

Data Agents: Levels, State of the Art, and Open Problems

Yuyu Luo HKUST (GZ) Guangzhou, China yuyuluo@hkust-gz.edu.cn	Guoliang Li Tsinghua University Beijing, China liguoliang@tsinghua.edu.cn	Ju Fan Renmin University of China Beijing, China fanj@ruc.edu.cn	Nan Tang HKUST (GZ) Guangzhou, China nantang@hkust-gz.edu.cn
---	--	---	---

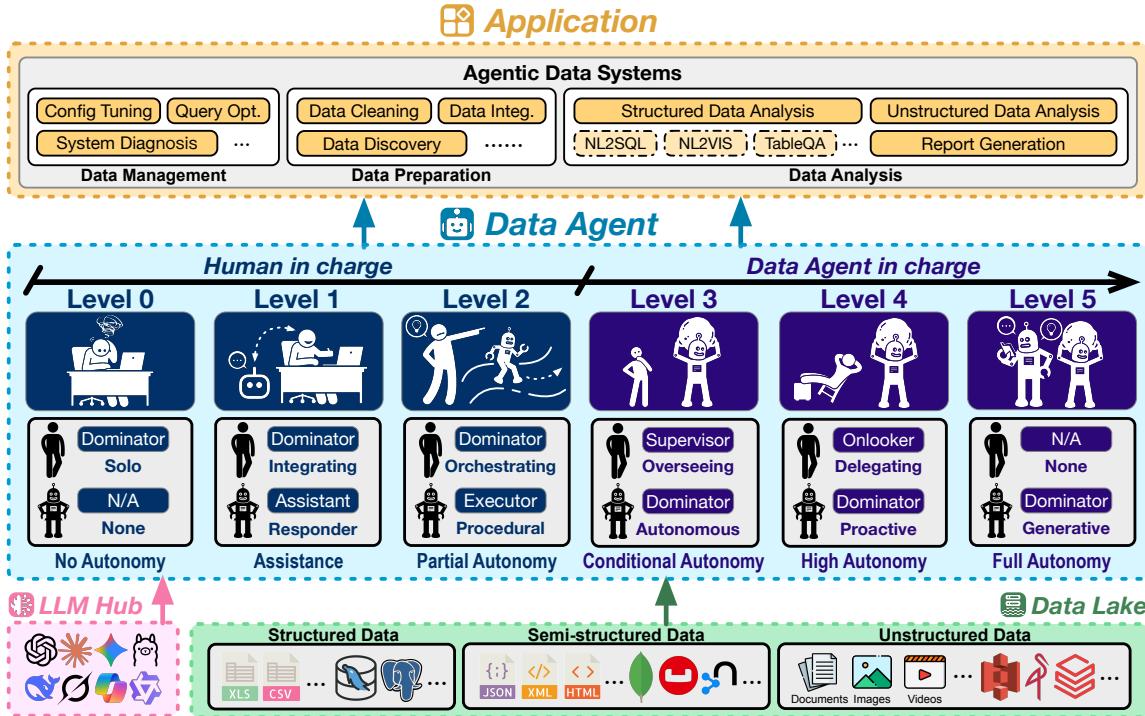


Figure 1: An Overview of Data Agents. (<https://github.com/HKUSTDial/awesome-data-agents>)

Abstract

Data agents are an emerging paradigm that leverages large language models (LLMs) and tool-using agents to automate data management, preparation, and analysis tasks. However, the term “data agent” is currently used inconsistently, conflating simple query responsive assistants with aspirational fully autonomous “data scientists”. This ambiguity blurs capability boundaries and accountability, making it difficult for users, system builders, and regulators to reason about what a “data agent” can and cannot do.

In this tutorial, we propose the first hierarchical taxonomy of data agents from Level 0 (L0, no autonomy) to Level 5 (L5, full autonomy). Building on this taxonomy, we will introduce a lifecycle- and level-driven view of data agents. We will (1) present the L0–L5

taxonomy and the key evolutionary leaps that separate simple assistants from truly autonomous data agents, (2) review representative L0–L2 systems across data management, preparation, and analysis, (3) highlight emerging Proto-L3 systems that strive to autonomously orchestrate end-to-end data workflows to tackle diverse and comprehensive data-related tasks under supervision, and (4) discuss forward-looking research challenges towards proactive (L4) and generative (L5) data agents. We aim to offer both a practical map of today’s systems and a research roadmap for the next decade of data-agent development.

1 Introduction

Modern data ecosystems are increasingly complex, spanning heterogeneous and multimodal data sources, evolving schemas, and tightly coupled Data+AI pipelines [31, 32, 86]. At the same time, LLM-based agents have demonstrated strong capabilities in tool use, planning, and iterative reasoning [37, 43, 67, 88]. As a result, the term *data agent* has rapidly gained popularity in both academia and industry [16, 61], with systems ranging from simple SQL or BI chatbots to ambitious products marketed as fully autonomous “AI data scientists”.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGMOD '26

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.
 ACM ISBN 978-x-xxxx-xxxx-x/YYY/MM
<https://doi.org/XXXXXX.XXXXXXX>

Without a shared vocabulary, however, fundamentally different systems are being conflated under a single, overloaded term. This leads to mismatched user expectations, ambiguous accountability when failures occur, and difficulty in objectively comparing different approaches. Similar challenges were previously faced by the driving-automation community, which motivated the SAE J3016 standard that introduced a six-level taxonomy of autonomy [53].

To address this confusion in data systems community, recent work proposes a hierarchical taxonomy of data agents [90], from Level 0 (L0, no autonomy) to Level 5 (L5, full autonomy), together with a structured survey of existing systems along this axis, which describes how task dominance and responsibility gradually shift from human operators to data agents as autonomy increases.

In this tutorial, we build on that survey and turn it into a *teaching-oriented* framework for SIGMOD attendees. Our goal is to help participants (1) understand what different levels of data agents can realistically do, (2) navigate the growing landscape of systems across the data lifecycle, and (3) identify key research challenges for advancing data agents towards higher autonomy.

1.1 Tutorial Overview

We will give a 3-hour tutorial consisting of a 140-minute lecture-style part (Parts I–IV) followed by a 40-minute *Data Agent Playground* (Part V) for hands-on exploration and discussion.

Part I: Problem Definition and Preliminaries (30 minutes). We begin by motivating data agents in modern Data+AI ecosystems and formalizing the notion of a data agent. We will: (i) introduce the motivation and problem definition of data agents, emphasizing why existing “data assistant” systems are insufficient and why autonomy and responsibility need to be made explicit; (ii) define data agents more formally and contrast them with general-purpose LLM agents along dimensions such as environment, data scale and structure, error propagation, and governance requirements, using a comparison table to highlight these differences; (iii) summarize key challenges (terminology ambiguity, lifecycle fragmentation, autonomy vs. governance, technical bottlenecks) and motivate the need for a level-based taxonomy of data agents.

Part II: L0–L2 Data Agents Across the Data Lifecycle (40 minutes). Next, we focus on the lower autonomy levels (L0–L2) and instantiate them in three phases of the data lifecycle: data management, data preparation, and data analysis. We will: (i) give an overview of how L0, L1, and L2 manifest in each phase and connect them to the roles of humans and agents illustrated in Figure 1; (ii) deep-dive into each phase: in data management, from manual DBAs (L0) to database tuning/diagnosis/query optimization copilots (L1) and L2 agents with direct access to DBMSs and monitoring signals; in data preparation, from scripts and rules (L0), to suggestion-style copilots to conduct data cleaning, integration, and discovery (L1), to L2 agents that invoke external tools and close the loop via execution feedback; in data analysis, from structured data analysis (Table QA / NL2SQL / NL2VIS), unstructured data analysis, and report generation with prompt-response paradigm (L1) to L2 environment-perceived analysis agents that maintain state and invoke SQL, plotting, and retrieval tools; (iii) use one or two running examples (e.g., database operations and BI analytics) to make the differences between L0, L1, and L2 concrete. We conclude

this part by summarizing recurring design patterns at L0–L2 and their reliability boundaries.

Part III: L3 Data Agents and Proto-L3 Systems (45 minutes). We then move to Level 3, the ongoing research frontier where data agents start to act as workflow orchestrators under human supervision. We will: (i) formally define L3 and explain the key evolutionary leap from L2 to L3; (ii) present representative Proto-L3 systems from academia that explore LLM orchestrators, semantic operators, task DAG optimization, and tool evolution to support versatile, cross-task workflows, and discuss their architectures, supported tasks, orchestration strategies, and limitations; (iii) analyze industrial “data agent” products in cloud data platforms and lakehouses, map them onto corresponding levels, and highlight common design patterns (e.g., DAG-based pipeline orchestration, planner–executor separation, multi-agent collaboration mechanism) and current bottlenecks (e.g., predefined operators/tools, limited causal/meta reasoning, constrained task coverage, strong reliance on human-crafted guardrails).

Part IV: Towards L4–L5 and Research Roadmap (25 minutes). Finally, we complete the lecture part by discussing the visionary Levels 4 and 5 and outlining a research roadmap. We will: (i) elaborate the vision of L4 data agents as proactive, long-lived, self-governing components that continuously monitor Data+AI ecosystems, autonomously discover issues and opportunities, and orchestrate pipelines without explicit instructions; (ii) introduce L5 data agents as generative data scientists that can invent new solutions, algorithms, and paradigms rather than only applying existing methods; (iii) summarize key open problems, including autonomous orchestration and versatility, causal and meta reasoning, intrinsic motivation and task discovery, long-horizon planning and trade-offs, safety and governance, and benchmarks for autonomy.

Part V: Data Agent Playground – Hands-on Exploration and Discussion (40 minutes). The final part is an interactive *Data Agent Playground* that increases audience engagement. We will walk through a few concrete data-agent workflows [11, 19, 55, 59, 80], show how L1/L2/Proto-L3 agents behave step by step, and invite attendees to try out our own data-agent prototypes. Participants will be encouraged to sketch or refine agents for their own settings, position them on the L0–L5 spectrum, and discuss key trade-offs in autonomy, governance, and reliability, followed by a brief Q&A that ties these insights back to the research roadmap in Part IV.

1.2 Our Scope and Goals

Our Distinction from Existing Tutorials. Existing tutorials and surveys on LLMs and data systems typically focus on specific aspects such as LLMs for databases and data analysis [31, 32, 38, 41, 66, 85], data management for machine learning [7, 14, 86], or general-purpose LLM agents and tool-using systems [59]. In contrast, our tutorial is distinguished by three aspects:

- (1) **Level-based view.** We adopt a *level-based* perspective on data agents (L0–L5) that explicitly links autonomy, capability, and responsibility, making it easier to reason about what a “data agent” at each level can and cannot do.
- (2) **Holistic lifecycle perspective.** We take a *holistic lifecycle* view, jointly covering data management, data preparation, and

Table 1: Comparison between General LLM Agents and Data Agents

Aspect	General LLM Agents	Data Agents
Primary Focus	Task and Content Centric: <i>Completing defined tasks or generating content.</i>	Data-Lifecycle Centric: <i>Data management, preparation, and analysis.</i>
Problem Scope	Self-contained and Static: <i>Acts on explicit instructions and a finite prompt.</i>	Exploratory and Dynamic: <i>Actively explores and navigates vast, dynamic data lakes.</i>
Input Data	Small-Scale and Ready-to-Use: <i>Typically receives manageable, clean inputs.</i>	Large-Scale and “Raw”: <i>Designed to handle heterogeneous, dynamic, and noisy raw data.</i>
Tool Invocation	General-Purpose Toolkit: <i>Web search, calculators, OCR, image generators, etc.</i>	Specialized Data Toolkit: <i>DB loaders, SQL equivalence checker, visualization libraries, etc.</i>
Primary Output	Generative Artifacts: <i>Human-consumable product: dialogues, reasoning, images, etc.</i>	Data Products and Insights: <i>Config, processed data, insights, visualizations, analytical report, etc.</i>
Error Consequence	Localized: <i>Typically affects limited to only the direct output.</i>	Cascading: <i>Errors can cascade, affecting downstream insights.</i>

data analysis under a unified data-agent framework, rather than treating individual tasks in isolation.

(3) **Evolutionary leaps and roadmap.** We emphasize the *evolutionary leaps* between levels, especially the crucial L2→L3 and L3→L4 transitions, and present a *research roadmap* towards proactive (L4) and generative (L5) data agents, instead of providing an exhaustive but flat catalogue of systems.

Target Audience and Learning Outcomes. This tutorial is intended for a broad SIGMOD audience, including researchers in databases, data mining, machine learning, AI agents, and data-centric AI; system developers and practitioners building data platforms, lakehouses, or enterprise data stacks; and students who wish to enter the emerging area of data agents. By the end of the tutorial, participants will be able to use the L0–L5 framework to position existing and future systems, distinguish data agents from general-purpose LLM agents, interpret and calibrate vendor claims about “data agents”, choose appropriate autonomy levels for their own applications, and reason about key design dimensions such as perception, planning, tools, memory, and governance. We assume familiarity with basic database concepts and LLM terminology; the tutorial itself will be self-contained.

2 Tutorial Outline

We first define what we mean by data agents and situate them in the broader landscape of data systems and LLM agents.

2.1 Background and Problem Definition

2.1.1 Problem Description: What is a Data Agent? Informally, a *data agent* is an LLM-based architecture that orchestrates a Data+AI ecosystem to perform data-related tasks such as configuration tuning, data cleaning, integration, exploration, and analysis [16, 61, 85]. Formally, we can define a data agent \mathcal{A} that operates on raw data \mathcal{D} within an environment \mathcal{E} (e.g., DBMS, code interpreters, APIs, etc.), utilizing LLMs \mathcal{M} , ultimately producing an output \mathcal{O} to tackle the data-related task \mathcal{T} , abstractly represented as:

$$\mathcal{A} : (\mathcal{T}, \mathcal{D}, \mathcal{E}, \mathcal{M}) \rightarrow \mathcal{O}.$$

This broad formulation captures a spectrum of systems, from simple assistants that suggest SQL queries to aspirational “AI data scientists” that autonomously manage and analyze data.

2.1.2 Task Landscape and Data Agents vs. General LLM Agents. Data agents operate within modern Data+AI ecosystems that span

relational databases, data warehouses and lakehouses, data lakes, ETL/ELT pipelines, BI tools, and ML services. Therefore, data agents must reason over large, heterogeneous, and often schema-rich data lakes without exhaustive ingestion [32]; interact with dynamic and noisy data and systems whose behavior changes over time [31]; and operate inside multi-stage pipelines where errors can silently propagate and amplify, rather than affecting only a single response.

Compared to general-purpose LLM agents, data agents thus face more constrained yet substantially more demanding environments. They also need to satisfy stringent requirements on reliability, governance, and reproducibility that are less prominent in many generic agent settings. Table 1 summarizes key differences between data agents and general LLM agents along these dimensions.

2.1.3 Key Challenges. These characteristics give rise to several challenges that motivate the need for a principled taxonomy for data agents:

- **Ambiguous terminology and overstated claims.** Without a shared vocabulary, systems with very different autonomy levels are all marketed as “data agents”, leading to hype, confusion, and misaligned expectations [16, 61].
- **Fragmentation across the data lifecycle.** Data agents must span data management, data preparation, and data analysis over heterogeneous, multi-modal data lakes [66, 85], yet most existing works focus on individual tasks or stages in isolation, making it difficult to reason about end-to-end capabilities and trade-offs.
- **Autonomy vs. governance.** As autonomy increases, assigning responsibility, defining safe operating regions, and providing guarantees become both more important and more challenging, especially when data agents can autonomously modify data, configurations, or pipelines.
- **Technical bottlenecks.** Advancing to higher autonomy levels requires progress in perception (over large, complex data and systems), long-horizon planning and orchestration, memory and continual adaptation, causal and meta reasoning, and robust interaction with dynamic environments.

To bring clarity, we adopt a level-based framework for data agents that explicitly links autonomy, capability, and responsibility.

2.2 The L0–L5 Hierarchy of Data Agents

Inspired by the SAE J3016 standard for driving automation [53], we adopt a six-level taxonomy of data agents from L0 to L5. As summarized in Figure 1, data agents are organized into six autonomy

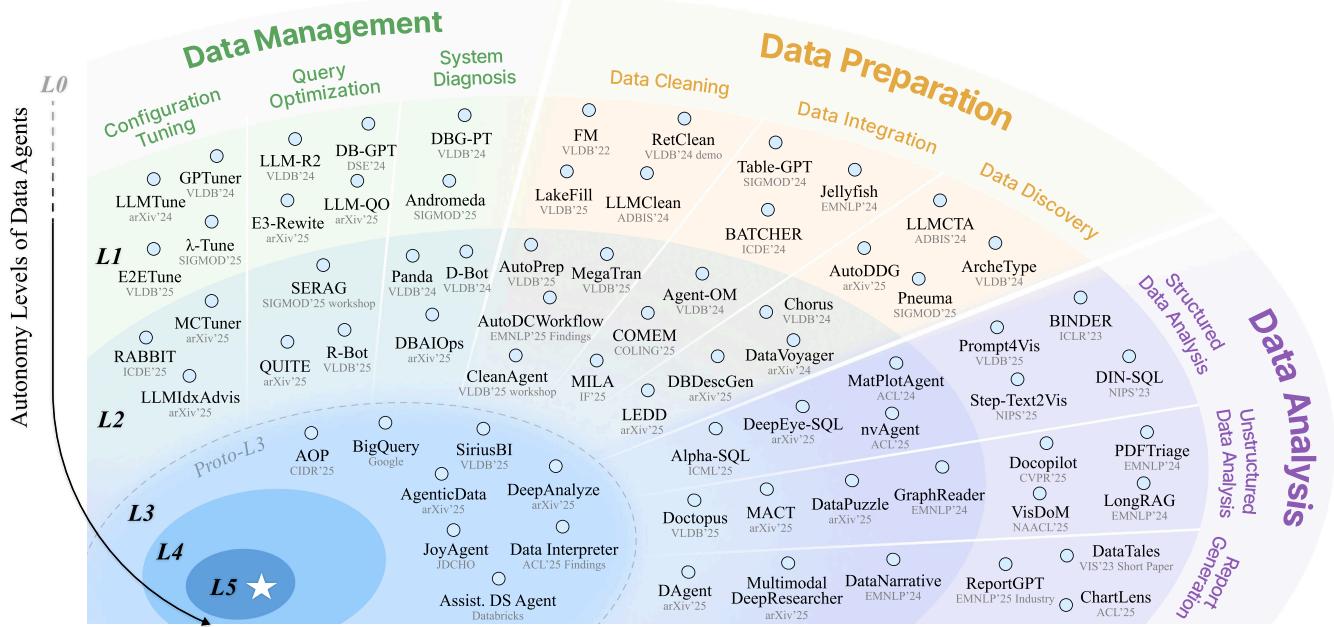


Figure 2: Representative Data Agents Across Different Levels.

levels, from Level 0 (L0) to Level 5 (L5). The figure indicates, for each level, who is in charge of the data-related task (human vs. data agent), what role data agent plays (e.g., responder, executor, orchestrator, proactive or generative component), and which parts of the data lifecycle (management, preparation, analysis) are involved. We briefly review these levels below.

2.2.1 Levels of Autonomy.

L0: No Autonomy. At L0, there is no data agent involvement. All tasks in data management, preparation, and analysis are performed manually by humans.

L1: Assistance. L1 data agents operate within a stateless, prompt-response framework. They can answer questions, generate code snippets, or suggest queries, but they do not perceive or interact with the environment. Humans remain fully responsible for executing and verifying any suggestions.

L2: Partial Autonomy. L2 data agents gain the ability to perceive and interact with their environment, including data lakes, DBMSs, code interpreters, and external APIs. They may possess memory and can invoke tools to autonomously execute task-specific procedures within human-orchestrated pipelines.

L3: Conditional Autonomy. L3 data agents are expected to autonomously orchestrate and execute tailored data pipelines for a wide range of tasks under human supervision. They interpret high-level user intentions and dominate the end-to-end workflow, while humans act as supervisors.

L4: High Autonomy. L4 data agents achieve high autonomy and reliability, eliminating the need for human supervision and explicit instructions. They are fully delegated to proactively monitor Data+AI ecosystems, autonomously discover issues and opportunities in data lakes, and orchestrate pipelines to address them.

L5: Full Autonomy. At L5, data agents are envisioned to innovate new solutions and paradigms beyond existing methods, acting as

fully autonomous and generative data scientists. Human involvement becomes unnecessary.

As an overview, Figure 2 positions representative systems from academia and industry across the L0–L5 levels and the three phases of the data lifecycle.

2.3 L0–L2: From Manual Workflows to Partial Autonomy

In this section, we review representative systems at L0–L2 across three phases of the data lifecycle: data management, data preparation, and data analysis.

2.3.1 Data Management. Data management includes configuration tuning, query optimization, and system diagnosis in database systems [83, 86]. At **L0**, DBAs manually tune knobs, index configurations, and execution plans, relying on expertise and trial-and-error [83]. At **L1**, LLMs are used as query-responsive assistants to generate tuning suggestions or rewritten queries. They operate in a prompt-response manner, returning recommendations that humans must integrate and validate [17, 26, 36]. For instance, λ -Tune [18] and E2ETune [21] use LLMs to recommend configuration candidates based on workload features, and Andromeda [8] generates diagnostic suggestions for configuration debugging. At **L2**, data agents gain direct access to the DBMS and monitoring information. They can observe workload statistics, execute tuning experiments, and adjust configurations or rewrite queries in a decision loop, while still following human-designed workflows [56, 73, 84]. Rabbit [60], R-Bot [62], D-Bot [87] exemplify this through utilizing environmental feedback in configuration tuning, query rewriting, and system diagnosis.

2.3.2 Data Preparation. Data preparation [14] covers data cleaning [4], integration [33], and discovery [15]. At **L1**, data agents primarily act as suggestion engines: RetClean [44] and LakeFill [74]

infer and impute missing values, LLMClean [4] generate rules for cleaning tasks, Narayan et al. [45] deploy LLMs to propose schema matches or entity correspondences, AutoDDG [79] and LLMCTA [25] produce dataset summaries, metadata, or column annotations. At **L2**, data agents go beyond query responder and directly interact with databases or data lakes to execute cleaning and transformation operations, verify constraints, and adjust their strategies based on execution feedback, and iteratively refine integration decisions as more data is explored [24, 50, 81]. Representative systems include CleanAgent [50], MegaTran [30] for data cleaning; SEED [9], Agent-OM [51] for data integration; LEDD [3] and DBDescGen [35] for data discovery.

2.3.3 Data Analysis. Data analysis includes structured and unstructured data analysis, as well as report generation. At **L1**, we mostly see LLM-driven question-answering assistants for Table QA [10, 57, 77], NL2SQL [28, 39, 49, 89], NL2VIS [34, 40, 42], textual or multimodal Document QA [13, 52, 63], which generate answers to response user questions over curated datasets, and report generators [6, 58, 64] that operate on input tables or documents. At **L2**, data agents move beyond static querying to dynamically engage with, verify, and refine multi-step analytical processes [48, 54, 68, 78]. They invoke tools such as SQL engines [27, 29], plotting libraries [46, 72, 76], or retrieval modules [23, 71, 82], and support iterative exploration and refinement of analyses [12, 47, 75].

2.4 L3: Striving for Autonomous Data Agents

We now turn to Level 3 (L3), which marks a crucial step from procedural executors to autonomous orchestrators.

2.4.1 From Executor to Dominator. At L2, humans design the overall pipelines, and data agents execute specific procedures within these human-prescribed workflows. At L3, by contrast, data agents are expected to interpret high-level user intent and autonomously orchestrate pipelines that span data management, preparation, and analysis. During execution, data agents adapt the pipeline based on feedback and intermediate results, while humans primarily act as supervisors who review plans and outcomes rather than as pipeline designers. In this sense, task dominance and primary responsibility shift from humans to data agents. Figure 3 illustrates the typical L3 data agent, highlighting its conditional autonomy in autonomous pipeline orchestration and optimization.

2.4.2 Proto-L3 Data Agents in Research. Recent research systems begin to exhibit partial L3 capabilities. They use LLM-based orchestrators [59], predefined operators [69, 70], workflow optimization [20, 69], and tool libraries [80] to orchestrate multi-step workflows over heterogeneous systems, cover multiple stages of the data lifecycle within a single agentic process, and maintain state across long-running interactions so that they can refine their plans and correct mistakes over time. These Proto-L3 agents typically operate in constrained environments with curated tools and data, but they provide concrete testbeds for studying the transition from execution-focused L2 agents to orchestration-centered L3 systems.

We will present several representative academic systems and discuss: (i) their pipeline representation, orchestration, and optimization strategies; (ii) their architectural choices (single vs. multi-agent, central vs. decentralized planners); (iii) their approach to

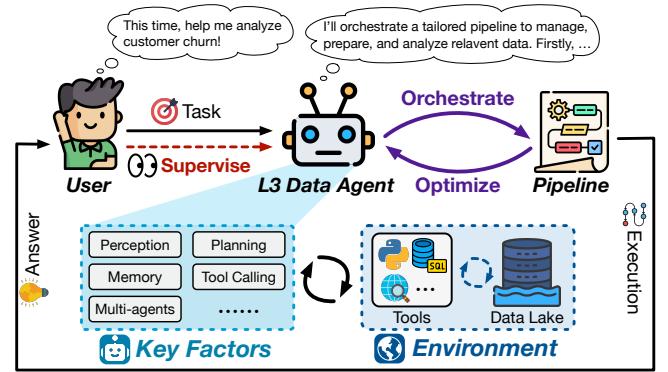


Figure 3: L3 Data Agents (Conditional Autonomy).

tool abstraction and composition; and (iv) their strategies for incorporating feedback and handling errors.

Table 2 compares representative Proto-L3 data agents from both academia and industrial products along dimensions such as tool flexibility, data complexity, data lifecycle coverage, and specific management, preparation, and analysis tasks they support.

2.4.3 Industrial Data-Agent Products. Industrial platforms (e.g., cloud data warehouses and lakehouses) have started to offer commercial “data agent” products [19, 55]. We analyze: (i) how these products map to the L0–L3 levels in practice; (ii) which guarantees they provide (e.g., human-in-the-loop confirmation, logging, and rollback); and (iii) common limitations and design trade-offs.

2.4.4 Current Bottlenecks and Gaps. The survey identifies several gaps preventing current systems from realizing full L3 autonomy:

- limited pipeline orchestration capabilities and reliance on pre-defined operators;
- inadequate higher-order, causal, and meta-reasoning to diagnose and avoid cascading errors;
- difficulty adapting to dynamic environments with changing data and workloads;
- heavy reliance on human-crafted reinforcement learning setups for alignment and adoption.

These challenges motivate the need for new methods that go beyond straightforward tool-calling LLM agents.

2.5 L4–L5: Vision and Research Roadmap

Finally, we discuss the visionary Levels 4 and 5 and outline a research roadmap.

2.5.1 L4: Proactive, High-Autonomy Data Agents. At L4, data agents are envisioned as proactive, long-lived, and self-governing components of Data+AI ecosystems. Instead of merely reacting to explicit user requests, an L4 agent continuously monitors data lakes, systems, and models, detects phenomena such as data drift, performance regressions, and schema changes, and identifies opportunities such as beneficial materializations, missing indexes, or promising analytical workflows. It is expected to prioritize among these tasks, design and adapt pipelines to address them without explicit instructions, and operate within reliability, safety, and governance constraints even in the absence of human supervision. Typical scenarios include autonomous detection and mitigation of

Table 2: Comparison of Representative Proto-L3 Data Agents from Academia Research and Industry Products. Compares Open-source: availability; Undef Ops.: capabilities in utilizing unpredefined operators; data-related task coverage across data management, preparation, analysis; data complexity dimensions: Multi-source (Multis.), Heterogeneous (Hete.), and Multimodal (Multim.).

Years	Data Agent	Open-source	Undef Ops.	Data Complexity			Data Management			Data Preparation			Data Analysis		
				Multis.	Hete.	Multim.	Config Tun.	Query Opt.	Sys. Diag.	Data Clean.	Data Integ.	Data Disc.	Struct.	Unstruct.	Report Gen.
2025	AgenticData [59]	-	✗	✓	✓	-	-	✓	✓	✓	✓	✓	✓	✓	-
2025	DeepAnalyze [80]	✓	-	✓	✓	-	-	-	-	✓	✓	✓	✓	✓	✓
2025	AOP [69]	-	-	✓	✓	✓	-	✓	✓	✓	✓	✓	✓	✓	-
2025	iDataLake [70]	✓	-	✓	✓	✓	-	✓	-	✓	✓	✓	✓	✓	✓
2024	Data Interpreter [20]	✓	-	-	✓	✓	-	-	-	✓	-	✓	✓	✓	✓
2025	JoyAgent [2]	✗	✗	✓	✓	-	-	-	✓	✓	✓	✓	✓	✓	✓
2025	Assist. DS Agent [11]	-	-	✓	✓	-	-	✓	-	✓	✓	✓	✓	✓	✓
2025	TabTab [65]	-	-	✓	✓	-	-	-	-	✓	✓	-	✓	✓	✓
2025	ByteDance Data Agent [5]	-	-	✓	✓	-	-	-	-	✓	-	✓	✓	✓	✓
2025	BigQuery [19]	-	-	✓	✓	-	-	✓	-	✓	✓	✓	✓	-	-
2025	Cortex [55]	-	-	✓	✓	✓	-	-	-	✓	✓	✓	✓	✓	-
2025	Xata Agent [1]	-	-	✓	✓	-	✓	✓	✓	-	-	✓	-	-	-
2025	SiriusBI [22]	-	-	-	-	-	-	-	-	✓	-	✓	✓	-	✓

workload shifts, long-horizon management of indexes and materialized views, and continuous quality assurance for critical data assets. Realizing such capabilities not only raises questions about autonomous orchestration across the full data lifecycle but also calls for mechanisms for intrinsic motivation, task discovery in large data ecosystems, and long-horizon planning that reasons about cumulative cost, latency, and data-quality trade-offs.

2.5.2 L5: Generative Data Agents. L5 data agents go beyond deploying existing techniques and are conceived as autonomous, generative data scientists. An L5 data agent is expected to identify gaps in current methods, hypothesize new algorithms or representations when existing approaches are insufficient, design and analyze experiments to test these hypotheses, and iteratively refine its own solutions over time. In this vision, the data agent is not only a user of database and ML systems, but also an active contributor to their evolution. Moving towards L5 requires abstractions that allow data agents to manipulate high-level design choices—such as physical designs, query rewrite strategies, data cleaning policies, or learning procedures—while staying grounded in executable systems, as well as causal and meta reasoning supporting principled diagnosis, comparison, and improvement of alternative designs, even pioneering of innovative solutions, novel theories, and new paradigm.

Although fully realized L4 and L5 data agents remain speculative, articulating these levels helps delineate a research roadmap. In the near term, the most pressing challenges lie in making L2 and Proto-L3 agents more robust, transparent, and governable; in the medium term, progress toward L4 will depend on advances in autonomous orchestration, task discovery, and long-horizon decision making under multi-objective constraints; and in the longer term, movement toward L5 will hinge on integrating causal and meta reasoning with agent-driven experimentation and on developing evaluation methodologies that capture autonomy, adaptability, and safety beyond traditional task-level accuracy.

2.5.3 Research Opportunities. The L0–L5 hierarchy suggests several research directions that are closely tied to core data management problems. A central question is how data agents should perceive and act over large, heterogeneous data lakes: which indexes,

materialized views, summaries, or learned representations should serve as their “senses”, how these structures are exposed as tools, and how agents can orchestrate complex pipelines across management, preparation, and analysis while preserving performance, data quality, and governance guarantees.

A second theme concerns how data agents are trained and evaluated in realistic environments. Here, operational logs, configuration histories, and telemetry can form the basis for constructing training corpora, adapting agent policies over time, and supporting causal and meta reasoning about failures and improvements. This, in turn, calls for benchmarks and methodologies that go beyond task-level accuracy to capture autonomy, robustness, adaptability, and safety on realistic data-management workloads.

3 BIOGRAPHY

Yuyu Luo is an Assistant Professor at The Hong Kong University of Science and Technology (Guangzhou), with an affiliated position at the HKUST. He received his PhD from Tsinghua University in 2023. His research interests include Data Agents, AI4DB, and Data-centric AI. He has received the Best-of-SIGMOD 2023 Papers.

Guoliang Li is a full professor in the Department of Computer Science, Tsinghua University. His research interests mainly include data cleaning and integration, and machine learning for databases. He got the VLDB 2017 early research contribution award, TCDE 2014 Early Career Award, VLDB 2023 Industry Best Paper Runner-up, Best of SIGMOD 2023, SIGMOD 2023 research highlight award, DASFAA 2023 Best Paper Award, and CIKM 2017 Best Paper Award.

Ju Fan is a Professor at the DEKE Lab, MOE China, and the School of Information, Renmin University of China. He received his PhD from Tsinghua University in 2013 and received the ACM China Rising Star Award and the 2023 SIGMOD Research Highlight Award. Dr. Fan’s main research interests are AI4DB and database systems.

Nan Tang is an Associate Professor at The Hong Kong University of Science and Technology (Guangzhou), with an affiliated position at the HKUST. He has received the VLDB 2010 Best Paper Award, the 2023 SIGMOD Research Highlight Award, and the Best-of-SIGMOD 2023. His main research interests are AI4DB and data-centric AI.

References

- [1] [n. d.]. *Xata Agent*.
- [2] Agent Team at JDCHO. [n. d.]. *JoyAgent-JDGenie*.
- [3] Qi An, Chihua Ying, Yuting 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).
- [4] Fabian Biester, Mohamed Abdelaal, and Daniel Del Gaudio. 2024. LLM-Clean: Context-Aware Tabular Data Cleaning via LLM-Generated OFDs. *arXiv:2404.18681* [cs.DB].
- [5] ByteDance Volcengine. [n. d.]. *Data Agent*.
- [6] Lucas Cecchi and Petr Babkin. [n. d.]. ReportGPT: Human-in-the-loop Verifiable Table-to-Text Generation. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track*, Franck Dernoncourt, Daniel Preoțiuc-Pietro, and Anastasia Shmorina (Eds.).
- [7] Chengliang Chai, Jiayi Wang, Yuyu Luo, Zeping Niu, and Guoliang Li. 2023. Data Management for Machine Learning: A Survey. *IEEE Transactions on Knowledge and Data Engineering* (2023).
- [8] Sibei Chen, Ju Fan, Bin Wu, Nan Tang, Chao Deng, Pengyi Wang, Ye Li, Jian Tan, Feifei Li, Jingren Zhou, et al. 2025. Automatic database configuration debugging using retrieval-augmented language models. *Proceedings of the ACM on Management of Data* (2025).
- [9] Zui Chen, Lei Cao, Sam Madden, Tim Kraska, Zeyuan Shang, Ju Fan, Nan Tang, Zihui Gu, Chunwei Liu, and Michael Cafarella. 2023. SEED: Domain-specific data curation with large language models. *arXiv preprint arXiv:2310.00749* (2023).
- [10] Zhoujun Cheng, Tianbao Xie, Peng Shi, Chengzu Li, Rahul Nadkarni, Yushi Hu, Caiming Xiong, Dragomir Radev, Mari Ostendorf, Luke Zettlemoyer, Noah A. Smith, and Tao Yu. 2023. Binding Language Models in Symbolic Languages. *arXiv:2210.02875* [cs].
- [11] Databricks. [n. d.]. *Assistant Data Science Agent*.
- [12] Minghang Deng, Ashwin Ramachandran, Canwen Xu, Lanxiang Hu, Zhewei Yao, Anupam Datta, and Hao Zhang. [n. d.]. ReFoRCE: A Text-to-SQL Agent with Self-Refinement, Format Restriction, and Column Exploration. *arXiv preprint arXiv:2502.00675* [n. d.].
- [13] Yuchen Duan, Zhe Chen, Yusong Hu, Weiyun Wang, Shenglong Ye, Botian Shi, Lewei Lu, Qibin Hou, Tong Lu, Hongsheng Li, Jifeng Dai, and Wenhui Wang. 2025. Docopilot: Improving Multimodal Models for Document-Level Understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- [14] Alvaro AA Fernandes, Martin Koehler, and et al. 2023. Data preparation: A technological perspective and review. *SN Computer Science* (2023).
- [15] Benjamin Feuer, Yurong Liu, Chinmay Hegde, and Juliana Freire. 2024. ArcheType: A Novel Framework for Open-Source Column Type Annotation Using Large Language Models. *Proceedings of the VLDB Endowment* (2024).
- [16] Yanjie Fu, Dongjie Wang, Wangyang Ying, Xiangliang Zhang, Huan Liu, and Jian Pei. 2025. Autonomous Data Agents: A New Opportunity for Smart Data. *arXiv preprint arXiv:2509.18710* (2025).
- [17] Victor Giannakouris and Immanuel Trummer. 2024. Dbg-pt: A large language model assisted query performance regression debugger. *Proceedings of the VLDB Endowment* (2024).
- [18] Victor Giannakouris and Immanuel Trummer. 2025. λ -tune: Harnessing large language models for automated database system tuning. *Proceedings of the ACM on Management of Data* (2025).
- [19] Google Cloud. [n. d.]. *BigQuery*.
- [20] Sirui Hong, Yizhang Lin, and et al. Bang Liu. 2025. Data Interpreter: An LLM Agent for Data Science. In *Findings of the Association for Computational Linguistics*.
- [21] Xinmei Huang, Haoyang Li, Jing Zhang, Xinxin Zhao, Zhiming Yao, Yiyi Li, Tieying Zhang, Jianjun Chen, Hong Chen, and Cuiping Li. 2025. E2Etune: End-to-end knob tuning via fine-tuned generative language model. *Proceedings of the VLDB Endowment* (2025).
- [22] Jie Jiang, Haining Xie, Siqi Shen, Yu Shen, and et al. 2025. SiriusBI: A Comprehensive LLM-powered Solution for Data Analytics in Business Intelligence. *Proc. VLDB Endow.* (2025).
- [23] Zhengbao Jiang, Frank Xu, Luyu Gao, Zhiqing Sun, Qian Liu, Jane Dwivedi-Yu, Yiming Yang, Jamie Callan, and Graham Neubig. [n. d.]. Active Retrieval Augmented Generation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*.
- [24] Moe Kayali, Anton Lykov, Ilias Fountalis, Nikolaos Vasiloglou, Dan Olteanu, and Dan Suciu. 2023. CHORUS: foundation models for unified data discovery and exploration. *arXiv preprint arXiv:2306.09610* (2023).
- [25] Keti Korini and Christian Bizer. 2025. Evaluating knowledge generation and self-refinement strategies for llm-based column type annotation. *arXiv preprint arXiv:2503.02718* (2025).
- [26] Jiale Lao, Yibo Wang, Yufei Li, Jianping Wang, Yunjia Zhang, Zhiyuan Cheng, Wanghu Chen, Mingjie Tang, and Jianguo Wang. 2024. GPTuner: A Manual-Reading Database Tuning System via GPT-Guided Bayesian Optimization. *Proc. VLDB Endow.* (2024).
- [27] Boyan Li, Chong Chen, Zhujun Xue, Yinan Mei, and Yuyu Luo. 2025. DeepEye-SQL: A Software-Engineering-Inspired Text-to-SQL Framework. *arXiv preprint arXiv:2510.17586* (2025).
- [28] Boyan Li, Yuyu Luo, Chengliang Chai, Guoliang Li, and Nan Tang. 2024. The Dawn of Natural Language to SQL: Are We Fully Ready? *Proc. VLDB Endow.* (2024).
- [29] Boyan Li, Jiayi Zhang, Ju Fan, Yanwei Xu, Chong Chen, Nan Tang, and Yuyu Luo. 2025. Alpha-SQL: Zero-Shot Text-to-SQL using Monte Carlo Tree Search. In *Forty-second International Conference on Machine Learning*.
- [30] Changlun Li, Chenyu Yang, Yuyu Luo, Ju Fan, and Nan Tang. 2025. Weak-to-Strong Prompts with Lightweight-to-Powerful LLMs for High-Accuracy, Low-Cost, and Explainable Data Transformation. *Proceedings of the VLDB Endowment* (2025).
- [31] Guoliang Li, Jiayi Wang, Chenyang Zhang, and Jiannan Wang. 2025. Data+AI: LLM4Data and Data4LLM. In *Companion of the 2025 International Conference on Management of Data*.
- [32] Guoliang Li, Xuanhe Zhou, and Xinyang Zhao. 2024. LLM for data management. *Proceedings of the VLDB Endowment* (2024).
- [33] Peng Li, Yeye He, Dror Yashar, Weiwei Cui, Song Ge, Haidong Zhang, Danielle Rifinski Fainman, Dongmei Zhang, and Surajit Chaudhuri. 2024. Table-gpt: Table fine-tuned gpt for diverse table tasks. *Proceedings of the ACM on Management of Data* (2024).
- [34] Shuaimin Li, Xuanang Chen, Yuanfeng Song, Yunze Song, Chen Jason Zhang, Fei Hao, and Lei Chen. 2025. prompt4vis: prompting large language models with example mining for tabular data visualization. *The VLDB Journal* (2025).
- [35] Yuchen Li, Kai Wang, Zhiqiang Sun, et al. 2025. Automatic Database Description Generation for Text-to-SQL. *arXiv preprint arXiv:2502.20657* (2025).
- [36] Zhaodonghui Li, Haitao Yuan, Huiming Wang, Gao Cong, and Lidong Bing. 2024. LLM-R2: A Large Language Model Enhanced Rule-Based Rewrite System for Boosting Query Efficiency. *Proceedings of the VLDB Endowment* (2024).
- [37] Bang Liu, Xinfeng Li, Jiayi Zhang, Jinlin Wang, Tanjin He, Sirui Hong, Hongzhang Liu, Shaokun Zhang, Kaitao Song, Kunlun Zhu, et al. 2025. Advances and challenges in foundation agents: From brain-inspired intelligence to evolutionary, collaborative, and safe systems. *arXiv preprint arXiv:2504.01990* (2025).
- [38] Xinyu Liu, Shuyu Shen, Boyan Li, Peixian Ma, Runzhi Jiang, Yuxin Zhang, Ju Fan, Guoliang Li, Nan Tang, and Yuyu Luo. 2025. A Survey of Text-to-SQL in the Era of LLMs: Where are we, and where are we going? *IEEE Transactions on Knowledge and Data Engineering* (2025).
- [39] Xinyu Liu, Shuyu Shen, Boyan Li, Nan Tang, and Yuyu Luo. 2025. Nl2sql-bugs: A benchmark for detecting semantic errors in nl2sql translation. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*.
- [40] Tianqi Luo, Chuhuan Huang, Leixian Shen, Boyan Li, Shuyu Shen, Wei Zeng, Nan Tang, and Yuyu Luo. 2025. nvBench 2.0: Resolving Ambiguity in Text-to-Visualization through Stepwise Reasoning. *arXiv:2503.12880* [cs.CL].
- [41] Yuyu Luo, Guoliang Li, Ju Fan, Chengliang Chai, and Nan Tang. 2025. Natural language to sql: State of the art and open problems. *Proceedings of the VLDB Endowment* (2025).
- [42] Paula Maddigan and Teo Sosnaj. 2023. Chat2VIS: Generating Data Visualizations via Natural Language Using ChatGPT, Codex and GPT-3 Large Language Models. *IEEE Access* (oct 2023).
- [43] Sherwin Minaee, Tomas Mikolov, Narjes Nikzad, Meysam Chenaghlu, Richard Socher, Xavier Amatriain, and Jianfeng Gao. 2024. Large language models: A survey. *arXiv preprint arXiv:2402.06196* (2024).
- [44] Zan Ahmad Naeem, Mohammad Shahmeer Ahmad, Mohamed Eltabakh, Mourad Ouzzani, and Nan Tang. 2024. RetClean: Retrieval-Based Data Cleaning Using LLMs and Data Lakes. *Proc. VLDB Endow.* (2024).
- [45] Avanika Narayan, Ines Chami, Laurel Orr, and Christopher Ré. 2022. Can Foundation Models Wrangle Your Data? *Proceedings of the VLDB Endowment* (2022).
- [46] Geling Ouyang, Jingyao Chen, Zhihe Nie, Yi Gui, Yao Wan, Hongyu Zhang, and Dongping Chen. 2025. nvAgent: Automated Data Visualization from Natural Language via Collaborative Agent Workflow. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics*.
- [47] Fatemeh Pesaran Zadeh, Juyeon Kim, Jin-Hwa Kim, and Gunhee Kim. [n. d.]. Text2Chart31: Instruction Tuning for Chart Generation with Automatic Feedback. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*.
- [48] Mohammadreza Pourreza, Hailong Li, Ruoxi Sun, Yeounoh Chung, Shayan Talaee, Gaurav Tarlok Kakkar, Yu Gan, Amin Saberi, Fatma Ozcan, and Sercan O Arik. 2025. CHASE-SQL: Multi-Path Reasoning and Preference Optimized Candidate Selection in Text-to-SQL. In *The Thirteenth International Conference on Learning Representations*.
- [49] Mohammadreza Pourreza and Davood Rafie. 2023. DIN-SQL: Decomposed In-Context Learning of Text-to-SQL with Self-Correction. In *Advances in Neural Information Processing Systems 36*, Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine (Eds.).
- [50] Danrui Qi, Zhengjie Miao, and Jiannan Wang. 2025. CleanAgent: Automating Data Standardization with LLM-based Agents. *arXiv:2403.08291*
- [51] Zhangcheng Qiang, Weiqing Wang, and Kerry Taylor. 2024. Agent-OM: Leveraging LLM Agents for Ontology Matching. *Proceedings of the VLDB Endowment* (2024).

[52] Jon Saad-Falcon, Joe Barrow, Alexa Siu, Ani Nenkova, Seunghyun Yoon, Ryan A. Rossi, and Franck Dernoncourt. [n. d.]. PDFTriage: Question Answering over Long, Structured Documents. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track*.

[53] Elisabeth Shi, Tom Michael Gasser, Andre Seeck, and Rico Auerswald. 2020. The Principles of Operation Framework: A Comprehensive Classification Concept for Automated Driving Functions. *SAE International Journal of Connected and Automated Vehicles* (2020).

[54] Zhihao Shuai, Boyan Li, Siyu Yan, Yuyu Luo, and Weikai Yang. 2025. Deepvis: Bridging natural language and data visualization through step-wise reasoning. *arXiv preprint arXiv:2508.01700* (2025).

[55] Snowflake. [n. d.]. *Cortex Agents*.

[56] Yuyang Song, Hanxu Yan, Jiale Lao, Yibo Wang, Yufei Li, Yuanchun Zhou, Jianguo Wang, and Mingjie Tang. 2025. QUITE: A Query Rewrite System Beyond Rules with LLM Agents. *arXiv preprint arXiv:2506.07675* (2025).

[57] Yuan Sui, Mengyu Zhou, Mingjie Zhou, Shi Han, and Dongmei Zhang. 2024. Table Meets LLM: Can Large Language Models Understand Structured Table Data? A Benchmark and Empirical Study. *arXiv:2305.13062* [cs]

[58] Nicole Sultanum and Arjun Srinivasan. 2023. DATATALES: Investigating the use of Large Language Models for Authoring Data-Driven Articles. In *2023 IEEE Visualization and Visual Analytics (VIS)*.

[59] Ji Sun, Guoliang Li, Peiyao Zhou, Yihui Ma, Jingzhe Xu, and Yuan Li. 2025. AgenticData: An Agentic Data Analytics System for Heterogeneous Data. *CoRR* (2025).

[60] Wenwen Sun, Zhicheng Pan, Zirui Hu, Yu Liu, Chengcheng Yang, Rong Zhang, and Xuan Zhou. [n. d.]. Rabbit: Retrieval-Augmented Generation Enables Better Automatic Database Knob Tuning. In *2025 IEEE 41st International Conference on Data Engineering (ICDE)*. IEEE.

[61] Zhaoyan Sun, Jiayi Wang, Xinyang Zhao, Jiachi Wang, and Guoliang Li. 2025. Data Agent: A Holistic Architecture for Orchestrating Data+ AI Ecosystems. *arXiv preprint arXiv:2507.01599* (2025).

[62] Zhaoyan Sun, Xuanhe Zhou, Guoliang Li, Xiang Yu, Jianhua Feng, and Zhangm Yong. 2025. R-Bot: An LLM-based Query Rewrite System. *Proceedings of the VLDB Endowment* (2025).

[63] Manan Suri, Puneet Mathur, Franck Dernoncourt, Kanika Goswami, Ryan A. Rossi, and Dinesh Manocha. [n. d.]. VisDoM: Multi-Document QA with Visually Rich Elements Using Multimodal Retrieval-Augmented Generation. In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies*.

[64] Manan Suri, Puneet Mathur, Nedim Lipka, Franck Dernoncourt, Ryan A. Rossi, and Dinesh Manocha. 2025. ChartLens: Fine-grained Visual Attribution in Charts. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics*.

[65] TabTab AI. [n. d.]. *TabTab*.

[66] Zirui Tang, Weizheng Wang, Zihang Zhou, Yang Jiao, Bangrui Xu, Boyu Niu, Xuanhe Zhou, Guoliang Li, Yeye He, Wei Zhou, et al. 2025. LLM/Agent-as-Data-Analyst: A Survey. *arXiv preprint arXiv:2509.23988* (2025).

[67] Fengwei Teng, Zhaoyang Yu, Quan Shi, Jiayi Zhang, Chenglin Wu, and Yuyu Luo. 2025. Atom of Thoughts for Markov LLM Test-Time Scaling. *CoRR* (2025).

[68] Fen Wang, Bomiao Wang, Xueli Shu, Zhen Liu, Zekai Shao, Chao Liu, and Siming Chen. 2025. ChartInsighter: An Approach for Mitigating Hallucination in Time-series Chart Summary Generation with A Benchmark Dataset. *arXiv:2501.09349* [cs.CL]

[69] Jiayi Wang and Guoliang Li. 2025. AOP: Automated and interactive llm pipeline orchestration for answering complex queries. *CIDR*.

[70] Jiayi Wang, Guoliang Li, and Jianhua Feng. 2025. iDataLake: An llm-powered analytics system on data lakes. *Data Engineering* (2025).

[71] Yuhao Wang, Ruiyang Ren, Junyi Li, Xin Zhao, Jing Liu, and Ji-Rong Wen. [n. d.]. REAR: A Relevance-Aware Retrieval-Augmented Framework for Open-Domain Question Answering. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*.

[72] Wenyi Xu, Yuren Mao, Xiaolu Zhang, Chao Zhang, Xuemei Dong, Mengfei Zhang, and Yunjun Gao. 2025. DAgent: A Relational Database-Driven Data Analysis Report Generation Agent. *arXiv:2503.13269* [cs.DB]

[73] Zihan Yan, Rui Xi, and Mengshu Hou. 2025. MCTuner: Spatial Decomposition-Enhanced Database Tuning via LLM-Guided Exploration. *arXiv preprint arXiv:2509.06298* (2025).

[74] Chenyu Yang, Yuyu Luo, Chuanxuan Cui, Ju Fan, Chengliang Chai, and Nan Tang. 2025. Data Imputation with Limited Data Redundancy Using Data Lakes. *Proc. VLDB Endow.* (2025). doi:10.14778/3748191.3748200

[75] Zhaorui Yang, Bo Pan, Han Wang, Yiyao Wang, Xingyu Liu, Minfeng Zhu, Bo Zhang, and Wei Chen. 2025. Multimodal DeepResearcher: Generating Text-Chart Interleaved Reports From Scratch with Agentic Framework. *arXiv:2506.02454* [cs.CL]

[76] Zhiyuan Yang, Zihan Zhou, Shuo Wang, Xin Cong, Xu Han, Yukun Yan, Zhenghao Liu, Zhixing Tan, Pengyuan Liu, Dong Yu, Zhiyuan Liu, Xiaodong Shi, and Maosong Sun. [n. d.]. MatPlotAgent: Method and Evaluation for LLM-Based Agentic Scientific Data Visualization. In *Findings of the Association for Computational Linguistics 2024*.

[77] Yunhu Ye, Binyuan Hui, Min Yang, Binhu Li, Fei Huang, and Yongbin Li. 2023. Large Language Models Are Versatile Decomposers: Decomposing Evidence and Questions for Table-based Reasoning. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*.

[78] Peiyong Yu, Guoxin Chen, and Jingjing Wang. 2025. Table-Critic: A Multi-Agent Framework for Collaborative Criticism and Refinement in Table Reasoning. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics*.

[79] Haoxiang Zhang, Yurong Liu, Aécio Santos, Juliana Freire, et al. 2025. Autoddg: Automated dataset description generation using large language models. *arXiv preprint arXiv:2502.01050* (2025).

[80] Shaolei Zhang, Ju Fan, Meihao Fan, Guoliang Li, and Xiaoyong Du. 2025. DeepAnalyze: Agentic Large Language Models for Autonomous Data Science. *arXiv:2510.16872* [cs.AI]

[81] Yunfan Zhang, Changlun Li, Yuyu Luo, and Nan Tang. 2024. Sketchfill: Sketch-guided code generation for imputing derived missing values. *arXiv preprint arXiv:2412.19113* (2024).

[82] Zhengxuan Zhang, Zhuowen Liang, Yin Wu, Teng Lin, Yuyu Luo, and Nan Tang. 2025. DataPuzzle: Breaking Free from the Hallucinated Promise of LLMs in Data Analysis. *arXiv preprint arXiv:2504.10036* (2025).

[83] Xinyang Zhao, Xuanhe Zhou, and Guoliang Li. 2023. Automatic database knob tuning: A survey. *IEEE Transactions on Knowledge and Data Engineering* (2023).

[84] Wei Zhou, Ji Sun, Xuanhe Zhou, Guoliang Li, Luyang Liu, Hao Wu, and Tianyuan Wang. 2025. GaussMaster: An LLM-based Database Copilot System. *arXiv preprint arXiv:2506.23322* (2025).

[85] Xuanhe Zhou, Junxuan He, Wei Zhou, Haodong Chen, Zirui Tang, Haoyu Zhao, Xin Tong, Guoliang Li, Youmin Chen, Jun Zhou, et al. 2025. A Survey of LLM×DATA. *arXiv preprint arXiv:2505.18458* (2025).

[86] Xuanhe Zhou, Guoliang Li, and Zhiyuan Liu. 2023. LLM as DBA. *arXiv preprint arXiv:2308.05481* (2023).

[87] Xuanhe Zhou, Guoliang Li, Zhaoyan Sun, Zhiyuan Liu, Weize Chen, Jianming Wu, Jiesi Liu, Ruohang Feng, and Guoyang Zeng. 2024. D-Bot: Database Diagnosis System using Large Language Models. *Proceedings of the VLDB Endowment* (2024).

[88] Yizhang Zhu, Shiyin Du, Boyan Li, Yuyu Luo, and Nan Tang. 2024. Are large language models good statisticians? *Advances in Neural Information Processing Systems* (2024).

[89] Yizhang Zhu, Runzhi JIANG, Boyan Li, Nan Tang, and Yuyu Luo. 2025. EllieSQL: Cost-Efficient Text-to-SQL with Complexity-Aware Routing. In *Second Conference on Language Modeling*.

[90] Yizhang Zhu, Liangwei Wang, Chenyu Yang, Xiaotian Lin, Boyan Li, Wei Zhou, Xinyu Liu, Zhangyang Peng, Tianqi Luo, Yu Li, Chengliang Chai, Chong Chen, Shimin Di, Ju Fan, Ji Sun, Nan Tang, Fugee Tsung, Jannan Wang, Chenglin Wu, Yanwei Xu, Shaolei Zhang, Yong Zhang, Xuanhe Zhou, Guoliang Li, and Yuyu Luo. 2025. A Survey of Data Agents: Emerging Paradigm or Overstated Hype? *arXiv:2510.23587* [cs.DB] <https://arxiv.org/abs/2510.23587>