinking mode), LoRA fine-tuning, and training data (WikiTQ) as Operation-R1.

- **Table-R1 (Wu et al.)** [35] proposes a region-based reinforcement learning framework built on the Table Instruct-RE dataset ($\sim 20k$), where table regions are automatically generated by DeepSeek-R1 to capture the minimal evidence areas necessary for solving each problem. It adopts two-stage training that first guides LLMs with SFT and then optimizes via Table-Aware Group Relative Policy Optimization (TARPO) with a mixed reward balancing region precision and correctness.

Evaluation Metrics. We use *accuracy* to evaluate model performance on both datasets. For WikiTQ and TableBench, accuracy is computed as the percentage of questions where the predicted answer matches the ground truth. Following the evaluator from Binder [7], we apply answer normalization to handle unit variations and list-based answers. For TabFact, accuracy measures the percentage of statements correctly classified as True or False. For FeTaQA, we report BLEU score.

Implementation details. Refer to Appendix A in the extended version [16].

4.2 Overall Performance

We conduct experiments on WikiTQ and TabFact across multiple backbones. As shown in Table 1, Operation-R1 consistently outperforms all compared baselines under each backbone setting.

Operation-R1-4B improves over baselines by 8.83% on WikiTQ and 4.44% on TabFact on average; the smaller Operation-R1-1.7B variant surpasses baselines by 6.12% and 2.71% respectively. Traditional program-based methods (Dater, Binder) have no clear advantage over No-Prep due to execution brittleness. Operation-based methods (CoTable, AutoPrep) exhibit strong backbone dependency: on smaller LLMs, they show no consistent advantage over No-Prep and occasionally underperform it—AutoPrep in particular suffers catastrophic degradation (e.g., Qwen3-4B/TabFact: 53.36% vs. 74.95% No-Prep) as smaller models generate buggy code that traps the system in futile self-debugging loops. With DeepSeek-V3.2 ($>600B$) as backbone, baselines benefit substantially from the stronger model—Dater and AutoPrep improve to around 74.3% on WikiTQ—yet Operation-R1-4B still surpasses them by over 9 points (83.75%), confirming that our training strategy contribute gains orthogonal to backbone capability: the operator orchestration policy is model-agnostic, and a stronger backbone directly translates into stronger Operation-R1 performance. Under Qwen3-4B thinking mode, No-Prep already reaches 77.46%/92.89% with test-time scaling, leaving limited headroom; Operation-R1-4B nonetheless delivers consistent improvements (+4.17%/+0.58%). Thinking mode is not evaluated for multi-step baselines as adapting their prompting pipelines requires non-trivial redesign and incurs $\times 6.8$ average overhead (more statistics are provided in Appendix B.5 [16]).

Operation-R1’s stable performance stems from framing table preprocessing as a trainable operator orchestration task. By learning to compose operator sequences via ORPO, the system achieves robust results across backbones without the brittleness of code execution or the overhead of iterative agent collaboration.

Table 2 compares Operation-R1 against state-of-the-art end-to-end training approaches. Despite using only 4B parameters and 5K

Table 1: Comparison of Table Understanding Methods under Unified Model Settings. ([†] Thinking mode is not evaluated for all baselines due to non-trivial overhead. More details and statistics are referred to extend version [16].)

Model	Llama3.1-8B		Qwen2.5-7B		Qwen3-4B		Qwen3-4B(Thk) [†]		Qwen3-8B		GPT-4.1 mini		DeepSeek-V3.2	
	WTQ	TF	WTQ	TF	WTQ	TF	WTQ	TF	WTQ	TF	WTQ	TF	WTQ	TF
No-Prep	49.10	77.08	50.69	71.99	46.20	74.95	77.46	92.89	46.87	76.58	67.84	88.41	69.94	92.44
Dater	29.66	72.63	46.87	66.45	49.85	82.07	-	-	53.27	68.13	53.91	64.52	74.28	82.91
Binder	50.44	74.87	36.34	77.83	50.14	79.13	-	-	55.31	81.93	60.01	83.20	54.26	86.19
ReAcTable	42.94	-	47.33	-	53.38	-	-	-	52.30	-	60.15	-	70.95	-
CoTable	40.56	75.59	46.20	80.43	57.50	81.72	62.27	89.62	52.21	<u>91.25</u>	75.04	88.69	72.12	90.96
AutoPrep	17.05	49.52	55.64	74.99	59.68	53.36	-	-	63.53	78.28	72.81	88.24	74.26	85.95
Operation-R1-1.7B	<u>60.31</u>	<u>79.84</u>	<u>65.68</u>	<u>84.54</u>	<u>68.16</u>	<u>90.46</u>	<u>79.10</u>	<u>93.03</u>	<u>68.39</u>	90.42	<u>76.06</u>	<u>92.19</u>	<u>81.24</u>	<u>93.33</u>
Operation-R1-4B	64.27	84.24	69.66	88.44	70.72	91.35	81.63	93.47	71.06	92.14	76.77	93.87	83.75	94.17

Table 2: Comparison of End2End Training-based QA Methods. ([†]: results from the original studies due to unavailable model releases or GPU limits).

Method	Base Model	Train Data Size	Finetuning Method	WikiTQ	TabFact
TableGPT2 [†]	Qwen2.5-72B	2.36M	Full-Parameter	71.5	82.5
Table-R1-SFT (Yang et al.) [†]	Llama-3.1-8B-Instruct	34K	Full-Parameter	83.8	91.1
Table-R1-SFT (Yang et al.) [†]	Qwen2.5-7B-Instruct	34K	Full-Parameter	<u>81.9</u>	89.9
Table-R1-ZERO (Yang et al.) [†]	Llama-3.1-8B-Instruct	49K	Full-Parameter	81.2	87.6
Table-R1-ZERO (Yang et al.) [†]	Qwen2.5-7B-Instruct	49K	Full-Parameter	79.8	87.6
Table-R1 (Wu et al.) [†]	Qwen3-8B (Thinking mode)	20K	Full-Parameter	73.87	-
Table-R1-GRPO	Qwen3-4B (Thinking mode)	11K	LoRA	78.78	91.97
Table-R1-DAPO	Qwen3-4B (Thinking mode)	11K	LoRA	79.08	92.24
Operation-R1	Qwen3-4B (Thinking mode)	5K	LoRA	<u>81.63</u>	93.47

training samples, Operation-R1 achieves competitive WikiTQ accuracy (81.63%) and state-of-the-art TabFact performance (93.47%). *Operator Generation vs. Natural Language Reasoning.* Against large-scale supervised methods, Operation-R1 outperforms TableGPT2 (72B, 2.36M samples) by 10.12 points on WikiTQ and 10.97 points on TabFact, demonstrating the superiority of targeted operator-based preprocessing over massive generic pre-training. Table-R1-SFT with Llama-3.1-8B achieves the highest WikiTQ score (83.8%) but requires 34K curated reasoning traces from DeepSeek-R1. Operation-R1 maintains robust performance with 6.8× less training data while surpassing all SFT variants on TabFact (+2.37 points). This provide empirical evidence that structured operator composition can be easier to learn than free-form textual reasoning in our setting. *Preprocessing-then-QA vs. End-to-End RL.* Our controlled comparison with Table-R1-GRPO and Table-R1-DAPO—which train end-to-end QA on the same Qwen3-4B backbone with 11K samples—yields 78.78%/91.97% and 79.08%/92.24%. Operation-R1 outperforms both by over 2 points on WikiTQ using only 5K samples, *despite being trained solely on operator generation rather than the QA task itself.* This counterintuitive result stems from a fundamental difference in supervision quality: end-to-end QA conflates table understanding, information selection, and multi-step reasoning into a single objective with sparse, binary reward signals. Our two-stage paradigm provides deterministic, step-wise feedback at each transformation.




Figure 3: Efficiency-Accuracy-Cost Trade-off Comparison.

4.3 Efficiency Analysis

To ensure fair comparison, we evaluate efficiency on the full test sets under controlled settings. We measure end-to-end latency and actual API cost rather than token count, as API cost reflects LLM-side behavior such as reasoning effort and cache utilization. For latency, due to API stability concerns, we uniformly deploy Qwen3-8B across all baselines and Qwen3-4B (Thinking mode) for No-op and Operation-R1; for cost, we report actual API spending incurred by GPT-4.1-mini, while Operation-R1 is costed using the Qwen3 API pricing on the DashScope platform.

Fig. 3 presents a comprehensive efficiency-accuracy-cost analysis across WikiTQ and TabFact. Operation-R1 consistently achieves a favorable trade-off: in non-thinking mode, both the 1.7B and 4B variants outperform all multi-step agentic baselines in accuracy while incurring substantially lower latency and cost. Among the baselines, multi-step agentic baselines (AutoPrep, ReAcTable) impose the heaviest overhead due to their iterative LLM invocation loops, yet still fall short of Operation-R1’s accuracy. Program-based methods (Binder, Dater) scatter across the low-accuracy region without offering meaningful efficiency advantages.

Enabling thinking mode roughly doubles latency but delivers consistent accuracy gains, representing a worthwhile test-time scaling trade-off that remains more cost-efficient than multi-step agentic alternatives at comparable accuracy levels. It is also worth noting that since Qwen3-1.7B and Qwen3-4B share identical per-token pricing on DashScope, local deployment of the 1.7B model would yield even lower inference cost in practice.

Table 3: Ablation Studies on Two Backbone Models

Method	Qwen2.5-7B		Qwen3-4B(Thk)	
	WikiTQ	TabFact	WikiTQ	TabFact
No-Prep	50.69	71.99	77.46	92.89
Operation-R1	69.27	88.34	81.63	93.47
– ORPO & VGR strategy	63.49	85.67	73.87	89.53
– VGR strategy	68.67	86.35	80.57	93.28
– Operation Merge	68.16	85.67	76.84	88.49
– Adaptive Rollback	64.24	85.13	76.70	87.35

4.4 Ablation Studies

To validate each core component of Operation-R1, we conduct ablation studies on Operation-R1 with two backbone models. The results are presented in Table 3.

Completely removing ORPO training leads to an average accuracy drop of 6.8% accuracy drops across both backbones: $-5.78/-2.67$ on Qwen2.5-7B and $-7.76/-3.94$ on Qwen3-4B (WikiTQ/TabFact), confirming that ORPO reduces operator generation errors at the modeling level. Disabling Operation Merge results in substantial performance drops of 4.2% on the two backbones respectively ($-1.11/-2.67$ and $-4.79/-4.98$), confirming that consensus-based operator aggregation is critical for filtering spurious single-inference outputs. Similarly, removing Adaptive Rollback decreases accuracy by 5.8% on average, demonstrating its role in handling over-processing failures at runtime.

When only VGR is disabled, accuracy drops by 0.60/1.99 on Qwen2.5-7B and 1.06/0.19 on Qwen3-4B, indicating its importance in maintaining training stability. As shown in Fig. 4, removing VGR leads to a pronounced early decline in reward curves and slower subsequent recovery, whereas the full VGR strategy yields smoother, steadily increasing curves and higher final rewards.

4.5 In-depth Analysis

To gain deeper insights into how Operation-R1 achieves its performance improvements, we analyze the distribution and effectiveness of operator sequences across the WikiTQ and TabFact datasets.

Table 4: Performance comparison on additional types of questions on TableBench and FeTaQA.

Dataset	Question Type	Method	Performance
TableBench	Fact checking	No Prep	57.29
		CoTable	41.67
		AutoPrep	65.62
		Operation-R1	77.08
	Numerical reasoning	No Prep.	11.59
		CoTable	29.22
		AutoPrep	50.76
		Operation-R1	44.87
FeTaQA	Free-form table QA	No Prep.	21.10
		CoTable	19.89
		AutoPrep-CoT	16.05
		Operation-R1	23.78

Generalization to Unseen Datasets. Table 4 reports cross-dataset results on TableBench and FeTaQA without any additional fine-tuning. On TableBench, Operation-R1 achieves the best fact-checking accuracy (77.08%). On numerical reasoning, it surpasses No-Prep and CoTable but trails AutoPrep (44.87% vs. 50.76%), attributable to the absence of an aggregate operator. Introducing aggregate yields a +32.06% gain (Table 7), lifting overall NR accuracy to 61.96%. On FeTaQA, Operation-R1 achieves the highest BLEU score (23.78), while multi-step methods degrade below No-Prep. FeTaQA’s multi-intent questions require joint reasoning over multiple sub-queries; sequential execution biases the pipeline toward one misdirection, whereas single-step generation captures global intent and avoids premature commitment.

Performance on Cell-Focused vs. Non-Cell-Focused Tasks. Although Operation-R1 is trained exclusively on cell-focused (CF) instances, Table 5 shows that it yields consistent gains on non-cell-focused (NCF) tasks as well: on WikiTQ, the improvement over No-Prep on (+23.02%) even exceeds that on CF (+14.22%), indicating that operator-based preprocessing effectively reduces reasoning burden regardless of whether answers are directly extractable as cell values. This trend is further confirmed on TabFact, which is inherently NCF, where Operation-R1 surpasses all baselines by substantial margins. These results confirm that the preprocessing policies learned from CF supervision transfer robustly to NCF settings.

Analysis on LLM invocation. Table 6 breaks down the number of LLM invocations across three stages of TableQA: (1) *Prepare*: generating data preparation pipelines; (2) *Extra*: executing semantic operators such as `add_column`; and (3) *QA*: answering questions based on the curated table. Operation-R1 reduces total LLM calls by over 46% on both WikiTQ and TabFact compared to CoTable and AutoPrep, while maintaining superior accuracy. This reduction stems directly from our one-step pipeline generation: rather than invoking LLMs iteratively per operator as in loop-wise agentic methods, Operation-R1 generates the entire pipeline in a single forward pass, keeping *Prepare* calls equal to the number of test instances. The remaining overhead arises from a small number of parallel candidates used for Operation Merge and occasional calls for semantic operators, both of which are limited in frequency.




Figure 4: Training curves of Op-R1-1.7B with and without VGR strategy. Removing VGR causes a pronounced early drop in reward and slower subsequent improvement, whereas the VGR yields smoother, steadily increasing curves and higher rewards.

Table 5: Performance comparison on Cell-Focused (CF) vs. Non-Cell-Focused (NCF) tasks.

Dataset	Type	Prop.	Method	Accuracy (%)
WikiTQ	CF	46.1%	No-Prep	63.29
			CoTable	56.72
			AutoPrep	60.77
	NCF	53.9%	Operation-R1	77.51
			No-Prep	39.92
			CoTable	48.28
TabFact	NCF	100.0%	AutoPrep	54.80
			Operation-R1	62.94
			No-Prep	71.99
			CoTable	80.43

Table 6: LLM invocation comparison across Operation-R1 and baselines.

Dataset	Method	Prepare	Extra	QA	Total Calls
WikiTQ	CoTable	12537	197	4344	17078
	AutoPrep	9691	8109	5248	20965
	Operation-R1	4344	134	4831	9309
TabFact	CoTable	6258	44	2024	8326
	AutoPrep	3593	3780	2318	9691
	Operation-R1	2024	56	2376	4456

Additional results. We report additional results including the prevalence of cell-focused instances across popular TQA benchmarks, a comparison of rollback strategies, a sensitivity analysis on the number of parallel inferences for Operation Merge, and scalability with table size in [16] in the interest of space.

Table 7: Accuracy comparison on subsets of datasets where the additional operator unseen in training is required for TableQA.

Dataset	Usage (%)	Accuracy (%)		
		No Prep	ours	ours w/ additional <i>op</i>
<i>Subset requiring aggregate operator (op = op_{aggr})</i>				
TableBench-NR	52.64	5.26	37.80	69.86 (+32.06)
TableBench-FV	13.54	46.16	61.54	69.23 (+7.69)
<i>Subset requiring clean_column operator (op = op_{clean})</i>				
WikiTQ	13.63	45.43	52.36	56.93 (+4.57)
TabFact	17.98	82.14	84.61	86.81 (+2.20)

4.6 Generalization Ability

4.6.1 Generalization to extended operator space. To explicitly test the scalability of our framework beyond the five operators seen during training, we add two operators only at inference time, targeting the plug-and-play extensibility of Operation-R1:

Clean_column. The `clean_column` operator performs simple value normalization such as fixing date formats or standardizing numeric units. The target column c and cleaning description d are treated as parameters of the operator, denoted `clean_columnc,d`. The operator takes only a table T as input:

$$\text{clean_column}_{c,d}(T)$$

and returns a table where column c has been normalized according to d . For example, `clean_columnDate, 'standardize date format'(T)` normalizes date values in the Date column. This operator requires semantic understanding and is implemented by invoking an LLM.

Aggregate supports group-by with numerical aggregations (e.g., sum, mean, count), parameterized by a grouping column c , an aggregation target column c' , and an aggregation function f :

$$\text{aggregate}_{c,c',f}(T)$$

As shown in Table 7, both operators yield consistent gains on the subsets where they are triggered. `aggregate`, required in 52.64% of TableBench-NR instances, delivers a +32.06% gain on that subset (lifting overall NR accuracy to 61.96%) and a +7.69% gain on

Table 8: The evaluation results of Operation-R1 trained on different datasets (Accuracy %).

Training set	Query Num.	WikiTQ Acc.	TabFact Acc.
WikiTQ	5403	69.66	88.44
TabFact	5000	69.77	88.49

TableBench-FV. `clean_column`, triggered in 13–18% of WikiTQ and TabFact instances, improves accuracy by +4.57% and +2.20% respectively. These results confirm that Operation-R1 generalizes robustly to unseen operators—leveraging them appropriately without any retraining—and that cell-focused rewarding effectively incentivizes operator-based evidence localization and preservation.

4.6.2 Generalization of the training regime to non-cell-focused tasks. To show the generalization of our training regime, we adapt it to TabFact and evaluate the model on WikiTQ. To construct cell-focused supervision on TabFact, we prompt DeepSeek-V3.2 to annotate each claim-table pair with the *terminal cells in the logical reasoning chain*—those cells whose values directly determine the verdict—as surrogate targets for the correctness reward. This preserves the per-operator verifiability while accommodating the non-extractive nature of the task. Following prior works [3, 29], we apply difficulty-based filtering to construct a 5,000-example training set, excluding instances that are trivially solvable (by Qwen3-1.7B) or intractably complex (failing for Qwen3-30B-A3B).

As shown in Table 8, Operation-R1 trained on this auto-annotated TabFact subset achieves performance nearly identical to its WikiTQ-trained counterpart on both benchmarks (69.77% vs. 69.66% on WikiTQ; 88.49% vs. 88.44% on TabFact). These results confirm that the cell-focused proxy is not a fundamental constraint but a *generalizable training proxy*: for any table QA task, terminal reasoning cells can be identified via LLM annotation, enabling a plug-and-play extension to diverse table reasoning tasks without requiring explicit ground-truth cell annotations.

5 RELATED WORK

5.1 Table Reasoning Paradigms

Text-based methods. These methods serialize tables into natural language and apply in-context learning [4] or chain-of-thought reasoning [33, 44]. Beyond generic prompting, extensions include decomposition over large tables [9, 37, 38] and structured linearization of complex headers [41, 44]. These methods struggle with row-column structure encoding, incur high token overhead, and lack guarantees on numerical operations (e.g., sorting, aggregation). **Program-based methods.** These methods prompt LLMs to produce executable SQL or Python code to guarantee numerical correctness. Early work such as Binder [7] adapts Program-of-Thought by binding LLM-generated symbolic programs to external verifiers. ReAct-style systems [40] iteratively construct and execute partial programs with tool augmentation. Decomposition-based pipelines [21, 31, 38] split complex queries into structured sub-programs before execution, while multi-step planning methods [43] coordinate SQL generation across multiple reasoning stages and retrieval-augmented methods [12] integrate external knowledge for

program synthesis. Despite their strengths, these methods suffer from fragile program synthesis.

Operator-centric methods. These methods mediate between the two extremes by manipulating tables through composable operations, preserving structural reasoning while remaining simpler to generate than full programs. CoTable [32] treats iterative table transformation as the reasoning chain, while AutoPrep [8] organizes a multi-agent (planner/programmer/executor) pipeline to produce a question-aware preprocessed table passed to a separate QA model. Compared to text2sql, such paradigms fit more in semi-structured and dirty tabular data and provide better transparency and explainability. However, they require multiple LLM calls (up to 10) and super large models (≥ 70 B parameters), raising latency and cost concerns for real-time use.

5.2 Learning Paradigms for Table Reasoning

Supervised End-to-End Training. SFT-based methods optimize end-to-end table QA/fact verification with large curated or synthetic traces. Representative lines include TableGPT2 [27] and Table-R1-SFT [36], which achieve strong accuracy but require full-parameter training and tens of thousands of training signals (Table 2).

RLVR for Table Reasoning. RLVR methods [19, 25] optimize LLMs using automatically verifiable rewards. Specifically, (i) Yang et al. [36] and Jin et al. [13] applied variants of GRPO to short-form QA, fact verification, and free-form QA with task-specific rewards evaluating format adherence and answer correctness, enabling inference-time scaling without human annotations. (ii) Wu et al. [35] and Lei et al. [15] introduced region-based reward mechanisms that compute the Intersection over Union (IoU) between predicted and ground-truth table regions (cells/columns/rows), using GRPO to optimize both evidence localization accuracy and final answer correctness. However, current RLVR approaches either rely on sparse binary correctness signals [36] or require manually annotated ground-truth regions and evaluate reasoning quality indirectly by checking whether models mention correct areas in natural language [15, 35].

6 CONCLUSION

In this work, we present Operation-R1, enabling replacing multi-step data preparation pipeline assembly with one-step pipeline generation with a lightweight LLM (e.g., Qwen-1.7B/4B). To advance LLM’s capability in pipeline generation, we propose ORPO, a LLM post-training algorithm with a novel self-supervised pipeline rewarding mechanism and a dynamic variance-aware group resampling strategy. Moreover, we introduce inference-time operation merge and adaptive rollback modules in Operation-R1 to enhance its robustness. Extensive experiments on 4 datasets show that Operation-R1 improves TQA accuracy by 8.83% and 4.44% over multi-step preparation baselines with only a light RL training cost, as well as reducing inference latency by up to 51%, and 2.2× monetary cost.

ACKNOWLEDGMENTS

This work was supported in part by the NSFC under Grants No. 62472377. Lu Chen is the corresponding author of the work.

REFERENCES

- [1] Jayetri Bardhan, Anthony Colas, Kirk Roberts, and Daisy Zhe Wang. 2022. DrugEHRQA: A question answering dataset on structured and unstructured electronic health records for medicine related queries. *arXiv preprint arXiv:2205.01290* (2022).
- [2] Bradley Brown, Jordan Juravsky, Ryan Ehrlich, Ronald Clark, Quoc V Le, Christopher Ré, and Azalia Mirhoseini. 2024. Large language monkeys: Scaling inference compute with repeated sampling. *arXiv preprint arXiv:2407.21787* (2024).
- [3] Mingrui Chen, Haogeng Liu, Hao Liang, Huaibo Huang, Wentao Zhang, and Ran He. 2025. Unlocking the Potential of Difficulty Prior in RL-based Multimodal Reasoning. *arXiv:2505.13261* [cs.CV]
- [4] Wenhui Chen. 2023. Large language models are few (1)-shot table reasoners. In *Findings of the association for computational linguistics: EACL 2023*. 1120–1130.
- [5] Wenhui Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyu Zhou, and William Yang Wang. 2019. TabFact: A large-scale dataset for table-based fact verification. *arXiv preprint arXiv:1909.02164* (2019).
- [6] Yang Chen, Zhuolin Yang, Zihan Liu, Chankyu Lee, Peng Xu, Mohammad Shoeybi, Bryan Catanzaro, and Wei Ping. 2025. AceReason-Nemotron: Advancing math and code reasoning through reinforcement learning. *arXiv preprint arXiv:2505.16400* (2025).
- [7] Zhoujun Cheng, Tianbao Xie, Peng Shi, Chengzu Li, Rahul Nadkarni, Yushi Hu, Caiming Xiong, Dragomir Radev, Mari Ostendorf, Luke Zettlemoyer, et al. 2022. Binding language models in symbolic languages. *arXiv preprint arXiv:2210.02875* (2022).
- [8] Meihao Fan, Ju Fan, Nan Tang, Lei Cao, Guoliang Li, and Xiaoyong Du. 2025. AutoPrep: Natural language question-aware data preparation with a multi-agent framework. *Proceedings of the VLDB Endowment* 18, 10 (2025), 3504–3517.
- [9] Che Guan, Mengyu Huang, and Peng Zhang. 2024. MFORT-QA: Multi-hop few-shot open rich table question answering. In *Proceedings of the International Conference on Computing and Artificial Intelligence*. 434–442.
- [10] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. DeepSeek-R1: Incentivizing reasoning capability in LLMs via reinforcement learning. *arXiv preprint arXiv:2501.12948* (2025).
- [11] Jujie He, Jiakai Liu, Chris Yuhao Liu, Rui Yan, Chaojie Wang, Peng Cheng, Xiaoyu Zhang, Fuxiang Zhang, Jiacheng Xu, Wei Shen, et al. 2025. Skywork open reasoner 1 technical report. *arXiv preprint arXiv:2505.22312* (2025).
- [12] Ruochun Jin, Xiyue Wang, Dong Wang, Haoqi Zheng, Yunpeng Qi, Silin Yang, and Meng Zhang. 2025. TALON: A multi-agent framework for long-table exploration and question answering. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*. 27385–27401.
- [13] Rihui Jin, Zheyu Xin, Xing Xie, Zuoyi Li, Guilin Qi, Yongrui Chen, Xinbang Dai, Tongtong Wu, and Gholamreza Haffari. 2025. Table-R1: Self-supervised and reinforcement learning for program-based table reasoning in small language models. *arXiv preprint arXiv:2506.06137* (2025).
- [14] Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, et al. 2024. Tulu 3: Pushing frontiers in open language model post-training. *arXiv preprint arXiv:2411.15124* (2024).
- [15] Fangyu Lei, Jinxiang Meng, Yiming Huang, Tinghong Chen, Yun Zhang, Shizhu He, Jun Zhao, and Kang Liu. 2025. Reasoning-Table: Exploring reinforcement learning for table reasoning. *arXiv preprint arXiv:2506.01710* (2025).
- [16] Fengyu Li, Junhao Zhu, Kaishi Song, Lu Chen, Zhongming Yao, Tianyi Li, and Christian S. Jensen. 2026. Replacing Multi-Step Assembly of Data Preparation Pipelines with One-Step LLM Pipeline Generation for Table QA. *arXiv:2602.22721* [cs.CV]
- [17] Shih-Yang Liu, Xin Dong, Ximing Lu, Shizhe Diao, Peter Belcak, Mingjie Liu, Min-Hung Chen, Hongxu Yin, Yu-Chiang Frank Wang, Kwang-Ting Cheng, Yejin Choi, Jan Kautz, and Pavlo Molchanov. 2026. GDPO: Group reward-Decoupled Normalization Policy Optimization for Multi-reward RL Optimization. *arXiv preprint arXiv:2601.05242* (2026).
- [18] Youssef Mroueh. 2025. Reinforcement Learning with Verifiable Rewards: GRPO’s Effective Loss, D dynamics, and Success Amplification. *arXiv preprint arXiv:2503.06639* (2025).
- [19] Youssef Mroueh. 2025. Reinforcement Learning with Verifiable Rewards: GRPO’s Effective Loss, Dynamics, and Success Amplification. *arXiv preprint arXiv:2503.06639* (2025).
- [20] Linyong Nan, Chiachun Hsieh, Ziming Mao, Xi Victoria Lin, Neha Verma, Rui Zhang, Wojciech Kryscinski, Hailey Schoelkopf, Riley Kong, Xiangru Tang, Muthia Mutuma, Ben Rosand, Isabel Trindade, Renusree Bandaru, Jacob Cunningham, Caiming Xiong, and Dragomir R. Radev. 2022. FeTaQA: Free-form Table Question Answering. *Trans. Assoc. Comput. Linguistics* 10 (2022), 35–49.
- [21] Giang Nguyen, Ivan Brugere, Shubham Sharma, Sanjay Kariyappa, Anh Totti Nguyen, and Freddy Lecue. 2025. Interpretable LLM-based table question answering. *arXiv preprint arXiv:2412.12386* (2025).
- [22] Panupong Pasupat and Percy Liang. 2015. Compositional semantic parsing on semi-structured tables. *arXiv preprint arXiv:1508.00305* (2015).
- [23] Shraman Pramanick, Rama Chellappa, and Subhashini Venugopalan. 2024. SPIQA: A dataset for multimodal question answering on scientific papers. *Advances in Neural Information Processing Systems* 37 (2024), 118807–118833.
- [24] Nitarshan Rajkumar, Raymond Li, and Dzmitry Bahdanau. 2022. Evaluating the text-to-SQL capabilities of large language models. *arXiv preprint arXiv:2204.00498* (2022).
- [25] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. 2024. DeepSeekMath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300* (2024).
- [26] Guijin Son, Hyunwoo Ko, Hoyoung Lee, Yewon Kim, and Seunghyeok Hong. 2024. LLM-as-a-judge & reward model: What they can and cannot do. *arXiv preprint arXiv:2409.11239* (2024).
- [27] Aofeng Su, Aowen Wang, Chao Ye, Chen Zhou, Ga Zhang, Gang Chen, Guangcheng Zhu, Haobo Wang, Haokai Xu, Hao Chen, et al. 2024. TableGPT2: A large multimodal model with tabular data integration. *arXiv preprint arXiv:2411.02059* (2024).
- [28] Yang Sui, Yu-Neng Chuang, Guanchu Wang, Jiamu Zhang, Tianyi Zhang, Jiayi Yuan, Hongyi Liu, Andrew Wen, Shaochen Zhong, Na Zou, et al. 2025. Stop overthinking: A survey on efficient reasoning for large language models. *arXiv preprint arXiv:2503.16419* (2025).
- [29] Yifan Sun, Jingyan Shen, Yibin Wang, Tianyu Chen, Zhendong Wang, Mingyuan Zhou, and Huan Zhang. 2026. Improving Data Efficiency for LLM Reinforcement Fine-tuning Through Difficulty-targeted Online Data Selection and Rollout Replay. *arXiv:2506.05316* [cs.LG]
- [30] Wenqing Wang, Muhammad Asif Ali, Ali Shoker, Ruohan Yang, Junyang Chen, Ying Sha, and Huan Wang. 2025. A-IPO: Adaptive intent-driven preference optimization. *arXiv preprint arXiv:2510.10077* (2025).
- [31] Yuxiang Wang, Junhao Gan, and Jianzhong Qi. 2025. TabSD: Large free-form table question answering with SQL-based table decomposition. *arXiv preprint arXiv:2502.13422* (2025).
- [32] Zilong Wang, Hao Zhang, Chun-Liang Li, Julian Martin Eisenschlos, Vincent Perot, Zifeng Wang, Lesly Miculicich, Yasuhisa Fujii, Jingbo Shang, Chen-Yu Lee, et al. 2024. Chain-of-Table: Evolving tables in the reasoning chain for table understanding. *arXiv preprint arXiv:2401.04398* (2024).
- [33] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems* 35 (2022), 24824–24837.
- [34] Xianjie Wu, Jian Yang, Linzheng Chai, Ge Zhang, Jiaheng Liu, Xeron Du, Di Liang, Daixin Shu, Xianfu Cheng, Tianzhen Sun, et al. 2025. TableBench: A comprehensive and complex benchmark for table question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 39. 25497–25506.
- [35] Zhenhe Wu, Jian Yang, Jiaheng Liu, Xianjie Wu, Changzai Pan, Jie Zhang, Yu Zhao, Shuangyong Song, Yongxiang Li, and Zhoujun Li. 2025. Table-R1: Region-based reinforcement learning for table understanding. *arXiv preprint arXiv:2505.12415* (2025).
- [36] Zheyuan Yang, Lyuhao Chen, Arman Cohan, and Yilun Zhao. 2025. Table-R1: Inference-time scaling for table reasoning. *arXiv preprint arXiv:2505.23621* (2025).
- [37] Zhen Yang, Ziwei Du, Minghan Zhang, Wei Du, Jie Chen, Zhen Duan, and Shu Zhao. 2025. Triples as the Key: Structuring makes decomposition and verification easier in LLM-based TableQA. In *International Conference on Learning Representations*.
- [38] Yunhu Ye, Binyuan Hui, Min Yang, Binhua Li, Fei Huang, and Yongbin Li. 2023. Large language models are versatile decomposers: Decomposing evidence and questions for table-based reasoning. In *Proceedings of the international ACM SIGIR conference on research and development in information retrieval*. 174–184.
- [39] Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Weinan Dai, Tiantian Fan, Gaohong Liu, Lingjun Liu, et al. 2025. DAPO: An open-source LLM reinforcement learning system at scale. *arXiv preprint arXiv:2503.14476* (2025).
- [40] Yunjia Zhang, Jordan Henkel, Avriella Floratou, Joyce Cahoon, Shaleen Deep, and Jignesh M. Patel. 2024. ReActTable: Enhancing ReAct for Table Question Answering. *Proceedings of the VLDB Endowment* 17, 8 (2024), 1981–1994.
- [41] Bowen Zhao, Changkai Ji, Yuejie Zhang, Wen He, Yingwen Wang, Qing Wang, Rui Feng, and Xiaobo Zhang. 2023. Large language models are complex table parsers. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. 14786–14802.
- [42] Junnan Zhu, Jingyi Wang, Bohan Yu, Xiaoyu Wu, Junbo Li, Lei Wang, and Nan Xu. 2025. TableEval: A Real-World Benchmark for Complex, Multilingual, and Multi-Structured Table Question Answering. *arXiv preprint arXiv:2506.03949* (2025).
- [43] Jun-Peng Zhu, Peng Cai, Kai Xu, Li Li, Yishen Sun, Shuai Zhou, Haihuang Su, Liu Tang, and Qi Liu. 2024. AutoTQA: Towards Autonomous Tabular Question Answering through Multi-Agent Large Language Models. *Proceedings of the VLDB Endowment* 17, 12 (2024), 3920–3933.
- [44] Jin Ziqi and Wei Lu. 2023. Tab-CoT: Zero-shot tabular chain of thought. In *Findings of the Association for Computational Linguistics: ACL 2023*. 10259–10277.