



QUEST: Query Optimization in Unstructured Document Analysis

Zhaoze Sun
sunzhaoze@bit.edu.cn
Beijing Institute of Technology
China

Chengliang Chai*
ccl@bit.edu.cn
Beijing Institute of Technology
China

Qiyang Deng
Kaisen Jin
qiyangdeng@bit.edu.cn
jks@bit.edu.cn
Beijing Institute of Technology
China

Xinyu Guo
Han Han
xinyuguo@arizona.edu
hanhan1@arizona.edu
University of Arizona
United States

Ye Yuan
Guoren Wang
yuan-ye@bit.edu.cn
wanggr@bit.edu.cn
Beijing Institute of Technology
China

Lei Cao
caolei@arizona.edu
University of Arizona
United States

ABSTRACT

Most recently, researchers have started building large language models (LLMs) powered data systems that allow users to analyze unstructured text documents like working with a database because LLMs are very effective in extracting attributes from documents. In such systems, LLM-based extraction operations constitute the performance bottleneck of query execution due to the high monetary cost and slow LLM inference. Existing systems typically borrow the query optimization principles popular in relational databases to produce query execution plans, which unfortunately are ineffective in minimizing LLM cost. To fill this gap, we propose QUEST, which features a bunch of novel optimization strategies for unstructured document analysis. First, we introduce an index-based strategy to minimize the cost of each extraction operation. With this index, QUEST quickly retrieves the text segments relevant to the target attributes and only feeds them to LLMs. Furthermore, we design an evidence-augmented retrieval strategy to reduce the possibility of missing relevant segments. Moreover, we develop an instance-optimized query execution strategy: because the attribute extraction cost could vary significantly document by document, QUEST produces different plans for different documents. For each document, QUEST produces a plan to minimize the frequency of attribute extraction. The innovations include LLM cost-aware operator ordering strategies and an optimized join execution approach that transforms joins into filters. Extensive experiments on 3 real-world datasets demonstrate the superiority of QUEST, achieving 30%-6× cost savings while improving the F1 score by 10% -27% compared with state-of-the-art baselines.

PVLDB Reference Format:

Zhaoze Sun, Chengliang Chai, Qiyang Deng, Kaisen Jin, Xinyu Guo, Han Han, Ye Yuan, Guoren Wang, and Lei Cao. QUEST: Query Optimization in Unstructured Document Analysis. PVLDB, 18(11): 4560 - 4573, 2025.
doi:10.14778/3749646.3749713

*Chengliang Chai is the corresponding author.

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. 18, No. 11 ISSN 2150-8097.

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/qiyangdeng/QUEST>.

1 INTRODUCTION

Modern corporations often maintain a large amount of unstructured data including text documents such as web pages. In fact, according to IDC research [2], unstructured data accounts for 80%-90% of the data. Recently, researchers have started building LLMs-powered systems, such as UQE [9], ZenDB [29], Lotus [35], and Palimpsest [30], to analyze the valuable information hidden in text documents. These systems allow a user to select a set of documents, specify some attributes that can be extracted from them, and apply some database analytical operations on these attributes, e.g., filter, aggregation, or join. The core is to leverage LLMs [4, 36, 39] to effectively extract out the attribute values that users are interested in because recent research [3] shows that LLMs are remarkably good at data extraction. These efforts, if successful, can turn unstructured data into actionable insights. For instance, a lawyer may employ this system to swiftly locate legal cases about murder with a minimum of three charges and a 15-year sentence.

Similar to traditional databases, a successful unstructured data analysis system relies on a query optimizer that automatically produces a plan, minimizing query execution costs. However, LLMs play a central role in such a system, raising unique challenges and optimization opportunities. More specifically, compared to traditional database operations, LLM inferences are much more expensive in both execution time and monetary cost [30], no matter whether using commercial LLM services or deploying open source LLMs on high-performance thus expensive GPUs. Because unstructured data analysis relies on LLMs to extract attributes, the extraction operation thus constitutes its performance bottleneck. Therefore, the key optimization objective in this scenario is to minimize the LLM cost incurred during extraction, equivalent to (1) minimizing the LLM cost of each extraction operation, which depends on the number of input tokens to an LLM, and (2) minimizing the frequency of invoking the data extraction operations.

doi:10.14778/3749646.3749713

To achieve the above optimization objective, we propose QUEST with two key components: (1) an index-based attribute extraction strategy to minimize the *number of input tokens* per extraction; (2) an instance-optimized query execution strategy to minimize the *frequency* of attribute extraction. Crucially, these optimization strategies are model-agnostic. Although the cost and performance may vary across different LLMs, the core principles of QUEST remain effective in reducing the number of input tokens and extraction frequency. QUEST thus offers an affordable and scalable approach that allows users to select thousands of documents from a large document collection, specify any attributes out of tens or even hundreds of attributes that these documents potentially contain, and analyze these attributes.

(1) Index-based Attribute Extraction. QUEST designs an index for LLMs to efficiently and effectively extract attributes from documents. Rather than using LLMs to scan each document routinely, it only feeds LLMs the relevant text segments, thus significantly reducing the number of input tokens. This retrieval augmented generation (RAG) inspired solution features two innovative designs: a two-level index and evidence augmented retrieval.

Two-level Index. First, QUEST summarizes the subject of each document and indexes this information to filter documents irrelevant to the target attributes. Then QUEST builds a segment level index to retrieve relevant segments from the remaining documents. Compared to existing document analysis systems such as ZenDB [29], which filter data only at the segment level, our two-level strategy more accurately identifies the relevant segments because it avoids erroneously recognizing some segments as relevant when they are from documents in different subjects.

Evidence Augmented Retrieval. Second, when searching the index, existing RAG style solutions tend to miss relevant segments, because the query, which typically contains the attributes and their text description, is often not informative enough. To solve this problem, one can rely on users to manually enhance text descriptions like prompt engineering, but it is tedious and time-consuming. We propose a sampling-based approach to automatically collect the *evidence* on what the relevant segments should look like and use this evidence to *augment* the retrieval.

(2) Instance-optimized Query Execution. Rather than first extracting out the attribute values involved in a query and then employing the database query optimizer to produce an execution plan, QUEST adopts a *lazy extraction* strategy. That is, it interleaves attribute extraction and analytical operations (e.g., filter or join), only extracting an attribute when an analytical operation has to evaluate it. In this way, QUEST opens optimization opportunities to minimize the frequency of data extraction. For example, if a legal case in a document is not about murder, QUEST does not have to extract other attributes, e.g., the number of charges or the years of the sentence. Intuitively, similar to relational databases, appropriately ordering the filters or the joins in a query could effectively avoid extracting the values that have no chance to appear in the final query results.

Guided by the optimization objective of QUEST, we revisit the query optimization principles in databases, e.g., *filter ordering*, *predicate pushdown*, and *join ordering*. As detailed in Sec. 3, rather than simply ordering the filters based on their selectivities and applying the same order to process all documents, QUEST produces an

instance-optimized order for each individual document, leveraging the observation that extracting the same attribute could incur significantly different cost across documents.

To accurately estimate the extraction cost per document – the key of this instance optimized strategy, QUEST adopts an “optimize at execution time” architecture. That is, unlike traditional database optimizers, which produce optimized plans before query execution, QUEST estimates the cost during query execution time and produces a plan per document on the fly.

Moreover, we show that pushing down the predicates does not always yield the optimal plan. Instead, QUEST introduces a join transformation method that converts a join into a specialized filter and orders it with other filters in the query. We theoretically show that the plan produced in this way is guaranteed to be better than predicate pushdown. For a multi-way join, QUEST faces challenges in estimating join selectivities as data records are not available in advance. To address this, we introduce a dynamic join ordering strategy that determines which tables to join during query execution, leveraging our *optimize at execution time* architecture.

To summarize, we make the following contributions.

- (1) We propose an execution-time optimizer that produces query plans instance-optimized w.r.t. different documents.
- (2) We propose a join transformation technique that is guaranteed to outperform the classical predicate pushdown strategy.
- (3) We propose a two-level index to effectively filter irrelevant segments, reducing extraction costs without compromising accuracy.
- (4) Extensive experiments demonstrate the superiority of QUEST, achieving 0.3–6× cost savings while improving the F1-score by 10%–27% compared with the state-of-the-art baselines.

2 QUEST OVERVIEW

In this section, we first introduce how a user specifies a query and the types of queries that QUEST supports. We then describe QUEST’s overall architecture and highlight the novel designs.

2.1 User Queries

In QUEST, a user could formulate a query in SQL semantics similar to ZenDB [29]. That is, the user selects a subset of documents from her data sources, e.g., the legal documents produced in the last two years. She then specifies some attributes that can potentially be extracted from these documents and applies some analytical operations such as filtering and joining on these attributes. Note that the key techniques of QUEST are orthogonal to the query language and thus are compatible with other systems such as Palimpzest [30], which uses Python-style interfaces. Moreover, users who favor querying in natural language can use QUEST, which can employ current NL2SQL techniques [15, 16, 31] to convert NL queries into SQL-like queries.

Supported Queries. In this work, we target optimizing Selection-Projection-Join (SPJ) queries over unstructured documents. We leave optimizations on other important types of queries as future work, such as aggregation.

We first introduce some notations. A denotes the set of attributes in the user-specified query Q , and a_i denotes each attribute in A . ϑ denotes the expression in the WHERE clause of Q , consisting of a series of filters, and θ_j corresponds to each filter in ϑ . QUEST

supports a broad range of filters. (1) QUEST supports queries that are conjunctions or disjunctions of any number of filters. As an example, in Figure 1, the query Q seeks to find NBA players who are over the age of 35 and have made more than 12 All-Star appearances. The expression ϑ is a conjunction of these two filters (e.g., $\vartheta : \text{age} > 35 \text{ AND } \text{all-stars} > 12$). (2) Any single filter θ_j for $a_i \in A$ can be an equality filter (e.g., $\theta_j : a_i = \text{"Kevin Durant"}$), an open range filter (e.g., $\theta_j : a_i > 27$) or a close range filter (e.g., $\theta_j : 25 \leq a_i \leq 30$). QUEST also supports join operations. We use $\mathcal{G} = \{\mathcal{T}, \mathcal{E}\}$ to denote the join graph, where nodes (\mathcal{T}) denote the set of tables and edges (\mathcal{E}) represent the join relations.

2.2 Overall Framework

Inspired by the principle of relational databases, QUEST first indexes the documents and conducts query optimization once receiving a query. However, different from traditional databases, which minimize query execution time, QUEST targets reducing LLM costs while ensuring analysis quality. We thus propose several innovative designs to achieve these objectives by (1) minimizing the cost of the attribute extraction via index and (2) minimizing the frequency of calling the extraction via query optimization.

Given a document collection $\mathcal{D} = \{d_1, d_2, \dots, d_n\}$, QUEST first builds a two-level index over all documents. Then, once the user selects a subset of documents, specifies some attributes in these documents, and composes a query on these attributes, QUEST will first *sample* some documents, extract the attribute values from them and collect some statistics used by the query optimizer, which aims to minimize the LLM costs of extracting attributes from documents. However, fundamentally different from traditional databases, QUEST optimizer does not produce a query plan before query execution. Instead, it uses these statistics and the LLM costs calculated during query execution to produce optimized plans *on the fly*. Moreover, given a query, it produces query plans at a *document by document* basis instead of presuming a uniform plan for all documents. This is because the attribute extraction cost could vary significantly from document to document. To quickly produce these plans, QUEST features novel optimization strategies, which, although lightweight, guarantee to produce optimal plans. The query execution engine uses index-based data extraction and other analytical operations such as join to execute query plans efficiently.

QUEST’s key components include: a two-level index-based strategy to minimize the cost of extracting attributes from the documents – the most fundamental operation in QUEST; and (2) a query optimizer that revisits the classical query optimization principles in relational databases, such as filter ordering, predicate pushdown, and join ordering, to avoid unnecessary extraction operations.

2.3 Index-based Attribute Extraction

The key idea is to use an index to quickly and accurately identify the text segments that potentially contain the target attributes and only feed these relevant segments to LLMs for extraction. Moreover, we leverage the common patterns shown in the sample segments that have the target attributes to improve the extraction quality. The details are presented in Sec. 4.1.

At a high level, the index-based tuple extraction resembles Retrieval Augmented Generation (RAG). That is, QUEST divides all

documents into segments, encodes each segment into an embedding, and loads the embeddings into a high dimensional vector index such as PQ [18] or HNSW [33]. When a query comes, QUEST encodes the name of the relevant attributes and their corresponding text descriptions into an embedding. It then uses the vector index to retrieve the segments with embeddings similar to the query embedding. Then QUEST only extracts tuples from these segments.

QUEST optimizes the two key steps of this basic RAG style strategy, namely the *indexing* and the *retrieval*.

Two-level Index. In addition to indexing documents at the segment level, QUEST constructs a document-level index to filter the documents irrelevant to the to-be-extracted attributes. The segment-level index will only be used to retrieve the relevant segments from the remaining documents, as shown in Figure 1. QUEST produces this document-level index by first extracting the *key sentences* from each document, encoding these sentences into an embedding $e(d_i)$ to represent the document, and then indexing these embeddings. Next, QUEST segments each document, encodes each segment into an embedding $e(s_i)$, and builds the segment-level index.

Evidence-Augmented Segment Retrieval. Although QUEST lets users explain attributes, embeddings from attribute names and descriptions often lack information to find all relevant segments. For example, when QUEST extracts the age attribute, simply embedding age alongside its description “Player’s age.” to search relevant segments may miss the segments such as “Wardell Stephen Curry II (born March 14, 1988) is an American professional basketball player and point guard ...”. This mismatch occurs because the segment does not explicitly mention “age”, but instead contains a birthdate and unrelated details like “point guard”. Therefore, its embedding is not necessarily similar to that of the query.

Rather than relying on users to provide such hints, which in fact is equivalent to a tedious prompt engineering process, QUEST automatically collects such valuable information during sampling, serving as *evidence* to augment segment retrieval. Specifically, when using LLMs to extract an attribute from document samples, QUEST records the segments where a corresponding attribute value is extracted from the sampled documents. Then QUEST encodes these segments into representative embeddings and uses each as evidence to retrieve the segments relevant to this attribute from the relevant documents. Finally, it merges all these segments, eliminates duplicates, and inputs them into the LLM for extraction.

2.4 Query Optimization

We revisit the classical query optimization principles in relational databases, such as filter ordering, predicate pushdown, and join ordering, to avoid unnecessary attribute extraction operations.

Filter Ordering. In QUEST, the LLM cost of extracting different attributes from a document collection could vary significantly because the QUEST index tends to discover various numbers of relevant text segments with respect to different attributes. Similarly, the extraction cost varies significantly document by document, even when extracting the same attribute. This *variant cost* observation guides us to propose a new strategy to order filters when multiple filters exist in one query.

First, unlike traditional databases, which simply order the filters based on their selectivities, QUEST orders the filters based on a cost

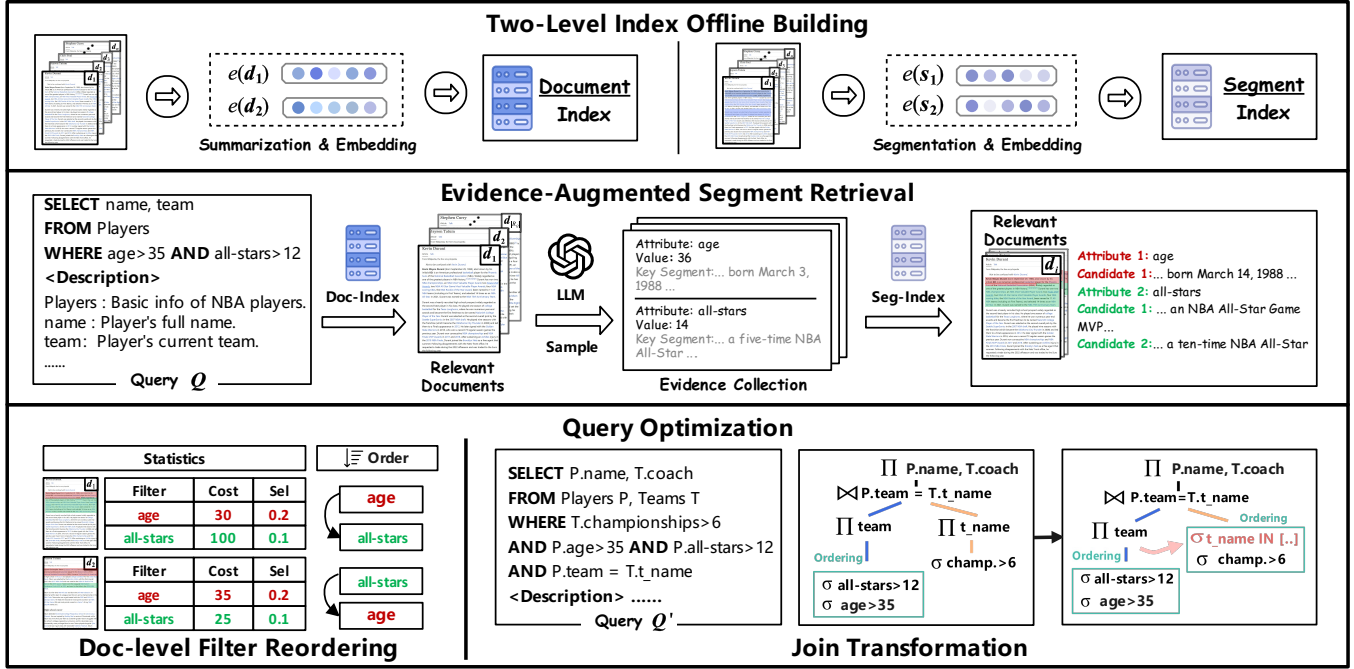


Figure 1: QUEST Framework

model that takes into consideration both the LLM cost – the number of tokens of the relevant segments and the selectivities estimated on sample documents. In Figure 1, taking query Q as an example, for document d_1 , the filter on age with a low selectivity but a small number of candidate tokens should have a higher priority than the filter on all_stars which although has a low selectivity, faces a large number of input tokens, thus incurring high LLM cost. Using this new cost model, we design a linear logarithmic time complexity algorithm that produces optimal plans.

Moreover, because extracting the same attribute from different documents could incur different costs, QUEST thoroughly abandons the “one single order for one query” methodology in traditional databases and instead produces different orders when evaluating tuples extracted from different documents, as shown in Figure 1.

EXAMPLE 1. For the query in Figure 1, the selectivity of all-stars, which is estimated on some sample documents, is smaller than that of age. However, the LLM costs vary between documents d_1 and d_2 . Therefore, when processing d_1 , QUEST will extract age before all-stars, because the cost of extracting all-stars is significantly higher. On the contrary, the cost of extracting all-stars from d_2 is slightly lower than that of extracting age. Thus, QUEST will prioritize all-stars over age due to its low selectivity.

Join Transformation. Although QUEST could execute a join by first extracting the attributes of the two tables involved in the join and then applying an existing join algorithm such as hash join, this is sub-optimal due to the prohibitive attribute extraction cost, while minimizing the extraction cost is QUEST’s key objective. Therefore, we propose a join transformation strategy that first extracts the join attribute of one table and then uses the extracted values as filters

to filter the other table (as shown in the bottom right of Figure 1). In other words, QUEST transforms a join into a filter. Treating this automatically generated filter equally to other filters, the QUEST optimizer uses the cost model discussed above to order these filters. In this way, QUEST might prioritize joins over filters to minimize the LLM cost, contradictory to the predicate pushdown principle in traditional databases.

EXAMPLE 2. In Figure 1, Players and Teams join on the team (t_name) attribute with the filters “P.age>35 AND P.all-stars>12” and “T.championships>6”. Suppose that after applying “P.age>35 AND P.all-stars>12”, only a few documents satisfy the filters on Players; and the values of the team attribute only include Warriors, Lakers and Celtics in these remaining documents. QUEST will add an IN filter “T.t_team IN [Warriors, Celtics, Lakers]” to Teams and order it with the existing filter “T.championships>6”. If the new filter has low selectivity and cost, it is likely to be executed prior to “T.championships>6”, thus saving LLM cost.

Dynamic Join Ordering. Different join orders may incur significantly different costs when multiple joins exist in a query. To reduce search space, traditional databases typically use dynamic programming to identify an effective order. However, accurately estimating join selectivities, the key to effectively ordering joins, is still an open problem. This is even worse in QUEST because the attribute values and the tables are not available beforehand. To address this issue, we introduce an algorithm that dynamically and progressively decides the join order during query execution. More specifically, QUEST first selects two tables to join based on our cost model, and it will determine the next join only after the first join finishes execution. This process iterates in a left-deep manner until

all joins have been executed. In this way, every time QUEST decides which table to join, the left table has already been extracted. This effectively alleviates the problem of estimating join selectivity.

In the remainder of this paper, we first introduce our query optimization techniques to minimize the frequency of attribute extraction (Sec. 3), which constitute our core technical novelties. Then, in Sec. 4, we introduce our index-based attribute extraction method, which minimizes the LLM cost of each extraction operation while improving the accuracy.

Remark. Although our primary focus is on optimizing the LLM cost for executing individual query, QUEST’s architecture naturally supports multiple concurrent queries. First, the two-level index built offline efficiently supports concurrent queries. Furthermore, the instance-optimized nature of query execution, processing documents either individually or in batches, enables straightforward intra-query parallelism: a query spanning numerous documents can be partitioned across multiple workers, each executing the optimized plan for its document subset. Crucially, QUEST is able to cache the results of LLM attribute extractions. By storing and reusing previously extracted tuples, redundant extraction operations can be eliminated among concurrent queries, significantly reducing overall cost and latency. These aspects, combined with standard database concurrency control mechanisms, ensure that QUEST is effective in serving simultaneous user requests. However, in-depth exploration of specific optimizations and benchmarking the performance of multiple query processing are beyond the scope of this paper and represent important directions for future work.

3 QUERY OPTIMIZATION

We first discuss the filter ordering optimization in Sec. 3.1. Then, we present the optimizations with respect to join in Sec. 3.2, including the optimization of one single join and join ordering.

3.1 Filter Ordering

As introduced in Section 2.4, QUEST differs from relational databases on filter ordering strategy, because QUEST’s key objective is to minimize the LLM cost in extracting tuples. In this section, we introduce our filter ordering method that takes both the number of input tokens to LLMs and the selectivities of the filters into consideration, yields optimal orders with respect to different combinations of filters, e.g., conjunction and disjunction.

3.1.1 The Filter Ordering Problem in QUEST

The filter ordering in QUEST aims to find an order of extracting attributes and evaluating the corresponding filters to minimize the LLM cost. Furthermore, as discussed in Section 2.4, because the extraction cost could vary significantly document by document, QUEST produces different orders with respect to different documents. Formally, we use $C_Q(o)$ to denote the total cost of executing query Q on a document with an order (denoted by o) of the filters in \mathcal{D} . The goal is to find the optimal order o^* :

$$o^* = \arg \min_{o \in \mathcal{O}} C_Q(o) \quad (1)$$

Next, we introduce our filter ordering method and analyze its optimality considering different combinations of filters, namely conjunction and disjunction.

3.1.2 Conjunctions

Consider a query Q that includes a WHERE clause containing a conjunction of filters. Let $A_s \subseteq A$, $A_w \subseteq A$ denote the set of attributes that appear in SELECT and WHERE clauses, respectively. For each filter $o[i]$ (i.e., the i -th filter in the order), the cost associated with extracting its relevant attribute is represented as $c^F[i]$, and its selectivity is given by $p[i]$. For each attribute $a_j \in A_s$, the generation cost is denoted by c_j^E . Therefore, for a given order o , the expected query cost can be represented as follows:

$$C_Q(o) = \sum_{i=1}^{|o|} c^F[i] \prod_{j=1}^{i-1} p[j] + \left(\sum_{j=1}^{|A_s|} c_j^E \right) \prod_{i=1}^{|o|} p[i] \quad (2)$$

where $\prod_{j=1}^{i-1} p[j]$ represents the likelihood that filter $c^F[i]$ has to be processed by an LLM, given that all its preceding filters in o return True. $\prod_{i=1}^{|o|} p[i]$ represents the probability of all filters returning True. Only in this case, QUEST has to extract the attributes in the SELECT clause, leading to a cost of $\sum_{j=1}^{|A_s|} c_j^E$.

Remark. Note that for conjunctions, A_s will have to be extracted only if all filters return True. Therefore, in this scenario, QUEST should always extract the attributes in A_w first, followed by A_s . As a result, when QUEST computes the optimal order for conjunctions, it only has to consider the cost of extracting attributes in A_w , i.e., only the first term in Equation 2.

Optimal Order. To find the optimal order with respect to each document, the brute-force method is to enumerate all possible orders, which has a time complexity of $O(|D_Q| \times |A_w|!)$ and therefore is too costly. Our key insight here is that using the cost model in Eq. 3, QUEST is able to find the optimal order P^* in linear logarithmic time.

First, given a filter $\theta_k \in \mathcal{D}$, QUEST uses the index to retrieve the segments relevant to the corresponding attribute and estimates its cost c_k which is proportional to the number of tokens in these segments. QUEST then uses its selectivity p_k estimated on the sampled table, and the cost c_k to compute a priority score, which determines the position of each filter θ_k in the optimal order.

LEMMA 1. *Sorting filters in descending order based on the following priority score minimizes the expected query cost.*

$$\text{priority}(\theta_k) = \frac{1 - p_k}{c_k}, \theta_k \in \mathcal{D} \quad (3)$$

Intuitively, Eq. 3 prioritizes a filter that is more likely to return False and has a low LLM cost. The proof of this lemma can be found in our technical report [1].

3.1.3 Disjunctions

For disjunctions, a better order should instead prioritize a filter that is more likely to return True. QUEST thus has a better chance to short-circuit other filters. Accordingly, we modify the cost model from Eq. 2 to the equation below:

$$C_Q(o) = \sum_{i=1}^{|o|} c^F[i] \prod_{j=1}^{i-1} (1 - p[j]) + \left(\sum_{j=1}^{|A_s|} c_j^E \right) \left[1 - \prod_{i=1}^{|o|} (1 - p[i]) \right] \quad (4)$$

Optimal Order. In this scenario, QUEST can still produce the optimal order by sorting the filters by their priority scores in descending order if the score is computed slightly different from Eq. 3.

$$\text{priority}(\theta_k) = \frac{p_k}{c_k}, \theta_k \in \mathcal{D}. \quad (5)$$

However, in the disjunction scenario, the attributes A_s in SELECT clause have to be handled differently, especially when an attribute exists in both the WHERE and SELECT clauses, e.g., $A_s \cap A_w \neq \emptyset$. In this scenario, to correctly evaluate the query, these attributes must be extracted for the following reasons: (1) if one filter in A_w returns true, QUEST has to extract them as the output of the query; (2) if no filter in A_w returns true, although QUEST does not produce any output, it still has to extract these attributes. This is because in such circumstances QUEST has to examine all filters in A_w , while these attributes are also part of A_s .

Therefore, if $A_s \cap A_w \neq \emptyset$, QUEST first extracts the attributes in $A_s \cap A_w$, followed by sorting and executing the filters in $A_s \setminus (A_s \cap A_w)$ based on their priority scores, which still guarantees the optimal result. The proof is analogous to the way we prove the optimality of conjunction and, hence, is omitted here.

3.2 Query Optimization For Join

Next, we introduce the optimization techniques with respect to join queries. We begin with optimizing one single join, and then extend our discussion to queries involving multiple joins.

3.2.1 Single Join: Joining Two Tables

Consider a query that joins two tables T_1 and T_2 on attribute a (i.e., $T_1.a = T_2.a'$). The query also has multiple filters ϑ . We abuse θ_1 and θ_2 a little to denote the filters on T_1 and T_2 respectively.

Typically, relational databases apply filters before joins to reduce the number of tuples in costly join operations. However, to minimize LLM costs, we find that converting joins into filters and optimizing them with other filters is often more effective in our scenario.

Taking the query in Figure 2 as an example, θ_1 corresponds to the filter on T_1 (the Teams table), where $T.championships > 6$, and θ_2 corresponds to the filter on T_2 (the Players table), where $P.age > 35$. A traditional query optimizer first pushes down θ_1 and θ_2 to T_1 and T_2 , respectively. It then performs a join on the returned documents (i.e., the tuples underlined in Figure 2-c). We establish a cost model to compute the expected cost under this optimization. [Plan \mathbb{D} : *Push θ_1 to T_1 , θ_2 to T_2 and join*] The expected cost under the optimal order can be calculated as:

$$Cost(\theta_1(T_1) \bowtie \theta_2(T_2)) = \sum_{i=1}^{|T_1|} C_1^i + p_1 \sum_{i=1}^{|T_1|} c_a^i + \sum_{i=1}^{|T_2|} C_2^i + p_2 \sum_{i=1}^{|T_2|} c_{a'}^i \quad (6)$$

where C_1^i is the expected cost of executing θ_1 on the i -th document of T_1 . p_1 is the likelihood of having to extract a after performing θ_1 , and c_a^i is the cost of extracting a from the i -th document in T_1 . The expected cost of $\theta_2(T_2)$ is calculated in the same way.

According to Figure 2, the cost of executing θ_1 on T_1 can be computed as $\sum_{i=1}^{|T_1|} C_1^i = |T_1| \times c_1^i = 30 \times 50 = 1500$, while the cost of extracting a on T_1 can be computed as $p_1 \sum_{i=1}^{|T_1|} c_a^i = p_1 \times |T_1| \times c_a^i = 0.1 \times 30 \times 30 = 90$. Similarly, we have $\sum_{i=1}^{|T_2|} C_2^i = 2550$ and $p_2 \sum_{i=1}^{|T_2|} c_{a'}^i = 459$ on T_2 . Finally, we have $Cost(\theta_1(T_1) \bowtie \theta_2(T_2)) = 4599$.

However, in our scenario, the optimal plan might not be achieved. We analyze the join operation on unstructured documents, consisting of two steps: extracting values of the join attribute from one table and finding matches in the other table. Since the optimization goal of QUEST is to minimize data extraction, a join may not be more costly than a filter operation and could have a higher priority. Similar to filters, the priority of a join operation relies on its own

SELECT P.player_name FROM Players P, Teams T WHERE P.age>35 AND T.championships>6 AND P.team name=T.team name	Filters		sel	cost
	T.championships>6		0.1	50
	T.team_name (join)		0.3	30
	P.age>35		0.3	50
	P.team_name (join)		0.1	30

(a) Query

team_name	championships
Lakers	17
Bulls	6
Celtics	17
Warriors	7
...	...

(c) Teams (30 tuples)

(b) Statistics

player_name	age	team_name
LeBron James	39	Lakers
Stephen Curry	36	Warriors
Rudy Gobert	32	Timberwolves
Chris Paul	39	Warriors
...

(d) Players (51 tuples)

Figure 2: Example for JOIN

data extraction cost and the potential extraction cost it may introduce to the other table, which in turn is determined by the number of matches that each tuple could potentially find, namely the join selectivity. Consequently, the priority of a join can be determined in the same manner as that of a filter. We are now prepared to order operations in a query with one join and multiple filters.

[Optimal Plan: *Sort Join and Filters Together.*] As shown in Eq. 7, the cost of the optimal order corresponds to:

$$Cost^*(\theta_1(T_1) \bowtie \theta_2(T_2)) = \sum_{i=1}^{|T_1|} \hat{C}_1^i + \sum_{i=1}^{|T_2|} \hat{C}_2^i \quad (7)$$

where \hat{C}_1^i represents the optimal expected cost obtained by sorting θ_1 with join on the i -th document in T_1 . The same applies to \hat{C}_2^i .

Unfortunately, producing this optimal solution requires an accurate estimation of the join selectivity, which is known to be a notoriously hard problem in databases. It is even worse in our scenario, where the tables are in fact not available beforehand.

To tackle this challenge, we propose an approach that transforms a join operation into a filter operation and progressively orders the operations during query execution. First, it chooses one table and executes the respective operations, i.e., pushing down the filters on it and then extracting the join attribute. Now, QUEST has acquired all the values of this join attribute that could potentially produce the final query output. Therefore, it is able to convert the join operation into an IN filter and apply it to the other table. Using the samples from the second table, QUEST is able to estimate its selectivity. As a result, QUEST can order other filters along with this IN filter to minimize the expected cost. Essentially, QUEST might end up running the join operation ahead of the filters, contradicting the traditional database optimizers.

Taking Figure 2 as an example, suppose that QUEST chooses to transform the join into an IN operation as a filter on T_2 , i.e., "P.team_name IN [Lakers, Celtics, Warriors]". Assume that the selectivity of this filter is 0.1 as shown in Figure 2-b. This means that approximately 10% of the tuples in the Players table can be joined with the tuples in the Teams table that satisfy θ_1 . Then QUEST updates θ_2 to $\hat{\theta}_2$ = "P.team_name IN [Lakers, Celtics, Warriors] AND P.age>35". Because this newly generated filter has a relatively low selectivity and cost, QUEST will prioritize it over existing filters based on the principle discussed in Section 3.1. This filter prunes most documents before running other filters (i.e., P.age>35), thus significantly reducing the LLM cost.

Because we have two options to convert the join operation, namely an IN filter either on table T_1 or table T_2 , we establish two cost models respectively, for QUEST to make decision.

[Plan ②: *Push θ_1 to T_1 and transform the join to filter on T_2*] As discussed above, this plan executes θ_1 on T_1 first, and then transforms the join operation to a filter on T_2 . Its cost is estimated as below:

$$\text{Cost}(\theta_1(T_1) \rightarrow \hat{\theta}_2(T_2)) = \sum_{i=1}^{|T_1|} C_1^i + p_1 \sum_{i=1}^{|T_1|} c_a^i + \sum_{i=1}^{|T_2|} \hat{C}_2^i \quad (8)$$

[Plan ③: *Push θ_2 to T_2 and transform the join to filter on T_1*] This plan executes θ_2 on T_2 first, then transforms the join to a filter on T_1 . Its cost can be estimated in a similar way to that of Plan ②.

$$\text{Cost}(\theta_2(T_2) \rightarrow \hat{\theta}_1(T_1)) = \sum_{i=1}^{|T_2|} C_2^i + p_2 \sum_{i=1}^{|T_2|} c_a^i + \sum_{i=1}^{|T_1|} \hat{C}_1^i \quad (9)$$

For Plan ②, the cost on $|T_1|$ remains the same compared to Plan ①. However, the cost on T_2 changes since this plan prioritizes the filter on a' . Now the cost is $\sum_{i=1}^{|T_2|} \hat{C}_2^i = |T_2| \times (c_a^i + p_{a'} \times c_b^i) = 51 \times (30 + 0.1 \times 50) = 1785$. Here, $p_{a'}$ is the selectivity of the IN filter. Then the overall cost becomes $\text{Cost}(\theta_1(T_1) \rightarrow \hat{\theta}_2(T_2)) = 1590 + 1785 = 3375$.

Selecting a Plan. Next, we discuss how to use the above cost models to produce a query plan. First, we present a lemma showing that Plans ① and ② are at least as good as Plan ① in all cases. The proof can be found in our technical report [1]. Then, we show how QUEST picks one between Plan ② and Plan ③.

LEMMA 2. *Given a query Q containing a join operation, the expected cost of Plan ① is always greater than or equal to that of Plan ② and Plan ③, i.e., $\text{Cost}(\theta_1(T_1) \bowtie \theta_2(T_2)) \geq \text{Cost}(\theta_1(T_1) \rightarrow \hat{\theta}_2(T_2))$, and $\text{Cost}(\theta_1(T_1) \bowtie \theta_2(T_2)) \geq \text{Cost}(\theta_2(T_2) \rightarrow \hat{\theta}_1(T_1))$.*

Next, QUEST has to make a decision between Plan ② and Plan ③. The simplest approach is to calculate the cost using Equation 8 and 9. While the first two terms are straightforward, the challenge lies in determining the third term, specifically the selectivity of $\theta(\text{IN})$. This is essentially about estimating the join selectivity between T_1 and T_2 . Precise estimation of join selectivity is a known issue in databases. Worst yet, in QUEST, the optimizer only has access to samples of T_1 and T_2 , not the full tables, before executing the query.

Nonetheless, we state that it is typically sufficient to decide between Plans ② and ③ by only considering the first two terms. This is because a small sum of these two terms suggests that the number of tuples, either before or after applying the filters, is small. Consequently, the selectivity of $\theta(\text{IN})$ is low, which likely in turn leads to a small third term. Therefore, if a plan has a smaller cost on the first two terms than the other, its overall cost also tends to be smaller. Thus, given T_1, T_2 , if $\sum_{i=1}^{|T_1|} C_1^i + p_1 \sum_{i=1}^{|T_1|} c_a^i < \sum_{i=1}^{|T_2|} C_2^i + p_2 \sum_{i=1}^{|T_2|} c_a^i$, we choose Plan ②, otherwise we choose Plan ③.

Query Execution in QUEST: Mixing Query Optimization With Execution. After selecting a plan (e.g., Plan ②), QUEST extracts the values of the join attribute $T_1.a$, executes the filters on it if there are any, and then transforms the join into an IN filter on T_2 . Since QUEST has obtained all the values of $T_1.a$, it can more accurately estimate the selectivity of IN. QUEST then triggers the optimizer again and uses the filter ordering optimization described in Sec. 3.1 to produce the optimal order to execute the remaining filters.

3.2.2 Adaptive Join Ordering

Next, we discuss the join ordering strategy, which orders multiple joins involved in a query, a classical yet challenging problem in relational databases. Given a set of tables $\mathcal{T} = \{T_1, T_2, \dots, T_{|\mathcal{T}|}\}$, users could specify a join graph $\mathcal{G} = \{\mathcal{T}, \mathcal{E}\}$, where each edge indicates a join between two tables. In addition, $\theta_i, i \in [1, |\mathcal{T}|]$ denotes the corresponding filters on table T_i .

To produce an optimal join order, the classical database optimizers, e.g., the Selinger optimizer [7], utilize the dynamic programming algorithm [38] that depends on some key statistics, such as the cardinalities of the intermediate join results. Estimating these cardinalities becomes rather challenging in our scenario, as the attribute values have not been extracted from unstructured documents yet.

We introduce an adaptive join ordering approach to discover the optimal join order. First, it chooses one single join by iterating every edge $e \in \mathcal{E}$, i.e., every join in the query, estimating its cost, and choosing the join with the minimal cost. The cost of each join can be directly estimated using the cost model described in Section 3.2.1. Then it determines the join plan (Plan ② or ③) of this selected join operation and immediately executes this join. We use T' to denote the result of the first join. Next, it finds another table to join with T' , forming a left deep query plan. More specifically, we use $J(T')$ to denote the set of tables that can join with T' . QUEST selects the one in $J(T')$ that incurs the minimal cost. To this end, QUEST has to estimate the cost of $T' \bowtie T_j, T_j \in J(T')$. As discussed in Sec. 3.2.1, because T' has already been available, QUEST transforms this join to a IN filter operation and estimates its cost, i.e., $\sum_{i=1}^{|T_j|} \hat{C}_j^i$. Afterward, it joins T' with the selected table according to the join plan. QUEST repeats this process until all joins are conducted.

4 INDEX-BASED ATTRIBUTE EXTRACTION

In addition to avoiding unnecessary data extraction operations, we propose an index-based strategy to further reduce the cost of each extraction and improve the accuracy.

4.1 Two-level Index Construction

QUEST starts with constructing a document-level index. Given a document set \mathcal{D} as input, QUEST first uses the NLTK package to generate a document summary efficiently, which is then transformed into an embedding using a pretrained model. Here, we choose E5Model [44] due to its state-of-the-art performance on massive text embedding benchmarks covering diverse retrieval tasks.

QUEST then constructs the segment-level index. For each document $d \in \mathcal{D}$, QUEST dynamically splits the document into relatively small and semantically coherent segments. The goal is to ensure that each attribute can be extracted from a single segment. To achieve this, we employ the SemanticChunker function in LangChain to segment text by examining both its syntactic structure and semantic coherence. Initially, the document is divided into sentences. SemanticChunker then evaluates the embedding similarity of consecutive sentences. If sentences are semantically coherent, they merge; if not, they remain separate. This iterative procedure continues for all sentences, enhancing attribute extraction by maintaining semantic coherence.

Eventually, these segments are represented as a set \mathcal{S} , where each $s \in \mathcal{S}$ is a coherent, self-contained portion of the document's

content. QUEST then embeds each segment $s \in \mathcal{S}$ using E5Model. Finally, QUEST loads document and segment embeddings into two high-dimensional vector indexes for efficient data retrieval.

4.2 Searching the Index

Given a query, QUEST first uses the document-level index to search the relevant documents and then uses the segment-level index to identify the relevant segments from the returned documents. **Document Retrieval.** Once receiving a query, QUEST retrieves from \mathcal{D} the documents that potentially contain an attribute in the query. To achieve this, QUEST first converts attribute names and their descriptions into embeddings. Then it averages these embeddings to generate a final embedding $e(Q)$. Afterward, based on the document-level index \mathcal{I}_D , QUEST searches the documents with an embedding close to $e(Q)$, i.e., $D_Q = \{d_i | d_i \in \mathcal{D}, \text{dist}(e(d_i), e(Q)) < \tau\}$, where $\text{dist}()$ is the distance function. While cosine similarity is commonly used for embedding comparison, it is monotonically related to Euclidean distance when vectors are L2-normalized, i.e., $\|v_1 - v_2\|^2 = 2 - 2 \cdot \cos(v_1, v_2)$. Thus, minimizing normalized Euclidean distance is equivalent to maximizing cosine similarity for ranking. We adopt it here because it is natively supported in vector indexing libraries such as PQ [18] and is computationally efficient. τ is a threshold initially set as a high value to guarantee a high recall. However, this may return some irrelevant documents. To solve this problem, QUEST automatically adjusts this threshold in the next segment retrieval phase based on the sampled documents to obtain a more precise document set D_Q^* .

Evidence Augmented Segment Retrieval. As discussed in Section 2.3, to achieve accurate segment retrieval, we propose to sample a small subset (approximately 5% of D_Q) of documents that will be carefully analyzed by an LLM. To be specific, for the attributes in the attribute set A_Q of query Q , QUEST asks the LLM to return their values and the segments from which these values are extracted.

Next, we transform these segments into embeddings. The key observation here is that the segments containing the same attribute tend to show some common patterns. For example, in Stephen Curry’s profile, the segment about the “age” attribute includes “Wardell Stephen Curry II (born March 14, 1988) is an American professional basketball player ...”, while the corresponding segment in Kevin Durant’s profile mentions that “Kevin Wayne Durant (born September 29, 1988), also known by KD, is an American professional basketball player ...”. The patterns of these segments are remarkably similar. Using these patterns as additional evidence could enhance retrieval, addressing the issue where simply using the query embedding may miss attribute-relevant segments, as discussed in Sec. 2.3.

When no relevant segments are found for an attribute a_i in the sampled documents, QUEST leverages the LLM to synthesize evidence. It prompts the LLM with the attribute name, its description, and optional contextual information to generate a small number (e.g., 20) of representative text segments, which are then embedded.

However, using all these embeddings as evidence tends to unnecessarily introduce redundancy and, in turn, produce too many candidate segments. We thus propose to utilize the k -means algorithm to group these embeddings (with a relatively small k , such as 3) and only use the cluster centers as evidence. We use \bar{e}_i^j , $j \in [1, k]$

to denote one of the cluster centers, i.e., one piece of evidence for a_i (k pieces in total for each attribute).

Segment Retrieval. Next, we discuss how to use the collected evidence to retrieve the segments relevant to an attribute in the query, e.g., the attributes age and all-stars in the query Q shown in the middle block of Figure 1. For each document in D_Q^* , QUEST utilizes the segment-level index \mathcal{I}_S to find the relevant segments for a given attribute a_i based on each piece of evidence \bar{e}_i^j . Next, QUEST combines the retrieved segments in each document with respect to the evidence of a_i , removes the duplicate segments, and feeds them into the LLM for attribute extraction.

Setting the Threshold. QUEST uses two distance thresholds to determine whether a document or a segment is relevant to an attribute. Setting these thresholds appropriately is critical to ensure the quality of retrieval. Setting these distance thresholds too high tends to return many irrelevant segments and hence increase extraction cost, while setting them too small might miss relevant segments, in turn impacting the extraction quality. Clearly, relying on users to manually set these thresholds correctly is challenging. QUEST thus introduces an automatic thresholding method to solve this problem.

For the threshold τ used to find documents relevant to a query Q , QUEST initially sets it as a high value to avoid missing relevant documents. QUEST then adjusts it to an appropriate value by analyzing the LLM extraction results on the sampled documents. Specifically, QUEST first uses the high τ to obtain a set of documents denoted as D_Q . QUEST then samples a subset D_Q^S from D_Q and uses LLM to extract the attributes of a table from D_Q^S . Based on the result of the extraction, QUEST divides D_Q^S into two subsets: D_Q^n , which consists of documents lacking attribute information therefore deemed irrelevant to the query, and $D_Q^m = D_Q^S \setminus D_Q^n$, which contains relevant documents. The maximum Euclidean distance between the embeddings of documents in D_Q^m and the embedding $e(Q)$ serves as the threshold, i.e., $\tau = \max\{\text{dist}(e(d_i), e(Q)) | d_i \in D_Q^m\}$. Intuitively, this new τ threshold will exclude irrelevant documents.

Similarly, QUEST leverages the sampled documents to set the threshold γ_i to retrieve the segments containing an attribute a_i , $i \in [1, M]$. More specifically, γ_i is set as the maximal distance between any pair of segments containing the value of a_i , i.e., $\gamma_i = \max\{\text{dist}(E_i[x], E_i[y]) | \forall x, y \in [1, |E_i|], x \neq y\}$, where E_i represents the set of segments related to attribute a_i in D_Q^m . $E_i[x]$ and $E_i[y]$ denote the x -th and y -th embeddings of two segments in E_i . To be cautious, in implementation we increase γ_i by 0.1, that is, $\gamma_i = \gamma_i + 0.1$; and equally we adjust τ .

5 EXPERIMENTAL EVALUATION

5.1 Experimental Settings

Datasets. We use 3 datasets with 500 documents in total, covering diverse domains, and 50 queries in various types. We employ human evaluators verifying the attributes extracted by LLMs from documents, and thus establish the ground truth.

LCR [14] includes 3,000 case reports. We sample 100 documents from them. Each document averages 6,247 tokens and contains detailed information such as the court, judge, and legal reasoning.

WikiText. We crawl 200 Wikipedia pages across 10 domains, such as directors, cities, NBA players, companies, etc., some of which can be joined; the average number of tokens per document is 1,264. **SWDE** [19] is a dataset used in our baseline [3]. We sample 200 web pages, with each averaging 416 tokens. Despite the relatively short length of the documents, SWDE contains 16 attributes.

	Avg. #-Tokens	#-Doc	#-Attributes	#-Queries
LCR	6247	100	10	10
WikiText	1264	200	20	25
SWDE	416	200	16	15

Table 1: Datasets

The datasets vary in length, structure, and domain: WikiText has a hierarchical structure across various domains, SWDE includes short documents, and LCR features long documents from a single domain. They facilitate thorough evaluations in various scenarios. **Ground Truth Generation.** We organize each dataset into domains of similar documents, sample 5 documents per domain, and use LLMs to identify key attributes, as shown in Table 1. We utilize LLMs for attribute extraction from all documents, verified by 10 graduate students.

Query Construction. We create queries for single tables and join tables. The queries cover both range and equality filters. We first construct the filters in WHERE clause: (1) For each query, we randomly sample a certain number of attributes from the attribute set in the query to construct the filters; (2) For numerical attributes, we randomly create different types of filters, including =, ≤, and ≥, while for categorical attributes, we only generate equality filters; (3) We then use these single filters to construct conjunctions and disjunctions. Each of these three categories has roughly the same number of queries. Next, we randomly sample a certain number of attributes to form the SELECT clause. Finally, we ask graduate students to validate all queries and eliminate the unreasonable ones. **Baselines.** We compare QUEST with various baselines.

- (1) ZenDB [29] adopts a hierarchical semantic tree to extract tuples from documents and uses SQL-like queries for analysis.
- (2) PZ (Palimpsest) [30] allows users to convert and analyze unstructured data with a declarative language hosted in Python. The existing PZ prototype offers basic optimizations on the usage of LLMs, while in-depth optimizations are still ongoing.
- (3) Lotus [35] supports a bunch of LLM-powered operations to analyze documents. It features some basic optimizations to improve the accuracy and query latency.
- (4) RAG [26] embeds attributes and their descriptions for similarity search. In contrast to QUEST, it does not incorporate a document-level index and does not utilize evidence to enhance retrieval.
- (5) ClosedIE [23] uses a model that has been fine-tuned using a vast quantity of labeled (attribute, value) pairs to extract relevant information from a given context in response to a query.
- (6) Eva (Evaporate)[3] is an LLM-based data extraction method. Rather than routinely employing LLMs to extract values from each document, it instead uses LLMs to automatically synthesize code.
- (7) QUEST is our full-fledged solution.

Evaluation Metrics. We measure accuracy, cost, and latency with respect to all queries. For accuracy, we evaluate the average precision, recall, and F1-score across all queries. Given a query Q , the set of tuples returned by a method is denoted as $T(Q)$, and the ground truth is denoted as $GT(Q)$. A tuple $t \in T(Q)$ is considered correctly extracted only if all its cell values match the corresponding ground truth values. Therefore, we have $P = \frac{|T(Q) \cap GT(Q)|}{|T(Q)|}$, $R = \frac{|T(Q) \cap GT(Q)|}{|GT(Q)|}$ and $F1 = \frac{2 \cdot P \cdot R}{P + R}$. For LLM cost, we measure the average number of input and output tokens per document for each query, including the cost of sampling, to reflect the full processing of QUEST. For latency, we measure the mean execution time of queries per document.

5.2 Comparison with Baselines

Accuracy w.r.t. All Queries. By the results (including P , R and $F1$) shown in Table 2, these method are ranked as follows: QUEST \approx Lotus > PZ > ZenDB > RAG > Eva > ClosedIE.

QUEST and Lotus achieve the highest accuracy among all baselines. Lotus performs well because it uses an LLM to scan every piece of text of all documents, leading to extremely high LLM costs. QUEST is competitive with Lotus but with a much lower cost, because our two-level index and the evidence-augmented retrieval method accurately identify relevant segments for queries, feeding only these to LLMs. QUEST thus ensures quality with low LLM cost. On LCR dataset, QUEST achieves an F1-score of 0.7, much higher than that of Lotus (0.45). The reason is that the documents in LCR contain a large number of tokens, including much irrelevant information that misleads LLMs and causes hallucinations.

QUEST outperforms ZenDB because: (1) QUEST leverages the evidence to identify relevant segments, which is more accurate than ZenDB that simply uses the attribute description to find the most relevant sentence; (2) although ZenDB uses a semantic tree to locate each attribute, in practice, many documents, e.g., the documents in the LCR dataset, are not well-structured, making it hard for ZenDB to construct an effective tree. Therefore, the F1-score of ZenDB on LCR is only 0.52, much lower than that of QUEST (0.7). Similarly, QUEST outperforms PZ because its existing prototype lacks an effective text segment retrieval component.




Figure 3: F1-Score of Baselines (Different Query Groups)

As expected, the RAG-based method is less effective than QUEST because the embedding of attribute and its descriptions is often not informative enough, thus tending to miss segments that are highly relevant to the query but do not possess similar phrases. Our evidence augmented retrieval strategy successfully solves this issue. Evaporate uses LLMs to generate code for data extraction, aiming to reduce LLMs costs. However, it does not perform well on

		ClosedIE	Eva	RAG	PZ	ZenDB	Lotus	QUEST
Wiki	Precision	0.33	0.39	0.79	0.75	0.81	0.93	0.93
	Recall	0.19	0.29	0.24	0.81	0.78	0.84	0.82
	F1-score	0.24	0.33	0.37	0.78	0.79	0.90	0.87
SWDE	Precision	0.72	0.85	0.88	0.90	0.82	0.95	0.94
	Recall	0.51	0.59	0.78	0.83	0.86	0.98	0.97
	F1-score	0.59	0.70	0.83	0.86	0.84	0.96	0.95
LCR	Precision	0.11	0.16	0.42	0.43	0.45	0.45	0.63
	Recall	0.07	0.13	0.29	0.61	0.73	0.45	0.84
	F1-score	0.09	0.14	0.34	0.50	0.55	0.45	0.72

Table 2: Accuracy Comparison

Method	#-Token Cost			Latency (s)		
	Wiki	SWDE	LCR	Wiki	SWDE	LCR
Eva	-	-	-	0.06	0.05	0.08
ClosedIE	-	-	-	0.76	0.65	1.83
ZenDB	260	280	2530	2.08	1.77	2.73
PZ	400	320	2610	2.17	2.16	2.85
RAG	440	340	3500	2.57	2.65	3.01
Lotus	2520	1150	12480	2.66	2.78	3.36
QUEST	170	190	2030	1.12	1.21	2.68

Table 3: Cost and Latency Comparison




Figure 4: Cost of Baselines (Different Query Groups)

accuracy because code essentially corresponds to a limited number of rules, which tend to be less effective when handling complex documents. ClosedIE performs poorly as pre-trained NLP models lack generalizability across domains.

Overall Cost. Table 3 shows the LLM cost of the baselines. Costs for ClosedIE and Eva are excluded: ClosedIE does not employ LLM, and Eva uses minimal tokens for code generation, resulting in low accuracy for both. The other methods are ranked as follows: QUEST < ZenDB < PZ < RAG < Lotus. QUEST is the most cost-effective because (1) it has a two-level index that can precisely locate a small number of segments where an attribute can be extracted, and (2) the filter reorder optimization reduces the number of LLM calls by terminating the filter evaluations early when feasible.

Lotus is the most expensive method because it feeds the entire document to LLMs for each filter. For example, on dataset LCR, Lotus costs 6× more tokens than QUEST. RAG is cheaper than Lotus because it feeds only a subset of segments to LLMs rather than the entire document. PZ and ZenDB save more cost than RAG because they reorder filters based on selectivities. However, both methods are more costly than QUEST. For example, on dataset WikiText, PZ and RAG consume more than 2× tokens than QUEST because they process every document in the same order, while QUEST generates the optimal order per document.

Overall Latency. Table 3 shows the query latency. ClosedIE and Eva are the fastest, although their accuracy in general is low. This is because ClosedIE does not use LLMs, while Eva only spends a few tokens on code generation. For other LLM-based methods, QUEST is about 2× faster than PZ, RAG, Lotus and ZenDB. This is because, in these methods, the LLM inferences dominate the query execution time, while the fine-grained filter ordering strategy of QUEST reduces both the number of LLM calls and the number of tokens consumed per call. Lotus is the slowest because it has to send each document to the LLM when evaluating a filter.

Varying the Number of Filters. We evaluate QUEST’s performance with varying filter numbers, categorizing queries into: C1 with one

filter, C2 with 2-3 filters, and C3 with 4 or more. We can observe in Figure 3 that as the number of filters grows, the accuracy of all baselines decreases because more filters tend to introduce more errors during attribute extraction. The cost of almost all baselines increases because more filters invoke more LLM calls. However, the increase of QUEST is the slowest, thanks to our filter ordering that reduces unnecessary attribute extraction.

5.3 Comparison of Filter Ordering Strategies

We evaluate the following baselines: (1) Random: Filters are executed randomly; (2) Selectivity: Filters are ordered by selectivity; (3) Average_cost: Filters are ordered by selectivity and the estimated average cost of extracting each attribute from the sampled documents; (4) Exhaust: All potential orders are exhaustively enumerated to find the optimal one for each document.

In Figure 5, for queries in C1, the cost of all baselines is almost identical because there is only one filter and hence one order per query. For queries with more filters, these methods are ranked as follows by the LLM cost: QUEST ≈ Exhaust < Average_cost < Selectivity < Random. This shows that QUEST produces optimal orders in most cases. As the number of filters increases, QUEST’s performance improves due to increased optimization opportunities.

Furthermore, we evaluate the scalability of QUEST and Exhaust when handling queries with a relatively large number of filters. Figure 5 shows that as the number of filters increases, the run time of Exhaust explodes due to the exponential time complexity. Conversely, the run time of QUEST increases slowly, indicating that it is more efficient and effective for complex queries.

5.4 Evaluation of Join

Tables to be Joined. We define 4 tables (i.e., Player, Team, City, Owner) in WikiText dataset. Player and Team join on the team_name attribute. Team and City join on the location attribute. Team and Owner join on the owner_name attribute. We construct filters for each table using the method in Section 5.1.

Two-table Join. To demonstrate effectiveness, we compare QUEST with the typical optimization in relational databases, so-called Pushdown, which always pushes down the filters before performing the joins. We also compare QUEST with the Optimal plan, which is obtained by assuming the selectivity of each filter is known.

We create queries incorporating one of the three aforementioned joins, along with some random filters. In total, we construct 60 queries. We classify these queries into three categories based on the number of filters, as previously stated. We execute the set of queries (G1-G3) and present the mean token consumption. In Figure 6-a,

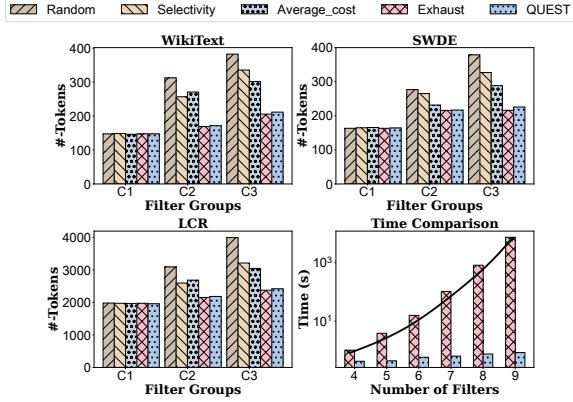


Figure 5: Comparison of Filter Reordering Strategies




Figure 6: Evaluation of Join

QUEST significantly reduces cost compared to Pushdown because it transforms a join to a filter operation and effectively orders the filters to minimize the frequency of LLM invocations. In this way, QUEST has the opportunity to run a join first if it incurs a small data extraction cost. In particular, QUEST only costs slightly more than Optimal due to its effective optimization strategies.

Next, we record the selectivity of each IN filter according to the plan chosen by QUEST for each query. We then classify the queries into three new categories ($E_1 - E_3$) based on the selectivities. The first group, E_1 , corresponds to the selectivities ranging from 0 – 0.3, while E_2 and E_3 correspond to 0.3 – 0.6 and 0.6 – 1, respectively. Moreover, we present the average cost associated with all queries. Figure 6-a illustrates that when the selectivity of the IN filter is low, it tends to be executed first, leading to lower cost than the traditional predicate pushdown strategy. Conversely, as the selectivity rises, the IN filter might fall back to the traditional predicate pushdown strategy, thus always finding a plan with the lowest cost.

Multi-table Join. Next, we evaluate QUEST on queries with multiple joins. Similarly, we create three groups of queries by the number of filters (F1-F3), and each group has 10 queries. The queries in each group apply the same filter operations, while queries in different groups apply different joins. To evaluate the effectiveness of QUEST, we compare it with (1) Random: each time we randomly select two tables to join using the above technique; (2) Pushdown: we push down all filters first and then run the joins and (3) Optimal: we assume that the selectivity of each filter is known and enumerate all possible join orders to obtain the optimal one.

In Figure 6-b, the system outperforms Random and Pushdown, comparable to the optimal plan in LLM cost. It outperforms Random

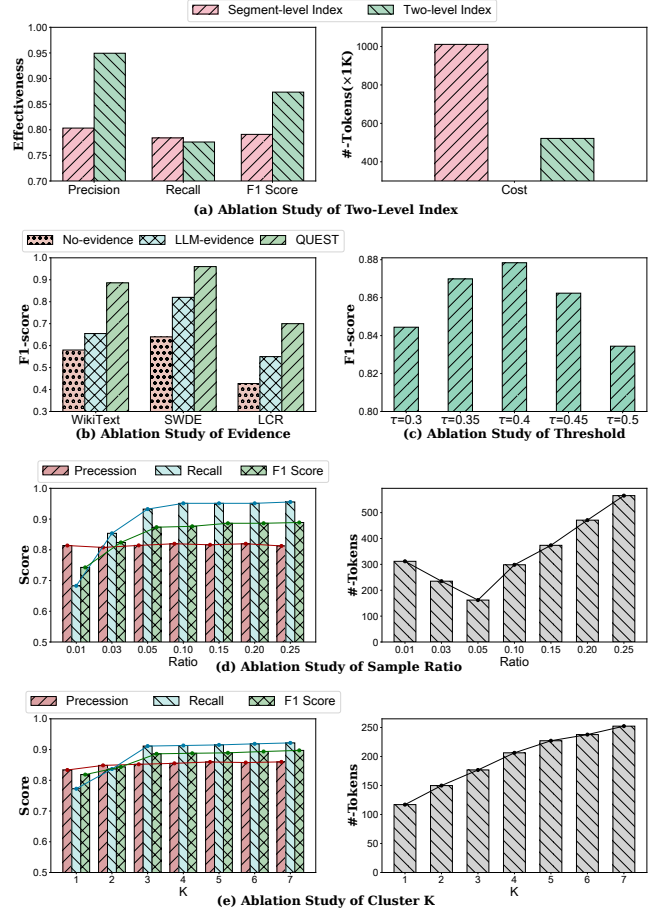


Figure 7: Ablation Studies

because during query execution, QUEST dynamically selects the join operation that leads to the lowest cost. QUEST outperforms Pushdown since in unstructured document analysis, pushing down filters first does not always improve the efficiency of join queries.

5.5 Ablation Studies

Ablation Study of Two-Level Index. We perform ablation studies on WikiText to validate our index’s effectiveness, comparing it with a baseline using only the segment-level index. We measure and average the total query cost across documents. As shown in Figure 7-a, our two-level index improves F1-score and reduces cost, as the segment-level index selects irrelevant documents, lowering precision and raising costs due to unnecessary document processing.

Ablation Study of Evidence. We evaluate QUEST against two baselines: No-evidence, relying solely on the attribute and description, and LLM-evidence, using LLM-generated synthetic text to enhance the query. Figure 7-b shows QUEST achieves higher accuracy by leveraging document-based evidence for better structure and semantics reflection, leading to more precise retrieval.

Ablation Study of τ . We evaluate the effectiveness of our strategy of setting the threshold τ automatically. Given a query about NBA

players, our adaptive strategy sets $\tau = 0.4$. Then, we vary τ around 0.4. In Figure 7-c, if τ is large, the accuracy decreases because more irrelevant documents are retrieved. If τ is small, the accuracy also decreases because of missing many relevant documents.

Ablation Study of Sample Rate. We perform ablation studies on the WikiText dataset, adjusting the sampling rate near the default 5%. Figure 7-d shows that accuracy initially rises with more samples but levels off quickly. Costs decrease first due to better selectivity estimation, then increase with excessive sampling due to LLM overhead. Overall, 5% efficiently balances quality and cost.

Ablation Study of Cluster K. We examine the WikiText dataset by varying the cluster count K near the default of 3. Figure 7-e demonstrates that while accuracy rises with additional clusters providing richer evidence, it soon levels off due to limited extra information. Precision remains stable, depending largely on extraction strategies. Costs rise with K as more clusters bring more evidence vectors and retrieved segments, increasing token usage.

6 RELATED WORK

Language Models for Multi-Modal Data Analysis. Many works focus on analyzing various types of data from diverse sources. Lotus [35] introduces several semantic operators to facilitate bulk semantic processing, including searching, extraction, and indexing, which can be used to build complex pipelines. However, it routinely utilizes LLMs to analyze the full text, incurring significant LLM costs. Some other works [8, 25, 41, 42] also apply LLMs or pre-trained language models to analyze unstructured documents. Like Lotus, they do not focus on optimizing the language model cost.

Palimpzest [30] analyzes unstructured data with a declarative language. Its optimizer produces an execution plan that uses LLMs to extract and analyze the data. However, its optimizations mainly focus on choosing suitable LLMs for tasks, code synthesis, or prompting strategies, which, although effective, are unrelated to classical query optimization principles such as filter ordering or join optimization. CAESURA [43] uses LLMs for natural language analysis over multi-modal data, decomposing queries into operators and invoking models like VisualQA (image-based) and TextQA (text-based) for question answering on different multi-modal tasks, without optimizing LLM costs. UQE [10] leverages LLMs for SQL-like analysis over multi-modal data, primarily enhancing aggregation queries with a sampling technique without considering filter ordering or join optimization.

Retrieval Augmented Generation. RAG [6, 13, 17, 27, 40] is a popular method for tasks like question answering and is well-suited for attribute extraction, a key operation in QUEST, because it retrieves relevant segments using an index. However, Sec. 5.2 shows that the typical RAG’s generic index-based retrieval is less effective than our customized two-level index and evidence augmented strategy.

Text-to-Table Extraction. Some works focus on extracting structured information from unstructured data. ZenDB [29] constructs a semantic hierarchical tree within each document to identify the sections that potentially contain a target attribute. Then, a single matching sentence, as well as several summaries within a section, are fed into an LLM for data extraction. It also uses several optimizations, including filter ordering and predicate pushdown. However, ZenDB lacks fine-grained document-level query optimization and

its index heavily relies on templated document structures. Unfortunately, in practice, abundant of documents lack a clear hierarchical structure. In particular, for documents with lengthy paragraphs, it is rather difficult to simply rely on several summaries and one sentence that matches the text of a query to identify the target attribute. EVAPORATE [3] uses LLMs to extract tables from HTML and PDF files through LLM-based code generation and adopts weak supervision to combine extraction functions. This aims to balance cost and quality. However, relying solely on LLM-generated code for complex documents is not highly accurate, as shown in Sec. 5.2.

Another line of work focuses on training models. Closed information extraction [21–23] uses language models to extract query-relevant information from context. Wu et al. [45] define text-to-table conversion as a sequence-to-sequence task, improving a pretrained model. Pietruszka et al. [37] uses a permutation-based decoder for text-to-table models, enhancing tasks such as entity extraction. Jiao et al. [24] finetune a pretrained model for instruction-following in text-to-table tasks. It extracts structured data, but lacks LLMs’ accuracy and cross-domain generalization [3].

LLMs for Data Preparation. Data preparation is a data processing sequence that transforms raw data into an analyzable format, involving tasks such as extraction, discovery, cleaning, integration, and labeling. Lately, LLMs have excelled in data preparation. Several works finetune LLMs to enhance their abilities in conducting data preparation [28, 47]. Other works utilize LLMs to address specific tasks, such as data extraction [3], schema matching [32, 34], data cleaning (including data imputation [11, 20], entity resolution [12], etc), and data labeling [5, 46].

7 CONCLUSION & FUTURE WORK

We propose QUEST, a cost-effective LLM-powered system that features novel query optimizations to support unstructured document analysis. By introducing a two-level index, an evidence augmented retrieval strategy, and instance-optimized query execution, QUEST effectively reduces the LLM cost while maintaining high accuracy. Our comprehensive experiments showcase the efficacy of QUEST, achieving 30%-6 \times cost saving, while improving F1-score much.

Future research can focus on extending QUEST to handle aggregation queries more efficiently. By leveraging approximate query processing, it is possible to estimate results using sampled documents. Moreover, creating summaries for various attributes could allow for direct aggregation from these summaries. These strategies would significantly enhance QUEST’s ability to analyze unstructured documents in depth.

ACKNOWLEDGMENTS

Chengliang Chai is supported by the National Key Research and Development Program of China (2024YFC3308200), the NSF of China (62472031), Beijing Nova Program, CCF-Baidu Open Fund (CCF-Baidu202402), and Huawei. Ye Yuan is supported by the Beijing Natural Science Foundation (L241010), the National Key Research and Development Program of China (2022YFB2702100), and the NSFC (61932004, 62225203, U21A20516). Guoren Wang is supported by the NSFC (62427808, U2001211), and the Liaoning Revitalization Talents Program (XLYC2204005). Lei Cao is supported by the NSF (DBI-2327954) and Amazon Research Awards.

REFERENCES

- [1] [n.d.]. https://anonymous.4open.science/r/QUEST/Full_version.pdf
- [2] 2019. <https://solutionsreview.com/data-management/80-percent-of-your-data-will-be-unstructured-in-five-years/>
- [3] Simran Arora, Brandon Yang, Sabri Eyuboglu, Avanika Narayan, Andrew Hojell, Immanuel Trummer, and Christopher Ré. 2023. Language Models Enable Simple Systems for Generating Structured Views of Heterogeneous Data Lakes. *Proc. VLDB Endow.* 17, 2 (Oct. 2023), 92–105. <https://doi.org/10.14778/3626292.3626294>
- [4] Dhananjay Ashok and Zachary C. Lipton. 2023. PromptNER: Prompting For Named Entity Recognition. (May 2023).
- [5] Parikshit Bansal and Amit Sharma. 2023. Large language models as annotators: Enhancing generalization of nlp models at minimal cost. *arXiv preprint arXiv:2306.15766* (2023).
- [6] Deng Cai, Yan Wang, Lemaou Liu, and Shuming Shi. 2022. Recent Advances in Retrieval-Augmented Text Generation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. <https://doi.org/10.1145/3477495.3532682>
- [7] Surajit Chaudhuri. 1998. An overview of query optimization in relational systems. In *Proceedings of the seventeenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*. 34–43.
- [8] Zui Chen, Zihui Gu, Lei Cao, Ju Fan, Sam Madden, and Nan Tang. [n.d.]. Symphony: Towards Natural Language Query Answering over Multi-modal Data Lakes. ([n.d.]).
- [9] Hanjun Dai, Bethany Yixin Wang, Xingchen Wan, Bo Dai, Sherry Yang, Azade Nova, Pengcheng Yin, Phitchaya Mangpo Phothilimthana, Charles Sutton, and Dale Schuurmans. 2024. UQE: A Query Engine for Unstructured Databases. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. <https://openreview.net/forum?id=t7SGOv5W5z>
- [10] Hanjun Dai, Bethany Yixin Wang, Xingchen Wan, Bo Dai, Sherry Yang, Azade Nova, Pengcheng Yin, Phitchaya Mangpo Phothilimthana, Charles Sutton, and Dale Schuurmans. 2024. UQE: A Query Engine for Unstructured Databases. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. <https://openreview.net/forum?id=t7SGOv5W5z>
- [11] Zhicheng Ding, Jiahao Tian, Zhenkai Wang, Jinman Zhao, and Siyang Li. 2024. Data imputation using large language model to accelerate recommendation system. *arXiv preprint arXiv:2407.10078* (2024).
- [12] Meihao Fan, Xiaoyue Han, Ju Fan, Chengliang Chai, Nan Tang, Guoliang Li, and Xiaoyong Du. 2024. Cost-effective in-context learning for entity resolution: A design space exploration. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*. IEEE, 3696–3709.
- [13] Wenqi Fan, Yujuan Ding, Liangbo Ning, Shijie Wang, Hengyun Li, Dawei Yin, Tat-Seng Chua, and Qing Li. 2024. A survey on rag meeting llms: Towards retrieval-augmented large language models. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 6491–6501.
- [14] Filippo Galgani and Achim Hoffmann. 2010. LEXA: Towards Automatic Legal Citation Classification. In *AI 2010: Advances in Artificial Intelligence (Lecture Notes in Computer Science)*, Jiuyong Li (Ed.), Vol. 6464. Springer Berlin Heidelberg, 445–454.
- [15] Dawei Gao, Haibin Wang, Yaliang Li, Xiuyu Sun, Yichen Qian, Bolin Ding, and Jingren Zhou. [n.d.]. Text-to-SQL Empowered by Large Language Models: A Benchmark Evaluation. ([n.d.]).
- [16] Yingqi Gao, Yifu Liu, Xiaoxia Li, Xiaorong Shi, Yin Zhu, Yiming Wang, Shiqi Li, Wei Li, Yuntao Hong, Zhiling Luo, et al. 2024. XiYan-SQL: A Multi-Generator Ensemble Framework for Text-to-SQL. *arXiv preprint arXiv:2411.08599* (2024).
- [17] Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and Haofen Wang. 2023. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997* (2023).
- [18] Tiezheng Ge, Kaiming He, Qifa Ke, and Jian Sun. 2013. Optimized product quantization. *IEEE transactions on pattern analysis and machine intelligence* 36, 4 (2013), 744–755.
- [19] Qiang Hao, Rui Cai, Yanwei Pang, and Lei Zhang. 2011. From one tree to a forest. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*. <https://doi.org/10.1145/2009916.2010020>
- [20] Ahatsham Hayat and Mohammad Rashedul Hasan. 2024. CLAIM Your Data: Enhancing Imputation Accuracy with Contextual Large Language Models. *arXiv preprint arXiv:2405.17712* (2024).
- [21] Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2021. DeBERTaV3: Improving deberta using electra-style pre-training with gradient-disentangled embedding sharing. *arXiv preprint arXiv:2111.09543* (2021).
- [22] Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. DeBERTa: Decoding-enhanced BERT with Disentangled Attention. *Cornell University - arXiv, Cornell University - arXiv* (Jun 2020).
- [23] Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. DEBERTA: DECODING-ENHANCED BERT WITH DISENTANGLED ATTENTION. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=XPZlaotutsD>
- [24] Yizhu Jiao, Ming Zhong, Sha Li, Ruining Zhao, Siru Ouyang, Heng Ji, and Jiawei Han. 2023. Instruct and extract: Instruction tuning for on-demand information extraction. *arXiv preprint arXiv:2310.16040* (2023).
- [25] Saehan Jo and Immanuel Trummer. 2024. ThalamusDB: Approximate Query Processing on Multi-Modal Data. *Proceedings of the ACM on Management of Data* 2, 3 (2024), 1–26.
- [26] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Kuttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems (Vancouver, BC, Canada) (NIPS '20)*. Curran Associates Inc., Red Hook, NY, USA, Article 793, 16 pages.
- [27] Huayang Li, Yixuan Su, Deng Cai, Yan Wang, and Lemaou Liu. [n.d.]. A Survey on Retrieval-Augmented Text Generation. ([n.d.]).
- [28] Peng Li, Yeye He, Dror Yashar, Weiwei Cui, Song Ge, Haidong Zhang, Danielle Rifinski Fainman, Dongmei Zhang, and Surajit Chaudhuri. 2023. Tablegpt: Table-tuned gpt for diverse table tasks. *arXiv preprint arXiv:2310.09263* (2023).
- [29] Yiming Lin, Madelon Hulsebos, Ruiying Ma, Shreya Shankar, Sepanta Zeigham, Aditya G Parameswaran, and Eugene Wu. 2024. Towards Accurate and Efficient Document Analytics with Large Language Models. *arXiv preprint arXiv:2405.04674* (2024).
- [30] Chunwei Liu, Matthew Russo, Michael Cafarella, Lei Cao, Peter Baile Chen, Zui Chen, Michael Franklin, Tim Kraska, Samuel Madden, Rana Shahout, and Gerardo Vitagliano. [n.d.]. Palimpsest: Optimizing AI-Powered Analytics with Declarative Query Processing. In *Proceedings of the Conference on Innovative Database Research (CIDR)* (2025).
- [31] Xinyu Liu, Shuyu Shen, Boyan Li, Peixian Ma, Runzhi Jiang, Yuxin Zhang, Ju Fan, Guoliang Li, Nan Tang, and Yuyu Luo. 2024. A Survey of NL2SQL with Large Language Models: Where are we, and where are we going? *arXiv preprint arXiv:2408.05109* (2024).
- [32] Yurong Liu, Eduardo Pena, Aecio Santos, Eden Wu, and Juliana Freire. 2024. Magneto: Combining Small and Large Language Models for Schema Matching. *arXiv:2412.08194 [cs.DB]* <https://arxiv.org/abs/2412.08194>
- [33] Yu A Malkov and Dmitry A Yashunin. 2018. Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. *IEEE transactions on pattern analysis and machine intelligence* 42, 4 (2018), 824–836.
- [34] Marcel Parciak, Brecht Vandevort, Frank Neven, Liesbet M. Peeters, and Stijn Vansummeren. 2024. Schema Matching with Large Language Models: an Experimental Study. *arXiv:2407.11852 [cs.DB]* <https://arxiv.org/abs/2407.11852>
- [35] Liana Patel, Siddharth Jha, Carlos Guestrin, and Matei Zaharia. 2024. LOTUS: Enabling Semantic Queries with LLMs Over Tables of Unstructured and Structured Data. *arXiv preprint arXiv:2407.11418* (2024).
- [36] PengLi and TianxiangSun et al. 2023. CodeIE: Large Code Generation Models are Better Few-Shot Information Extractors. (May 2023).
- [37] Michał Pietruszka, Michał Turski, Łukasz Borchmann, Tomasz Dwojak, Gabriela Palka, Karolina Szyndler, Dawid Jurkiewicz, and Łukasz Garnarek. 2022. Stable: Table generation framework for encoder-decoder models. *arXiv preprint arXiv:2206.04045* (2022).
- [38] Bruno R Preiss. 1999. *Data structures and algorithms*. John Wiley & Sons, Inc.
- [39] Oscar Sainz, Iker Garcia-Ferrero, Rodrigo Agerri, OierLopezde Lacalle, German Rigau, and Eneko Agirre. 2023. GoLLIE: Annotation Guidelines improve Zero-Shot Information-Extraction. (Oct 2023).
- [40] Parth Sarthi, Salman Abdullah, Aditi Tuli, Shubh Khanna, Anna Goldie, and ChristopherD. Manning. 2024. RAPTOR: Recursive Abstractive Processing for Tree-Organized Retrieval. (Jan 2024).
- [41] James Thorne, Majid Yazdani, Marzieh Saeidi, Fabrizio Silvestri, Sebastian Riedel, and Alon Halevy. 2021. From natural language processing to neural databases. *Proceedings of the VLDB Endowment* (Feb 2021), 1033–1039. <https://doi.org/10.14778/3447689.3447706>
- [42] Matthias Urban and Carsten Binnig. [n.d.]. Towards Multi-Modal DBMSs for Seamless Querying of Texts and Tables. ([n.d.]).
- [43] Matthias Urban and Carsten Binnig. 2023. CAESURA: Language Models as Multi-Modal Query Planners. *arXiv preprint arXiv:2308.03424* (2023).
- [44] Liang Wang, Nan Yang, Xiaolong Huang, Jiao Bingxing, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. 2022. Text Embeddings by Weakly-Supervised Contrastive Pre-training. *Cornell University - arXiv, Cornell University - arXiv* (Dec 2022).
- [45] Xueqing Wu, Jiacheng Zhang, and Hang Li. 2022. Text-to-Table: A New Way of Information Extraction. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. <https://doi.org/10.18653/v1/2022.acl-long.180>
- [46] Ruixuan Xiao, Yiwen Dong, Junbo Zhao, Runze Wu, Minmin Lin, Gang Chen, and Haobo Wang. 2023. FreeAL: Towards Human-Free Active Learning in the Era of Large Language Models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. 14520–14535.

- [47] Haochen Zhang, Yuyang Dong, Chuan Xiao, and Masafumi Oyamada. 2024. Jellyfish: Instruction-Tuning Local Large Language Models for Data Preprocessing. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (Eds.). Association for Computational Linguistics, Miami, Florida, USA, 8754–8782. <https://doi.org/10.18653/v1/2024.emnlp-main.497>