



# Revisiting Task-Oriented Dataset Search in the Era of Large Language Models: Challenges, Benchmark, and Solution

Zixin Wei

The Chinese University of Hong  
Kong, Shenzhen  
zixinwei1@link.cuhk.edu.cn

Yucan Guo\*

Institute of Computing Technology,  
Chinese Academy of Sciences  
guoyucan23z@ict.ac.cn

Jinyang Li

The University of Hong Kong  
jl0725@connect.hku.hk

Xiaolin Han<sup>†</sup>

The Northwestern Polytechnical  
University  
xiaolinhan@nwpu.edu.cn

Xiaolong Jin\*

Institute of Computing Technology,  
Chinese Academy of Sciences  
jinxiaolong@ict.ac.cn

Chenhao Ma<sup>†</sup>

The Chinese University of Hong  
Kong, Shenzhen  
machenhao@cuhk.edu.cn

## ABSTRACT

The search for suitable datasets is the critical “first step” in data-driven research, but it remains a great challenge. Researchers often need to search for datasets based on high-level task descriptions. However, existing search systems struggle with this task due to ambiguous user intent, task-to-dataset mapping and benchmark gaps, and entity ambiguity. To address these challenges, we introduce KATS, a novel end-to-end system for task-oriented dataset search from unstructured scientific literature. KATS consists of two key components, i.e., offline knowledge base construction and online query processing. The sophisticated offline pipeline automatically constructs a high-quality, dynamically updatable task-dataset knowledge graph by employing a collaborative multi-agent framework for information extraction, thereby filling the task-to-dataset mapping gap. To further address the challenge of entity ambiguity, a unique semantic-based mechanism is used for task entity linking and dataset entity resolution. For online retrieval, KATS utilizes a specialized hybrid query engine that combines vector search with graph-based ranking to generate highly relevant results. Additionally, we introduce CS-TDS, a tailored benchmark suite for evaluating task-oriented dataset search systems, addressing the critical gap in standardized evaluation. Experiments on our benchmark suite show that KATS significantly outperforms state-of-the-art retrieval-augmented generation frameworks in both effectiveness and efficiency, providing a robust blueprint for the next generation of dataset discovery systems.

## KEYWORDS

Dataset Search, Task-Oriented Search, Knowledge Graph

### PVLDB Reference Format:

Zixin Wei, Yucan Guo, Jinyang Li, Xiaolin Han, Xiaolong Jin, and Chenhao Ma. Revisiting Task-Oriented Dataset Search in the Era of Large Language Models: Challenges, Benchmark, and Solution. PVLDB, 19(5): 973 - 986, 2026.

doi:10.14778/3796195.3796209

\*Also affiliated with School of Computer Science and Technology, University of Chinese Academy of Sciences.

<sup>†</sup>Corresponding authors.

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](mailto:info@vldb.org). Copyright is held by the owner/author(s). Publication rights

### PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/starkersawz666/KATS>.

## 1 INTRODUCTION

The search for suitable datasets is a foundational yet often arduous step in any data-driven scientific investigation. In response to this pressing need, the problem of dataset discovery has attracted growing attention from both academia and industry, leading to the development of a range of data search systems [3, 5, 17]. However, the dominant paradigm among these systems remains keyword-based retrieval, which implicitly assumes that users can articulate their data needs using structured metadata or keyword fields. This approach presumes that users have already translated their high-level research objectives into a precise data specification.

In contrast, recent empirical studies reveal a stark misalignment between this assumption and real-world behavior: for 79% of data professionals, identifying a relevant dataset is the single most common and critical task [17]. This underscores the importance of *task-oriented* dataset search, i.e., retrieving datasets based on natural language task descriptions, which remains a long-standing and unsolved challenge [7].

We identify three fundamental obstacles that hinder progress in this domain:

- **Ambiguous User Intent.** A substantial “semantic gap” exists between users’ high-level goals and the low-level keyword queries typically required by search engines. This misalignment often yields irrelevant results [7].
- **Task-to-Dataset Mapping and Benchmark Gaps.** Even when a system can interpret user intent, there exists no explicit mapping between scientific tasks and their associated datasets. Moreover, the lack of standardized benchmarks has hindered rigorous evaluation and comparative study of task-oriented dataset search systems.
- **Entity Ambiguity.** Datasets are inconsistently referenced across the literature, e.g., by acronym, full name, or version, which leads to fragmented results and low recall [7].

licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 19, No. 5 ISSN 2150-8097.  
doi:10.14778/3796195.3796209

Fortunately, recent technological advances offer new opportunities to tackle these challenges. Large Language Models (LLMs) have demonstrated impressive capabilities in semantic understanding and reasoning [39], making them well-suited for interpreting complex task descriptions. In parallel, the Retrieval-Augmented Generation (RAG) paradigm [23] provides a flexible mechanism for grounding LLMs in external knowledge sources.

Despite these advances, existing RAG frameworks are typically designed for general-purpose question answering, not for task-oriented dataset discovery. For instance, HippoRAG focuses on multi-hop reasoning [14], while Raptor’s hierarchical design may miss cross-document associations [24]. Crucially, core issues such as entity ambiguity, task-dataset mapping, and relevance ranking in the context of dataset retrieval remain under-addressed. Furthermore, building and operating RAG systems at scale remains time- and resource-intensive.

To address these issues, we present **Knowledge graph-Augmented Task-oriented dataset Search (KATS)**, a purpose-built system for task-driven dataset discovery. KATS operates through a two-stage pipeline: (1) an offline component that constructs and updates a task-dataset Knowledge Graph (KG), and (2) an online hybrid query engine that integrates LLM-powered retrieval with graph-based reasoning to surface relevant datasets.

Importantly, we also introduce a standardized **benchmark** for task-oriented dataset search derived from scientific literature. Scientific papers are a natural source for this purpose, as they routinely describe both the task and the datasets used. This implicit alignment enables us to curate realistic task–dataset pairs at scale, offering a reproducible and representative framework for evaluating dataset discovery systems.

Our key contributions are summarized as follows:

- We develop a standardized benchmark suite for task-oriented dataset search, constructed from curated scientific literature. This addresses a longstanding evaluation gap and supports systematic assessment of dataset discovery methods.
- We design a novel end-to-end pipeline that automatically builds a task-dataset KG, augmented with a dedicated entity resolution module to combat dataset naming ambiguities.
- We propose a hybrid query engine that fuses dense retrieval with graph-based relevance ranking, yielding accurate and task-relevant dataset recommendations.
- Extensive experiments demonstrate that KATS outperforms existing LLM-based approaches in both effectiveness and computational efficiency.

The remainder of this paper is structured as follows: Section 2 surveys related work in dataset search and RAG systems. Section 3 defines the task-oriented dataset search problem and presents our benchmark. Section 4–6 detail the system architecture. Section 7 reports experimental results, followed by discussions and conclusions in Sections 9 and 10.

## 2 RELATED WORK

### 2.1 Dataset Search Systems

Dataset search has long been recognized as a fundamental yet challenging problem in both academia and industry. A recent survey shows that the most common goal of data search is to find the

initial dataset for a given task [17]. However, this process is challenging given the existing data quality and search methods. Existing search systems are often hampered by poor metadata quality [27]. Meanwhile, there is a semantic gap between user intent and system capabilities, including understanding user queries in the context of available datasets and matching semantically related but syntactically different datasets [30]. Our work is conducted within this extensive research context, aiming to promote the evolution of the dataset discovery paradigm from simple keyword searches to more intelligent and user-centric approaches.

*2.1.1 General and Web-scale Systems.* In the early stages of dataset search, researchers focused on building web-scale discovery tools, of which Google Dataset Search [5] is the most representative. This system relies on data owners to publish structured metadata in accordance with open metadata standards such as `schema.org` [34] and Croissant [1]. By aggregating this metadata, the system can index massive datasets. However, the fundamental limitation of this approach lies in its dependence on high-quality, manually curated metadata. When metadata is missing or of poor quality, search results are significantly degraded.

*2.1.2 Semantic-based Systems.* With the spread of the “data lake” concept [4], dataset discovery faces new challenges. It has become complicated to discover relevant datasets in an environment where large amounts of heterogeneous datasets are stored in raw format with missing or inconsistent pattern information. In response to these challenges, a series of advanced semantic-based systems has emerged. Aurum [32] is a representative semantic system. It builds an enterprise KG that systematically captures syntactic relationships between data columns, enabling users to perform flexible exploratory queries. Moreover, D3L [4] proposes a distance-based framework that combines multiple types of evidence to discover related tables in data lakes. In addition, Table Union Search [26] enables joint searches of table data in an open data environment, improving the quality of related table discovery and search efficiency. Similarly, Auctus [6] is a search engine that supports data enrichment queries, allowing users to find datasets that can be merged with a given table.

After the emergence of LLMs, systems with LLM integration have demonstrated the power of using LLMs in dataset retrieval. PNEUMA [3] combines LLMs with traditional information retrieval techniques to efficiently and accurately discover relevant tabular data from large collections based on natural language queries. LEDD [2] utilizes LLM to generate a hierarchical global directory automatically and supports semantic table search using natural language. Furthermore, SemExplorer [37] displays a user interface designed specifically for customized semantic dataset search, highlighting the importance of human-computer interaction in solving complex discovery problems. AUTOTUS [15] combines the contextual understanding capabilities of LLM with its multi-stage training strategies to achieve deep representation learning of complex relationships between tables, which significantly improves the accuracy and robustness of automatic joint table search. This series of research studies clearly demonstrates a paradigm shift from reliance on external metadata to a deeper understanding of data content. However, this is still insufficient to fully satisfy users’ needs to find datasets that serve their specific “tasks”.

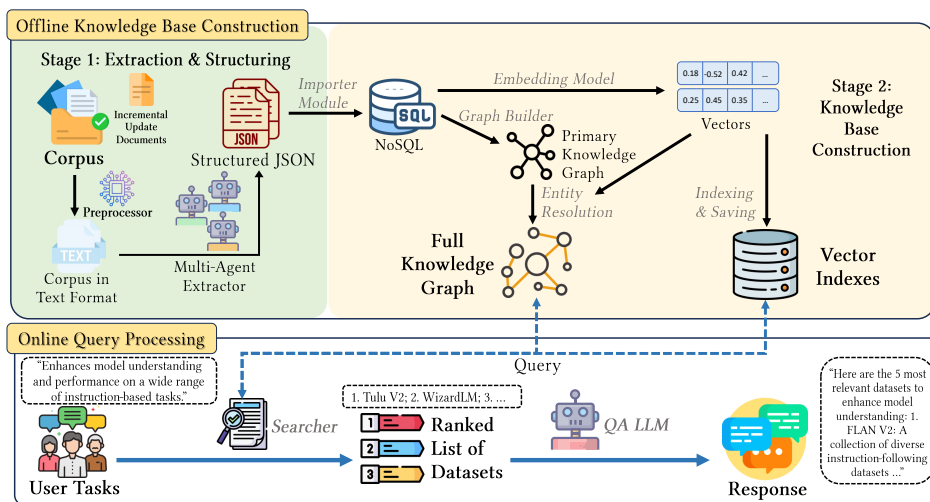


Figure 1: Overall architecture of the KATS system.

2.1.3 *Task-Driven Dataset Search.* The task-driven paradigm will be a key shift in the field of dataset search, motivated by the real-world work practices of data professionals. The recent survey of data professionals [17] identified task-driven search as one of the four key desiderata for next-generation dataset search systems. This highlights a critical need derived directly from user practice: users are often unclear about the name of the dataset they need, but rather seek datasets suitable for accomplishing a specific task.

Recent research has begun to explore this task-driven paradigm. For instance, METAM [12] proposes a goal-oriented framework where the “task” is a computable Machine Learning model. The system iteratively augments a dataset to improve a quantitative utility score, such as the F1-score. Similarly, Mileena [16] also defines the task as an ML model, but focuses on solving the critical challenges of search latency and data privacy by using semi-ring sketches.

While these systems are powerful and address important aspects of the problem, their definition of a “task” relies on a formal and quantifiable function. In contrast, our work tackles a more common scenario where a user’s task is expressed as an open-ended natural language description. This formulation requires a fundamentally different approach centered on a deep semantic understanding of the user’s intent and its relationship to the knowledge within the unstructured text. Addressing this challenge of task-driven search from natural language is the primary motivation for our work.

## 2.2 Knowledge Graph-based RAG

The advent of RAG has opened up new areas for knowledge-intensive tasks. Recent research has gone beyond simple vector-based retrieval by proposing more sophisticated architectures that allow the construction of knowledge corpora for more efficient retrieval. Therefore, these approaches are applicable to the task of dataset search from unstructured scientific literature. We position our work in relation to several of these state-of-the-art frameworks, which we also employ as strong baselines in our experiments.

These advanced systems represent different strategies for structuring knowledge. GraphRAG [11] constructs an entity-based KG

from the source documents using LLM and identifies communities to generate abstractive summaries for answering broad questions. The HippoRAG series of works presents a neurobiologically inspired framework for enhancing the reasoning capabilities of RAG. The original HippoRAG [14] introduced the core idea of using the Personalized PageRank on an LLM-constructed KG to enable multi-hop retrieval. Then, its successor HippoRAG 2 [20] builds upon this foundation to address several limitations of the original entity-centric approach. Specifically, HippoRAG 2 introduces passage nodes and a more sophisticated query-to-triple matching mechanism for deeper contextualization. It also incorporates a recognition memory module to improve the retrieval accuracy and efficiency. RAPTOR [33] recursively clusters and summarizes text chunks to build a hierarchy, allowing the system to retrieve information at varying levels of abstraction for general question answering over long documents.

While these systems are powerful general-purpose frameworks, they are not specifically optimized for the unique challenges of dataset discovery from scientific literature. In contrast, our proposed KATS system is an end-to-end pipeline tailored for this specific task. It incorporates specialized mechanisms, such as resolving entity ambiguity and datasets ranking, which are not the primary focus of these more general RAG architectures.

## 3 PROBLEM DEFINITION

This section provides a formal definition of our primary problem, Task-Oriented Dataset Search problem. Subsequently, we describe the construction and characteristics of our new benchmark suite, which was created to rigorously evaluate systems on this specific task.

### 3.1 The Task-Oriented Dataset Search Problem

Let  $C = \{c_1, c_2, \dots, c_n\}$  be a corpus of unstructured scholarly documents. Given a task query  $T$  in natural language from a user, the goal of a Task-Oriented Dataset Search system is to return a ranked list of datasets  $D = \{d_1, d_2, \dots, d_k\}$  such that each dataset

$d_i$  corresponds to a unique dataset entity within the corpus  $C$ , and is relevant to solving the task described by  $T$ . The objective of the system is to maximize the relevance of the returned list of datasets  $D$ . The relevance of a dataset to a task query is defined according to the ground-truth annotation criteria established in our benchmark suite (see Section 3.2).

### 3.2 Benchmark Suite

While several benchmarks exist for question answering (e.g., HotPotQA [38]) or general document retrieval (e.g., BEIR [35]), there is a lack of established benchmarks specifically designed for evaluating task-oriented dataset search over a large corpus of scientific documents. For a rigorous and reproducible evaluation, we constructed a new benchmark suite, which we refer to as the Computer Science Task Dataset Search (CS-TDS).

**3.2.1 Corpus Curation.** Our suite consists of two benchmarks built upon corpora of different scales to test system performance and scalability, with their detailed statistics summarized in Table 1. The first one, CS-TDS<sub>M</sub>, is constructed from a corpus of 628 computer science papers, and the second, larger benchmark CS-TDS<sub>L</sub> is built from 2101 papers. These papers were obtained using web crawlers from computer science conferences, e.g., NeurIPS.

**Table 1: Statistics of CS-TDS benchmark suite.**

Statistic	CS-TDS <sub>M</sub>	CS-TDS <sub>L</sub>
# Papers in Corpus	628	2,101
# Task Queries	47	204
# Datasets	~ 1,779	~ 7,525
# Unique Datasets	~ 1,217	~ 4,180
Year Span	2022, 2023	2024

**3.2.2 Query Generation.** The queries are generated as natural language descriptions of a task derived from real papers by an LLM and manually verified. 47 queries are selected for CS-TDS<sub>M</sub>, and 204 are selected for CS-TDS<sub>L</sub>. To ensure the evaluation tests a system’s generalization ability rather than its ability to simply retrieve the source document, the paper from which the query and ground-truth answer were derived is held out from the search corpus. Below is a representative example from our benchmark:

#### Benchmark Query Example

**Task Query:** “Assess HSIVI-SM’s ability to accelerate diffusion model sampling while maintaining sample quality, measured via Fréchet Inception Distance (FID).”

**Ground-Truth Dataset:** “CIFAR-10”, which contains 60,000 32 × 32 color images in 10 different classes [22].

**3.2.3 Ground-Truth Annotation.** For each query, the standard answer is set as the corresponding dataset actually used in the paper. To avoid “false negatives”, other reasonable answers besides the original answer are also considered correct answers.

Formally, a dataset  $d_i$  is considered *relevant* to a task query  $T$  if it satisfies one of the following conditions during ground-truth annotation:

- It is **identical** to the dataset used in the source paper for  $T$ .
- It is a **variant or alias** of the dataset (e.g., MIMIC-III vs. Medical Information Mart for Intensive Care III).
- It is a **functionally equivalent substitute**, typically from the same dataset family or with the same task objective (e.g., LAION-5B vs. LAION-400M for large-scale image-text pretraining).

These criteria and the corresponding annotations are manually validated and used to determine the binary *relevance label* during evaluation. Systems are scored based on how well they retrieve such relevant datasets for each query.

## 4 SYSTEM OVERVIEW

In this section, we introduce the design of our proposed KATS in detail. Firstly, we present the overall architecture and the function of each component. Then, we formally define the data model of our KG, which is the core data structure managed by the system.

### 4.1 Overall Architecture

As illustrated in Figure 1, the architecture of KATS is divided into two major phases: an offline phase for knowledge base construction and an online phase for query processing.

**4.1.1 Offline Knowledge Base Construction.** The offline phase begins with Stage 1: Extraction & Structuring. In this process, the input papers are transformed into a structured format for the following knowledge base construction. The Preprocessor module normalizes heterogeneous source documents (mostly in PDF format) into the plain text representation. Then, the downstream Extractor module uses an LLM to extract task and dataset information into JSON format.

Structured information is then passed to Stage 2: Knowledge Base Construction. It involves two interconnected workflows after structured information is imported into the database. The vector workflow generates dense vector representations for tasks and datasets, and stores them in the Vector Indexes. Meanwhile, in the graph workflow, the GraphBuilder module populates the initial KG. Subsequently, the Entity Resolution module further refines the KG with the previously generated vector embeddings.

**4.1.2 Online Query Processing.** The online phase is designed for real-time user interaction. When a user submits a Natural Language Task Query, the Hybrid Searcher interacts with the offline knowledge base. It queries the Vector Indexes for semantic similarity and the KG for relational task-task and task-dataset information. Combining these signals, the searcher produces a Ranked List of Datasets that are most relevant to the user’s task. This list is then passed to a downstream QA LLM, which can re-rank the results for better precision and generate a final natural language answer for the user.

### 4.2 Task-Dataset KG

The core data structure managed by KATS is the *Task-Dataset KG*. This graph is formally defined as  $G = (V, E)$ , where  $V$  represents

the set of nodes (vertices) and  $E$  represents the set of undirected edges.

**Table 2: Notations and meanings.**

Notation	Meaning
$G = (V, E, W)$	Undirected, weighted KG with vertex set $V = C \cup D \cup T$ , edge set $E = E_{C,D} \cup E_{D,T} \cup E_{T,T}$ , and weight function $W$ .
$C$	Set of document nodes.
$D$	Set of dataset nodes.
$T$	Set of task nodes.
$E_{C,D}$	Set of edges $(c, d)$ , where $c \in C, d \in D$ .
$E_{D,T}$	Set of edges $(d, t)$ , where $d \in D, t \in T$ .
$E_{T,T}$	Set of edges $(t_1, t_2)$ , where $t_1, t_2 \in T$ .

The set of vertices  $V$  consists of three distinct types of nodes:

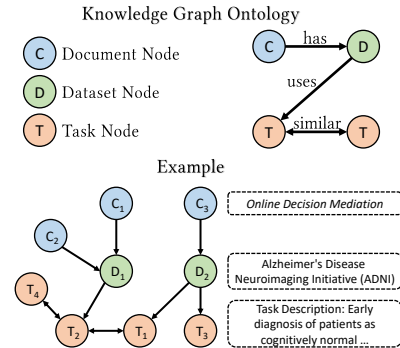
- **Document Node:** Represents a single document (paper) from the input. Each node stores metadata of the document, such as title and source file path.
- **Dataset Node:** Represents a dataset entity used in the document. Each node includes attributes such as the dataset’s name and brief description.
- **Task Node:** Represents a specific task for which a dataset is used, as described within a document. Each node contains attributes such as a textual description of the task.

The set of edges  $E$  captures the relationships between nodes:

- **Document-dataset Edge:** An edge  $(c, d)$  where  $c \in C, d \in D$  signifies that document  $c$  includes dataset  $d$ .
- **Dataset-task Edge:** An edge  $(d, t)$  where  $d \in D, t \in T$  signifies that dataset  $d$  is used for the purpose of task  $t$ .
- **Task-task Edge:** An edge  $(t_1, t_2)$  where  $t_1, t_2 \in T$  with weight  $w$  signifies that task  $t_1$  and  $t_2$  are semantically similar with cosine similarity  $w$ .

## 5 KNOWLEDGE BASE CONSTRUCTION

In this section, we will introduce the details of the offline pipeline for constructing the task-dataset knowledge base. While contemporary graph-based RAG systems also construct KGs from text, their methodology is often tailored for general-purpose summarization. However, for precise task-oriented dataset search, this approach can be suboptimal, as critical details such as the context of a task may be lost. Furthermore, these general frameworks often lack the dedicated mechanisms required to address the entity ambiguity problem. To overcome these limitations, our pipeline is specifically designed for the precise extraction, curation, and querying of dataset and task entities. The process starts with raw documents and produces a high-quality, queryable knowledge base within a series of steps. Furthermore, to ensure the long-term viability and scalability of the knowledge base, we designed the system to support efficient incremental updates, allowing it to remain up-to-date without costly periodic rebuilds. Each step will be further discussed in the following subsections. The ontology and an example of the task-dataset KG are demonstrated in Figure 2.



**Figure 2: Ontology and an example of the task-dataset KG.**

### 5.1 Document Normalization and Caching

The first step in the pipeline is to collect and normalize the source documents. The primary goal of this step is to convert a corpus of documents in heterogeneous formats into plain text. The main challenge lies in handling the complexity of PDF format papers to ensure that the LLM can accurately extract information in the subsequent step.

In our current implementation, we utilize the pypdf [31] library to employ direct text extraction, which is highly efficient. Since our target entities - datasets and their corresponding tasks - primarily appear in the narrative prose of the papers, non-textual elements such as figures and tables play a less critical role in our extraction strategy. Therefore, this method proves sufficient and effective for reliably extracting our target entities. To address more complex cases, such as when many texts are presented in the form of images or complex multi-column layouts, we could consider Optical Character Recognition (OCR) or intelligent document parsing services (e.g., LlamaParse [25]), which is computationally expensive. This step is completed in a completely separate module. Thus, the specific method can be replaced depending on the condition of the source documents to increase efficiency and save costs.

To optimize performance, especially when dealing with large, static corpora, we implement a caching strategy. Before processing each file, the system computes a unique hash value of the file’s content as its “fingerprint”. On subsequent runs, if a file with an identical fingerprint is encountered, the content will be retrieved directly from cache. This prevents costs associated with duplicate computing. We utilize the same caching approach in the subsequent information extraction step.

### 5.2 Information Extraction

Once the documents are normalized into plain text, the Extractor module is responsible for identifying and extracting structured information about datasets and their associated tasks. Some contemporary methods, such as GraphRAG [11], perform a one-shot, general-purpose entity extraction. Instead, we design a collaborative multi-agent framework specifically for our goal of task-oriented dataset discovery. The framework decomposes the complex extraction challenge into a sequence of sub-tasks, each handled by a specialized agent. This design not only improves the accuracy and

relevance of the extracted information but also enhances efficiency by avoiding deep processing of irrelevant documents.

**5.2.1 The Preliminary Filtering Agent.** The first agent in our framework acts as an efficient gatekeeper. Its primary responsibility is to rapidly identify documents that are likely to contain dataset mentions to avoid expensive LLM calls on irrelevant content. For each document, this agent queries the LLM with a concise instruction template designed to obtain a simple Boolean response of dataset relevance. Only documents identified to be relevant are passed to the subsequent agent.

**5.2.2 The Analyst Agent.** The Analyst Agent specializes in deep content analysis. It utilizes a detailed instruction template to guide the LLM’s careful reading of the text to help identify all mentions of datasets and their associated tasks. To ensure machine readability, this agent is constrained to generate its output in a predefined structured JSON format. This format consists of a list of objects, where each object represents a unique dataset-task relationship and contains relevant information. Table 3 presents a concrete example of an input text snippet in a paper [18] and the corresponding structured JSON object generated by the agent.

**5.2.3 The Enrichment Agent.** Finally, the structured JSON object is passed to the Enrichment Agent. This agent’s sole purpose is to enhance the semantic metadata of the extracted entities. For each task description, the agent uses a separate instruction template to generate a list of relevant keywords that summarize the core concepts of the task. These keywords will help graph curation in the subsequent stage.

**Table 3: Example of the information extraction process.**

Component	Content
<b>Text Input</b>	In Alzheimers, the task is to perform early diagnosis of patients in the Alzheimer’s Disease Neuroimaging Initiative study as cognitively normal, mildly impaired, or at risk of dementia.
<b>JSON Output</b>	{“dataset_name”: “Alzheimer’s Disease Neuroimaging Initiative”, “task_description”: “early diagnosis of patients as cognitively normal, mildly impaired, or at risk of dementia.”, “task_keywords”: [“early diagnosis”, “dementia”, “patient assessment”]}

### 5.3 Semantic Representation Learning

After extracting structured textual information, the next step is to convert this text into meaningful vector representations. This process maps textual descriptions into a high-dimensional vector space. These vector embeddings are fundamental to our system, for both quantitative analysis of similarity in graph building and the semantic search capabilities in online querying.

In our system, we generate embeddings for two key types of entities: tasks and datasets. For each task node, we embed its task description. For each dataset node, we concatenate its title and

description to create a comprehensive representation. In our implementation, we utilize OpenAI’s text-embedding-3-small model [29] via its API. We selected this model for its compelling balance of high embedding quality, cost-effectiveness, and ease of integration.

The generated embeddings,  $\{e_{task}\}$  and  $\{e_{dataset}\}$ , are then persisted and indexed for efficient retrieval. Specifically, we utilize the Faiss library [10] to build an index from the dense vectors. Faiss enables highly efficient approximate nearest neighbor search. It allows our system to find the most semantically similar items from a large collection in real-time. This indexing step ensures that both the online semantic search and the offline graph-building processes can perform similarity lookup with low latency.

### 5.4 Entity Resolution and Linking

The initial KG constructed from the extraction stage, while comprehensive, still contains ambiguity and redundancy inherent in natural language, especially for the task and dataset nodes:

- Semantically identical tasks with different phrasing, such as “image classification” and “classifying images”.
- Same datasets referred to by different names or acronyms, such as “COCO” and “MS-COCO”.

To enhance the quality and reliability of the knowledge base, we perform a graph curation stage consisting of two main processes: task linking and dataset resolution.

**5.4.1 Task Linking.** To address the ambiguity of task descriptions, our system identifies and connects task nodes that are semantically similar but are described using different phrasing. To achieve this goal, the TaskMerger module uses the vector embeddings  $\{e_{task}\}$ . This process involves Faiss to find pairs of tasks with high cosine similarity of the vectors. If the similarity score between two tasks exceeds a predefined threshold  $\theta_d$ , or the keyword overlap exceeds a threshold  $\theta_k$ , we consider them semantically closely related.

After identifying such pairs, a weighted edge is added to the KG between the corresponding task nodes. The weight of this edge is set to the calculated similarity score, representing the strength of the relationship. When querying, the system is able to broaden the search to a wider set of relevant tasks through the edges.

**5.4.2 Dataset Resolution.** The dataset resolution process tackles the entity redundancy problem, where the same dataset is referred to by different names. An all-pairs comparison of dataset nodes would be  $O(n^2)$ , which is computationally prohibitive. To address this problem, we designed an three-stage approach.

The first stage focuses on efficiently generating candidate pairs for potential merging. Instead of a brute-force comparison, we use the Faiss index. For each dataset, the system performs an approximate nearest neighbor search to retrieve a small set of the most semantically similar datasets. This process reduces the search space from all possible pairs to a manageable set of high-potential candidates. At the beginning, each dataset is stored as a separate node in a disjoint-set data structure.

In the second stage, the system performs the initial matching by normalizing the titles. Specifically, case, spaces, etc., are standardized. In addition, a dictionary is maintained to record the aliases of the dataset. After these processes, datasets with the same normalized title are merged into the disjoint set.

The above approach can partially reduce the problem of the same dataset with different names, such as “Image Net” and “imagenet”. For more complex situations, an LLM works as an expert judge to perform a rigorous verification process in the third stage. The prompt provides the names and descriptions of both datasets and instructs it to return a binary “True” or “False” answer on whether they refer to the same entity. For positive judgments, the system proceeds to merge the two nodes in the KG. Meanwhile, regardless of the judgment, the result is stored to avoid repeated LLM calls.

This hybrid approach combines the scalability of vector search with the reasoning capabilities of LLM. Therefore, KATS provides a robust, efficient solution to the dataset entity resolution problem.

### 5.5 Incremental Knowledge Base Updates

To address the limitation of a static knowledge base and ensure the system remains up-to-date, we designed an efficient incremental update pipeline. This mechanism allows KATS to integrate the newly extracted knowledge into the existing KG and vector indices without a full, costly reconstruction of the entire knowledge base.

The pipeline first identifies new documents by maintaining a manifest of all previously processed files. Each file is identified by its unique content-based fingerprint, which is generated by extracting text from the file and applying the SHA256 algorithm. When new documents are added to the corpus, the system computes their fingerprints and compares them against the manifest to create a queue of only the new files. Each new document undergoes the pre-processing and multi-agent extraction pipeline, and the extracted structured data is then ingested into the database. This process is inherently incremental since new entries are simply added to their respective tables without affecting existing data.

To update the semantic search indices, vector embeddings are generated for these new entries and are appended to the existing Faiss indices. Afterwards, new nodes, including documents, datasets, tasks, and their basic relationships, are added to the existing KG. Subsequently, the entity linking and resolution modules will operate in an incremental mode. For each new task and dataset node, the system computes its similarity against the entire index to establish new “similar\_task” edges and identify potential dataset merges. This approach ensures that new entities are correctly integrated into the existing graph structure while preserving all previously computed relationships. This entire pipeline ensures that KATS is not merely a static repository, but a dynamic system that evolves its knowledge over time.

Furthermore, to utilize broader dataset information beyond scientific literature, the incremental update feature also incorporates the ability to integrate datasets lacking explicit task descriptions. By embedding the new description and identifying the most similar dataset within the existing knowledge graph, KATS can automatically link the new entry to relevant tasks, thus enriching the knowledge base from a wider range of sources.

## 6 QUERY PROCESSING

Figure 3 illustrates the entire workflow of the online query processing pipeline, which is triggered when a user submits a natural

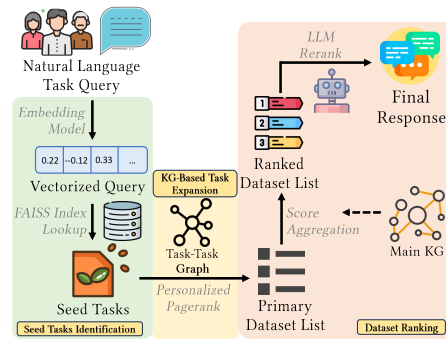


Figure 3: The query processing pipeline of KATS system.

language task query. Our approach is designed as a multi-stage funnel that progressively refines the search space from broad semantic matching to fine-grained graph-based ranking.

### 6.1 Seed Task Identification

The online query processing pipeline begins with the seed task identification stage. The primary objective of this stage is to efficiently bridge the semantic gap between the user’s natural language query and the structured entities within our KG. It achieves this by identifying an initial set of “seed” tasks that are most semantically relevant to the user’s query.

This stage involves two main steps. Firstly, the user’s task query  $T$  is transformed into a high-dimensional vector representation  $\mathbf{t}$ . This is accomplished using the same pre-trained embedding model as used during the offline knowledge base construction. This ensures that both the query and the knowledge base are in the same semantic space for meaningful comparisons. Next, the query vector  $\mathbf{t}$  is used to perform an Approximate Nearest Neighbor (ANN) search against the pre-built FAISS index of task embeddings. This search efficiently retrieves the top- $k$  task nodes from the KG, which we refer to as the “seed tasks” ( $T_{seed}$ ).  $T_{seed}$  serves as the crucial starting point for the subsequent KG-based task expansion to discover a wider range of relevant information.

### 6.2 KG-based Task Expansion

The initial set of seed tasks,  $T_{seed}$ , while semantically close to the query, may not be exhaustive. The KG-based task expansion stage uses the relational information in our task-task similarity graph to discover a wider set of relevant tasks. Our approach employs the Personalized PageRank (PPR) algorithm [19] for this expansion. PPR allows us to compute node relevance relative to a specific set of “source” nodes. In our context, the algorithm is executed on the task-task similarity graph. The previously identified seed tasks  $T_{seed}$  are used as the personalization vector for PPR, which means the “random walk” process is biased to teleport back to these seed nodes. Thus, the relevance scores from the initial candidates are effectively diffused throughout the graph. The parameter  $\alpha$  balances the exploration of the graph structure against the influence of the initial seed set.

The output of this stage is a globally-aware relevance score for every task node in the KG. This score reflects both the task’s direct

similarity and its interconnectedness within the broader network of related tasks. The enriched set of task scores forms the input for the final dataset aggregation and ranking stage.

### 6.3 Candidate Dataset Aggregation and Ranking

In the final stage, the system aggregates the task-level relevance scores to produce a final ranked list of datasets. This process involves two key steps: an initial score aggregation based on the graph structure, followed by a final re-ranking using an LLM to enhance precision.

To aggregate candidate datasets, the system identifies the set of all dataset nodes that are connected to the high-scoring task nodes in our main KG. To calculate a relevance score for each unique dataset, we employ an aggregation strategy. The score of a dataset is defined as the maximum relevance score among all the task nodes that are linked to it:

$$Score(d) = \max_{t \in \text{LinkedTasks}(d)} \{Score(t)\}$$

Based on these aggregated scores, all candidate datasets are then sorted in descending order to produce an initial ranked list.

To further enhance the precision of the top results, we introduce a final LLM-based re-ranking step. The system selects the top  $k_{rerank}$  datasets from the initial list. For each of these candidates, it retrieves their detailed descriptions from the knowledge base. This set of candidate descriptions, along with the original user query, is then formatted into a specialized instruction. The LLM is instructed to act as an expert evaluator and re-rank these few candidates based on their direct relevance to the user’s task. This re-ranking step utilizes the superior understanding and reasoning capabilities of the LLM to provide a highly refined and precise final output to the user.

## 7 EXPERIMENTAL EVALUATION

### 7.1 Experimental Setup

**7.1.1 Datasets and Benchmarks.** Our experiments are conducted on our benchmark suite CS-TDS. To ensure a fair and rigorous evaluation, we manually compare the ground truth with the output for each query to reduce false negatives. In addition to the original dataset, any other dataset that belongs to the same series (e.g., different scales of a dataset) and serves the same purpose was also marked as a correct answer.

**7.1.2 Baseline Systems.** To evaluate the effectiveness of KATS, we compare it against VanillaRAG [21] as a dense retrieval baseline, and GraphRAG [11], HippoRAG2 [20], and Raptor [33] (DIGIMON implementation [40]) as representative RAG frameworks. We adapt each method to return a ranked list of datasets under the same corpus, and evaluate all methods on CS-TDS under the same protocol.

**7.1.3 Evaluation Metrics.** We evaluate the performance of all systems from two perspectives, i.e., effectiveness and efficiency.

#### Effectiveness Metrics

- **Hit Rate@k:** The proportion of queries for which the correct dataset is found within the top- $k$  ranked results. We report for  $k$  in 1, 3, 5, 10.

- **Exact Match (EM):** For the top-1 result, this metric is 1 if the canonical name of the retrieved dataset exactly matches a ground-truth answer, and 0 otherwise.
- **F1-Score:** Also for the top-1 result, we compute the token-level F1-score between the predicted dataset name and the ground-truth name to account for partial matches.

#### Efficiency and Cost Metrics

- **Offline Build Time:** The total wall-clock time required to execute the entire offline knowledge base construction pipeline.
- **Storage Footprint:** The total disk space occupied by all generated artifacts, possibly including the KG, vector indexes, and any caches.
- **Build Token Cost:** The total number of input and output tokens consumed by all LLM API calls during the offline build phase.
- **Per-Query Time Cost:** The average wall-clock time for a single query during the online phase.
- **Per-Query Token Cost:** The average number of tokens consumed by LLM API calls for a query during the online phase.


**7.1.4 Implementation Details.** For all experiments, our KATS system was configured to use OpenAI’s text-embedding-3-small [29] as the embedding model. The LLM used for all extraction and judgment tasks was gpt-4o-mini [28]. Key hyperparameters for our KATS searcher are:

- **Initial Candidate Size ( $k$ ):** The number of initial candidate tasks retrieved from the FAISS index for each query is set to  $k = 2$ . This parameter controls the breadth of the initial semantic search.
- **PageRank Alpha ( $\alpha$ ):** The alpha parameter for the Personalized PageRank algorithm, which balances the influence of the initial query vector and the graph structure, is set to  $\alpha = 0.85$ .

All experiments were conducted on a server equipped with dual Intel Xeon Gold 6330 CPUs, an NVIDIA RTX A5000 GPU with 24GB of VRAM, and 1TB of RAM. The system was running on a Linux distribution with CUDA version 12.8.

### 7.2 End-to-End Effectiveness

In this section, we present the end-to-end effectiveness evaluation of our proposed KATS system against the baseline methods. Figure 4 visualizes the primary comparison in terms of Hit Rate@5. As the results clearly indicate, KATS consistently and significantly outperforms all baselines across both CS-TDS<sub>M</sub> and CS-TDS<sub>L</sub>.



**Figure 4: End-to-end effectiveness comparison in terms of Hit Rate@5 on CS-TDS<sub>M</sub> and CS-TDS<sub>L</sub>.**

**Table 4: Detailed Effectiveness Comparison on CS-TDS<sub>M</sub> and CS-TDS<sub>L</sub>. For each metric, the best result is in bold, and the second-best is underlined.**

System	Hit Rate@1		Hit Rate@3		Hit Rate@5		Hit Rate@10		EM (Top-1)		F1 (Top-1)	
	C <sub>M</sub>	C <sub>L</sub>	C <sub>M</sub>	C <sub>L</sub>	C <sub>M</sub>	C <sub>L</sub>	C <sub>M</sub>	C <sub>L</sub>	C <sub>M</sub>	C <sub>L</sub>	C <sub>M</sub>	C <sub>L</sub>
<b>KATS (ours)</b>	<b>0.447</b>	<b>0.446</b>	<b>0.638</b>	<b>0.672</b>	<b>0.723</b>	<b>0.730</b>	<b>0.787</b>	<b>0.804</b>	<u>0.208</u>	<b>0.206</b>	<b>0.350</b>	<b>0.312</b>
GraphRAG	<u>0.340</u>	<u>0.309</u>	<u>0.468</u>	0.475	<u>0.596</u>	0.534	0.617	0.601	0.064	0.034	<u>0.260</u>	0.192
HippoRAG 2	0.298	0.304	0.426	<u>0.515</u>	0.447	<u>0.608</u>	0.468	<u>0.662</u>	0.170	0.093	0.246	0.192
Raptor	<u>0.340</u>	0.284	0.404	0.431	0.511	0.520	<u>0.638</u>	0.583	<b>0.213</b>	<u>0.113</u>	0.256	<u>0.205</u>
VanillaRAG	0.255	0.221	0.404	0.358	0.426	0.422	0.468	0.471	0.149	0.064	0.206	0.149

On the larger dataset CS-TDS<sub>L</sub>, for instance, KATS achieves a Hit Rate@5 of 73.0%. This represents a substantial improvement of nearly 20% percentage points over the strongest baseline, HippoRAG 2 (60.8%). The performance gap is more significant when compared to the VanillaRAG baseline, which highlights the effectiveness of our KG-based hybrid search mechanism. We observe a similar trend on the CS-TDS<sub>M</sub> dataset, confirming the robustness of our approach across different corpus scales.


The significant performance advantage of KATS can be attributed to several key design choices. Firstly, unlike general-purpose frameworks, our pipeline constructs a highly specialized Task-Dataset KG, focusing only on the core entities and relationships relevant to our goal. Secondly, our novel LLM-based dataset entity resolution process effectively cleans the duplicated content in the graph by merging nodes with different names but identical semantics. This curation step is critical for consolidating knowledge and improving recall. Most importantly, our hybrid search mechanism performs a sophisticated task expansion using a PPR algorithm. This allows the system to understand user queries in a broader context and move beyond initial keyword or vector matches to discover datasets associated with a wider range of semantically related tasks. This directly contributes to its superior performance over more direct retrieval methods.

The superiority of KATS is not limited to the Hit Rate@5 metric alone. For a more comprehensive breakdown across all effectiveness metrics, Table 4 presents a detailed comparison. This demonstrates the ability of KATS to accurately identify the canonical name of the correct dataset. Even if not entirely precise, the search results are highly accurate in content. In summary, the combination of high recall and high precision provides strong evidence on multiple aspects of the end-to-end effectiveness of the KATS framework.

### 7.3 Efficiency and Cost Analysis

Beyond effectiveness, practical deployment depends on efficiency and operational cost. We evaluate KATS and baselines on CS-TDS<sub>M</sub> across time, LLM API tokens, and storage in Table 5, and summarize trade-offs between effectiveness and efficiency in Figure 5.

We first analyze time costs for offline construction and online retrieval, as shown in Table 5. For offline building, KATS takes 5.2 hours on CS-TDS<sub>M</sub>, slower than VanillaRAG (2.9h) but far faster than graph-based baselines such as Raptor (29.8h). This demonstrates that our specialized pipeline achieves a favorable balance between the richness of the constructed KG and the computational



**Figure 5: Effectiveness vs. efficiency trade-offs in KATS.**

**Table 5: Comparison of time, API, and storage costs during offline building (B.) and online retrieving (R.) on CS-TDS<sub>M</sub>.**

System	Time		API		Storage (MB)
	B. (h)	R. (s)	B. (M)	R.	
<b>KATS (ours)</b>	5.2	8.4	21	1,041	255
GraphRAG	99.4	124.3	151	20,011	8,228
HippoRAG 2	233.7	38.9	57	9,670	10,693
Raptor	29.8	10.0	30	6,067	1,017
VanillaRAG	2.9	27.1	0	1,723	418

cost of its construction. For online retrieval, KATS achieves 8.4s latency per query enabled by a pre-computed graph and a FAISS index. This ensures a responsive and interactive search experience.

In addition to time, LLM API token consumption is critical for deployment (Table 5). During building, KATS reduces token usage via our multi-agent extraction module and dataset resolution, which avoids deep processing over the entire corpus. For queries, the main cost comes from the final LLM reranking step, which only operates on a small candidate set and results in low per-query tokens.

Finally, we analyze the storage footprint. The results in Table 5 indicate that KATS maintains a moderate storage footprint. Unlike standard RAG, which embeds every chunk in the corpus, it embeds and indexes only the extracted relevant information. This strategy allows KATS to build a rich and targeted knowledge base while maintaining a compact storage footprint.

In summary, our efficiency and cost analysis demonstrates that KATS is a practical system. It achieves an effective balance of controlling reasonable one-time offline costs in terms of time, API calls, and storage to provide an efficient and cost-effective online query feature.

### 7.4 Incremental Update Efficiency

A crucial requirement for our system is the ability to efficiently incorporate new knowledge over time. To evaluate the efficiency and scalability of our incremental update pipeline, we conducted an experiment on the CS-TDS<sub>M</sub> corpus. First, an initial knowledge base was built using a base set of 518 documents. Subsequently, we simulated the arrival of new data by adding the remaining documents in batches of varying sizes, and measured the time required for each update operation.

As shown in Figure 6, there is a near-linear relationship between the number of new documents and the time required for the update process. The clear trend line indicates that the cost of an update is primarily dependent on the size of the incoming data. This confirms that our incremental approach is highly efficient and scalable, avoiding the prohibitive costs of performing a full system rebuild as the corpus grows.




Figure 6: The number of new documents vs. the time cost of the incremental update, along with the fitted trend line.

### 7.5 Ablation Studies

To understand the individual contribution of key components in our KATS system, we conduct a series of ablation studies. In these experiments, we systematically disable specific modules from our full pipeline and observe the impact on the end-to-end effectiveness. This allows us to quantify the value brought by each of our system design choices.

**7.5.1 LLM Rerank.** We first investigate the impact of the final LLM-based reranking step. The blue and orange folds in Figure 7 compare the performance of our full system with a variant where the reranking module is disabled. As illustrated, the inclusion of the LLM reranker brings a significant improvement in each level of hit rate. This demonstrates that while our graph-based ranking is effective at producing a strong set of candidates, the LLM’s reasoning capabilities are crucial for fine-tuning the precision of the top-ranked results.

**7.5.2 Dataset Resolution.** Next, we evaluate the contribution of our dataset resolution module, which is designed to combat entity ambiguity. To better demonstrate the effect of this module on the




Figure 7: Results of ablation experiments on CS-TDS<sub>M</sub>.

KG, instead of using LLM reranking, we directly evaluate the search results from the KG in our experiments. As shown in the orange and green folds in Figure 7, disabling this module has little effect on top-1 hit rates, but results in significant performance degradation of other hit rate metrics. This result strongly validates our hypothesis that resolving dataset naming ambiguities is a critical factor in real-world dataset search. By merging duplicate entities, our system is able to integrate fragmented information to form a more complete and accurate KG, thus directly improving recall and relevance.

### 7.6 Microbenchmarks

In addition to end-to-end performance, we conduct a series of microbenchmarks to analyze the robustness and sensitivity of our KATS system. Specifically, we investigate the system’s performance concerning its key hyperparameter and its dependency on LLM.




Figure 8: Sensitivity analysis of Hit Rate@3 with respect to the PageRank alpha (α) parameter on CS-TDS<sub>M</sub>.

**7.6.1 Hyperparameter Sensitivity.** We first analyze the system’s sensitivity to the choice of its key hyperparameter to ensure that its strong performance is not confined to a narrow setting. We focus on the PageRank alpha (α) that balances the influence of the initial semantic search and the subsequent graph traversal.

For this purpose, we fix all other parameters and run our system on the CS-TDS<sub>M</sub> corpus while varying the value of α:

$$\alpha \in \{0.75, 0.80, 0.85, 0.90, 0.95\}$$

The impact on the Hit Rate@3 metric is recorded.

As shown in Figure 8, the system maintains a high level of effectiveness across a reasonable range of  $\alpha$  values. A lower alpha value, which puts more emphasis on graph exploration, leads to a slight degradation. This suggests that the initial semantic signal from the seed tasks is highly valuable in the experiment. However, as the hyperparameter changes, the overall performance remains relatively stable. This result indicates that KATS is robust and not critically sensitive to the precise tuning of the hyperparameter  $\alpha$ .

**Table 6: Comparison of Hit Rate of KATS with different LLMs for building and QA on CS-TDS<sub>M</sub>.**

Building	QA	@1	@3	@5	@10
gpt-4o-mini	gpt-4o-mini	0.447	0.638	0.723	0.787
gemini-1.5-flash	gpt-4o-mini	0.426	0.690	0.746	0.787
gemini-1.5-flash	gemini-1.5-flash	0.404	0.532	0.617	0.702

**7.6.2 Robustness to LLMs.** To verify that the effectiveness of our framework is not dependent on a specific proprietary model, we evaluate its robustness using a different LLM. Specifically, for building and QA phases, we replaced the gpt-4o-mini [28] with a comparable model, gemini-1.5-flash [13], and reran the end-to-end evaluation on the CS-TDS<sub>M</sub> corpus.

Table 6 presents the results of this sensitivity analysis. Firstly, when only the LLM in the *building* phase is substituted, KATS’s end-to-end performance remains remarkably stable, with no significant degradation in Hit Rate metrics. This strongly suggests that our pipeline is not sensitive to the specific choice of the extraction model. Secondly, this stability indicates that the downstream components of our architecture, i.e., the hybrid entity resolution and graph-based ranking, are robust enough to handle the potential variations and minor errors introduced by different extraction models. While using a weaker model for the QA phase predictably impacts performance, the system still outperforms the baselines. This experiment validates that the core strength of KATS lies in its specialized architecture and strategy, rather than a dependency on any single proprietary model for its knowledge extraction process.

## 8 CASE STUDIES

### 8.1 Incremental Knowledge Discovery

To illustrate the practical utility of the incremental update feature, we present a case study on how KATS can dynamically incorporate new knowledge to answer previously unanswerable queries.

In the experiment, a specific paper from the CS-TDS<sub>M</sub> corpus that mentions the **CARLA Simulator Dataset** was deliberately withheld during the initial build process. Then, a query related to the dataset was issued as shown in Table 7. Since the relevant document was absent, the system was unable to find and recommend the CARLA Simulator Dataset. Afterwards, the withheld paper was introduced into the corpus through the incremental update process. The pipeline integrates the newly extracted information into the existing knowledge base. When the same query was issued, KATS successfully identified **CARLA** as a relevant result.

This real-world case study demonstrates that the incremental update mechanism is efficient and effective. It empowers KATS to

evolve by integrating new documents, thereby keeping its dataset recommendations current as new research emerges.

**Table 7: Comparison of results for the incremental knowledge discovery case study.**

<b>Query:</b>	<i>“Training models for urban driving scenarios using high-resolution visual inputs and expert actions.”</i>
<b>Before Update</b>	Unable to find the CARLA Simulator Dataset.
<b>After Update</b>	Recommend CARLA as the top-2 result.

### 8.2 Comparison with Online Research Agents

A relevant point of comparison for KATS is the emerging paradigm of online AI research agents, such as Gemini’s Deep Research feature. While both approaches aim to answer complex queries, they represent fundamentally different technical routes. We will illustrate these differences through the following case studies, as shown in Table 8. It is worth noting that this comparison inherently favors Deep Research, since the paper containing the ground-truth pair is hidden from KATS but accessible to the web-based agent.

**Table 8: Performance comparison of the online research agent case study.**

Task	Method	Latency	Result
Case 1	Gemini	392s	Ground truth within 6 answers
	KATS	< 20s	Ground truth ranks 5th
Case 2	Gemini	400s	Ground truth within 6-8 answers
	KATS	< 20s	Ground truth ranks 3rd

**Case 1: Physical Commonsense Reasoning Task.** Gemini Deep Research takes 392s and returns the ground-truth Kodak Dataset within its first 6 suggestions, while KATS responds in < 20s and ranks Kodak as the 5th result.

**Case 2: Image Compression Evaluation Task.** Gemini Deep Research takes 400s and returns the ground-truth PiQA within 6 to 8 suggestions across runs, while KATS responds in < 20s and consistently ranks PiQA as the 3rd result.

From the two cases above, it is evident that KATS and online research agent follow different paradigms. KATS follows an offline-KB paradigm: it preprocesses a corpus into a persistent knowledge base with low latency, cost efficiency, and reproducible results. In contrast, online research agents trade higher and more variable query-time latency for broader open-web coverage with no upfront build, but offer weaker privacy guarantees and less reproducibility.

### 8.3 Refined Task Disambiguation

Complex papers often evaluate multiple tasks on different datasets; if dataset mentions are linked to an overly broad task, the KG may contain spurious task-dataset associations. We test this setting with a paper from CS-TDS<sub>L</sub> that spans mathematical reasoning, natural language understanding, and code generation, evaluated on GSM8K [9], GLUE [36], and HumanEval [8], respectively. Table 9

shows that KATS extracts three independent tasks instead of collapsing all datasets under one task as a naive extraction would. Consequently, PPR retrieval for a query does not overscore datasets from the other two branches, which improves precision. This case study empirically validates that KATS’s architecture is capable of capturing context-specific relationships within a single document.

**Table 9: Results of the task disambiguation case study.**

Dataset	Extracted Task Description
GSM8K	<i>[...] multiple natural language understanding</i>
GLUE	<i>[...] multi-step mathematical word problems</i>
HumanEval	<i>[...] generate correct code for programming</i>

## 9 DISCUSSION

### 9.1 Findings and Implications

Our extensive experimental evaluation demonstrates that KATS, a system specifically designed for task-oriented dataset search from unstructured text, significantly outperforms state-of-the-art general-purpose RAG frameworks in both effectiveness and efficiency. This finding leads to several key implications for the future of dataset discovery.

*9.1.1 From Keywords to Intent: Bridging the Semantic Gap.* A primary challenge highlighted by prior work is the significant “semantic gap” between a user’s task-oriented need and the keyword-based queries that traditional systems require [17]. Our work demonstrates a successful paradigm that bridges this gap. By accepting open-ended natural language descriptions as input, KATS moves the cognitive burden of “query formulation” from the user to the system, which interprets complex task descriptions and maps them to relevant datasets. Future data discovery systems should focus less on optimizing keyword matching, but more on developing models that can analyze complex user intent.

*9.1.2 From General-purpose to Specialized RAG.* Our results also suggest that while general-purpose RAG architectures provide powerful and flexible frameworks, significant performance gains can be achieved by specialized systems for a specific problem. For task-oriented dataset search, this specialization involves not only retrieval but a deep integration of entity ambiguity and task-dataset mapping into the core pipeline. Developing specialized systems for high-value, vertical domains is a viable direction for RAG research.

### 9.2 Limitations

While our proposed KATS system demonstrates significant effectiveness for task-oriented dataset search, we also recognize several limitations to be addressed in the future.

*9.2.1 Domain Specificity of Knowledge Source.* Our current pipeline, particularly the multi-agent extraction framework, is tailored for the unique characteristics of scientific literature. A key premise of our approach is that these documents typically co-locate the mention of a dataset with a rich description of the specific task for which it was applied. However, the effectiveness of our extraction

strategy might vary when applied to other types of documents where this premise does not hold. Even though we have extended the system to integrate dataset descriptions without specific tasks, such as a dataset repository, the extraction process still requires a significant re-evaluation for other types of documents.

*9.2.2 Dependency on the Fidelity of LLM.* Like prior LLM-based knowledge extraction systems [11, 20, 33], KATS depends on the upstream LLM used in information extraction and entity resolution, and errors may propagate to the KG. To mitigate this risk, we decompose extraction into specialized agents and constrain the LLM’s role in entity resolution to binary verification over algorithmically generated candidate pairs. Table 6 shows stable end-to-end performance when substituting the building phase LLM, proving the effectiveness of these mitigation strategies. Nevertheless, the current system still lacks a specialized module for automatically verifying the authenticity of all LLM outputs and correcting possible errors based on the source documents, which remains a limitation.

## 10 CONCLUSION

Task-oriented search of datasets is a critical but unmet need for data professionals, encountering the challenges of a persistent semantic gap, lack of task-to-dataset mapping and benchmarks, and entity ambiguity. To address these challenges, our work makes several key contributions. Firstly, we developed the CS-TDS benchmark suite that addresses the critical gap in lacking evaluation benchmarks and enables the rigorous assessment of task-oriented dataset search systems. Moreover, we presented KATS, a novel end-to-end pipeline that automatically constructs a high-quality KG offline with a dedicated entity resolution mechanism, to resolve the challenge of lacking task-to-dataset mapping and entity ambiguity. Importantly, we equipped this pipeline with an efficient incremental update mechanism, transforming the static knowledge base into a dynamic and scalable one. To tackle the challenge of the persistent semantic gap, KATS is equipped with an online query processing pipeline that integrates vector search with graph-based ranking to generate highly relevant and task-consistent dataset recommendations based on the KG constructed before.

We believe that our work provides a solid foundation and a promising blueprint for the next generation of dataset discovery systems. Building on this foundation, there are several promising avenues for future work. A key direction is generalizing our framework beyond scientific literature to new problem domains, such as linking symptoms to diseases from medical texts. Further research could also focus on improving the technical capabilities, such as developing real-time streaming updates for the KG and incorporating automatic fact-checking modules to enhance its reliability.

## ACKNOWLEDGMENTS

Chen hao Ma was partially supported by NSFC under Grant 62302421, Basic and Applied Basic Research Fund in Guangdong Province under Grant 2025A1515010439, and the Guangdong Provincial Key Laboratory of Big Data Computing, The Chinese University of Hong Kong, Shenzhen. Xiaolin Han was supported by NSFC under Grant 62302397, the China Postdoctoral Science Foundation under Grant Number 2025M774364, and Shaanxi Post-doctoral Research Project (2025BSHSDZZ106).

## REFERENCES

- [1] Mubashara Akhtar, Omar Benjelloun, Costanza Conforti, Luca Foschini, Joan Giner-Miguel, Pieter Gijbbers, Sujata Goswami, Nitisha Jain, Michalis Karasousadakis, Michael Kuchnik, Satyapriya Krishna, Sylvain Lesage, Quentin Lhoest, Pierre Marcenac, Manil Maskey, Peter Mattson, Luis Oala, Hamidah Oderinwale, Pierre Ruyssen, Tim Santos, Rajat Shinde, Elena Simperl, Arjun Suresh, Geoffry Thomas, Slava Tykhonov, Joaquin Vanschoren, Susheel Varma, Jos van der Velde, Steffen Vogler, Carole-Jean Wu, and Luyao Zhang. 2024. Croissant: A Metadata Format for ML-Ready Datasets. *arXiv preprint arXiv:2403.19546* (2024). <https://arxiv.org/abs/2403.19546> Published at NeurIPS 2024 Datasets and Benchmark Track.
- [2] Qi An, Chihua Ying, Yuqing Zhu, Yihao Xu, Manwei Zhang, and Jianmin Wang. 2025. LEDD: Large Language Model-Empowered Data Discovery in Data Lakes. *arXiv preprint arXiv:2502.15182* (2025). <https://arxiv.org/abs/2502.15182>
- [3] Muhammad Imam Luthfi Balaka, David Alexander, Qiming Wang, Yue Gong, Adila Krisnadhi, and Raul Castro Fernandez. 2025. Pneuma: Leveraging LLMs for Tabular Data Representation and Retrieval in an End-to-End System. *Proc. ACM Manag. Data* 3, 3 (2025), 200:1–200:28. <https://doi.org/10.1145/3725337>
- [4] Alex Bogatu, Alvaro A. A. Fernandes, Norman W. Paton, and Nikolaos Konstantinou. 2020. Dataset Discovery in Data Lakes. In *Proceedings of the 36th IEEE International Conference on Data Engineering (ICDE)*. IEEE, 709–720. <https://doi.org/10.1109/ICDE48307.2020.00067>
- [5] Dan Brickley, Matthew Burgess, and Natasha Noy. 2019. Google Dataset Search: Building a search engine for datasets in an open Web ecosystem. In *Proceedings of the 2019 ACM Web Conference (WWW '19)*. ACM, New York, NY, USA, 1365–1375. <https://doi.org/10.1145/3308558.3313685>
- [6] Sonia Castelo, Rémi Rampin, Aécio Santos, Aline Bessa, Fernando Chirigati, and Juliana Freire. 2021. Auctus: A Dataset Search Engine for Data Discovery and Augmentation. In *Proceedings of the VLDB Endowment*, Vol. 14. VLDB Endowment, 2791–2794. <https://doi.org/10.14778/3476311.3476346>
- [7] Adriane P. Chapman, Elena Simperl, Laura Koesten, George Konstantinidis, Luis Daniel Ibáñez, Emilia Kacprzak, and Paul Groth. 2019. Dataset search: a survey. *The VLDB Journal* 29, 1 (2019), 251–272. <https://doi.org/10.1007/s00778-019-00564-x>
- [8] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Pondé de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgren Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Joshua Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021. Evaluating Large Language Models Trained on Code. *arXiv preprint arXiv:2107.03374* (2021). <https://arxiv.org/abs/2107.03374>
- [9] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training Verifiers to Solve Math Word Problems. *arXiv preprint arXiv:2110.14168* (2021). <https://arxiv.org/abs/2110.14168>
- [10] Matthijs Douze, Alexandr Krylov, Jeff Johnson, and Hervé Jégou. 2024. The Faiss library. *arXiv preprint arXiv:2401.08281* (2024). <https://arxiv.org/abs/2401.08281>
- [11] Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitan, Robert Osozuwa Ness, and Jonathan Larson. 2024. From Local to Global: A Graph RAG Approach to Query-Focused Summarization. *arXiv preprint arXiv:2404.16130* (2024).
- [12] Sainyam Galhotra, Yue Gong, and Raul Castro Fernandez. 2023. Metam: Goal-Oriented Data Discovery. In *39th IEEE International Conference on Data Engineering, ICDE 2023, Anaheim, CA, USA, April 3-7, 2023*. IEEE, 2780–2793. <https://doi.org/10.1109/ICDE55515.2023.00213>
- [13] Google. 2025. Gemini 1.5 Flash [Multimodal large language model]. <https://cloud.google.com/vertex-ai/generative-ai/docs/models/gemini/1-5-flash>.
- [14] Bernal Jiménez Gutiérrez, Yiheng Shu, Yu Gu, Michihiro Yasunaga, and Yu Su. 2024. HippoRAG: Neurobiologically Inspired Long-Term Memory for Large Language Models. In *Advances in Neural Information Processing Systems (NeurIPS)* 37. <https://arxiv.org/abs/2405.14831> NeurIPS 2024.
- [15] Xuming Hu, Shen Wang, Xiao Qin, Chuan Lei, Zhengyuan Shen, Asterios Katsifodimos, Christos Faloutsos, George Karypis, Lijie Wen, and Philip S. Yu. 2023. Automatic table union search with tabular representation learning. (2023). <https://www.amazon.science/publications/automatic-table-union-search-with-tabular-representation-learning>
- [16] Zezhou Huang, Jiaxiang Liu, Haonan Wang, and Eugene Wu. 2024. The Fast and the Private: Task-based Dataset Search. In *Conference on Innovative Data Systems Research (CIDR)*. <https://doi.org/10.48550/arXiv.2308.05637> arXiv:2308.05637 [cs.DB].
- [17] Madelon Hulsebos, Wenjing Lin, Shreya Shankar, and Aditya G. Parameswaran. 2024. “It Took Longer than I was Expecting.” Why is Dataset Search Still so Hard?. In *Proceedings of the 2024 Workshop on Human-In-the-Loop Data Analytics (HILDA '24)*. ACM, Santiago, Chile, 1–4. <https://doi.org/10.1145/3665939.3665959>
- [18] Daniel Jarrett, Alihan Hüyük, and Mihaela van der Schaar. 2022. Online Decision Mediation. In *Advances in Neural Information Processing Systems (NeurIPS)*, Vol. 35. 1790–1805. [http://papers.nips.cc/paper\\_files/paper/2022/hash/0bc795afae289ed465a65a3b4b1f4eb7-Abstract-Conference.html](http://papers.nips.cc/paper_files/paper/2022/hash/0bc795afae289ed465a65a3b4b1f4eb7-Abstract-Conference.html)
- [19] Glen Jeh and Jennifer Widom. 2003. Scaling personalized web search. In *Proceedings of the 12th International Conference on World Wide Web (Budapest, Hungary) (WWW '03)*. Association for Computing Machinery, New York, NY, USA, 271–279. <https://doi.org/10.1145/775152.775191>
- [20] Bernal Jiménez Gutiérrez, Yiheng Shu, Weijian Qi, Sizhe Zhou, and Yu Su. 2025. From RAG to Memory: Non-Parametric Continual Learning for Large Language Models. *arXiv preprint arXiv:2502.14802* (2025). <https://arxiv.org/abs/2502.14802>
- [21] Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen tau Yih. 2020. Dense Passage Retrieval for Open-Domain Question Answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, Online, 6769–6781. <https://doi.org/10.18653/v1/2020.emnlp-main.550>
- [22] Alex Krizhevsky. 2009. *Learning multiple layers of features from tiny images*. Technical Report. <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>
- [23] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. In *Advances in Neural Information Processing Systems 33 (NeurIPS 2020)*, Vol. 33. 9459–9474. <https://proceedings.neurips.cc/paper/2020/hash/6b493230205f780e1bc26945df7481e5-Abstract.html>
- [24] Hao Liu, Zhengren Wang, Xi Chen, Zhiyu Li, Feiyu Xiong, Qinhan Yu, and Wentao Zhang. 2025. HopRAG: Multi-Hop Reasoning for Logic-Aware Retrieval-Augmented Generation. *arXiv:2502.12442* [cs.IR] <https://arxiv.org/abs/2502.12442> arXiv preprint arXiv:2502.12442, version 2, last revised May 26, 2025.
- [25] LlamaIndex Team. 2025. LlamaParse. [https://github.com/run-llama/llama\\_cloud\\_services](https://github.com/run-llama/llama_cloud_services).
- [26] Fatemeh Nargesian, Erkang Zhu, Ken Q. Pu, and Renée J. Miller. 2018. Table Union Search on Open Data. *Proceedings of the VLDB Endowment* 11, 7 (2018), 813–825. <https://doi.org/10.14778/3192965.3192973>
- [27] Sebastian Neumaier, Jürgen Umbrich, and Axel Polleres. 2016. Automated Quality Assessment of Metadata across Open Data Portals. *ACM Journal of Data and Information Quality (JDIQ)* 8, 1 (Nov. 2016), 2. <https://doi.org/10.1145/2964909>
- [28] OpenAI. 2024. GPT-4o mini [Large language model]. <https://platform.openai.com/docs/models/gpt-4o-mini>.
- [29] OpenAI. 2024. text-embedding-3-small. <https://platform.openai.com/docs/guides/embeddings/embedding-models>.
- [30] Norman W. Paton, Jiaoyan Chen, and Zhenyu Wu. 2023. Dataset Discovery and Exploration: A Survey. *Comput. Surveys* 56, 4 (2023), 102:1–102:37. <https://doi.org/10.1145/3626521>
- [31] pypdf developers. [n.d.]. pypdf. <https://github.com/py-pdf/pypdf>.
- [32] Deepak Rajan, Michael J. Franklin, Stephan Hacker, Joseph M. Hellerstein, and Joseph E. Gonzalez. 2018. AURUM: A Data Discovery System. In *Proceedings of the 2018 International Conference on Management of Data (SIGMOD)*. ACM, 2109–2112. <https://doi.org/10.1145/3183713.3196927>
- [33] Parth Sarthi, Salman Abdullah, Aditi Tuli, Shubh Khanna, Anna Goldie, and Christopher D. Manning. 2024. RAPTOR: Recursive Abstractive Processing for Tree-Organized Retrieval. In *International Conference on Learning Representations (ICLR)*. <https://openreview.net/forum?id=GN921JHCRw>
- [34] Schema.org Community. 2011. Schema.org. <https://schema.org>.
- [35] Nandan Thakur, Nils Reimers, Andreas Rücklé, Abhishek Srivastava, and Iryna Gurevych. 2021. BEIR: A Heterogenous Benchmark for Zero-shot Evaluation of Information Retrieval Models. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*. <https://arxiv.org/abs/2104.08663>
- [36] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding, Vol. 3. 1786–1805. <https://openreview.net/forum?id=rJ4km2R5t7>
- [37] Zixin Wei, Jun Han, Xiaolin Han, and Chenhao Ma. 2025. SemExplorer: A User Interface for Semantic Approach to Customized Dataset Search. In *Companion of the 2025 International Conference on Management of Data (Berlin, Germany) (SIGMOD/PODS '25)*. Association for Computing Machinery, New York, NY, USA, 255–258. <https://doi.org/10.1145/3722212.3725133>
- [38] Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William W. Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. HotpotQA: A Dataset for Diverse, Explainable Multi-hop Question Answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for

Computational Linguistics, Brussels, Belgium, 2369–2380. <https://doi.org/10.18653/v1/D18-1259>

- [39] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu, Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2023. A Survey of Large

Language Models. <https://arxiv.org/abs/2303.18223>.

- [40] Yingli Zhou, Yaodong Su, Youran Sun, Shu Wang, Taotao Wang, Runyuan He, Yongwei Zhang, Sicong Liang, Xilin Liu, Yuchi Ma, et al. 2025. In-depth Analysis of Graph-based RAG in a Unified Framework. *arXiv preprint arXiv:2503.04338* (2025). <https://arxiv.org/abs/2503.04338>