

LLM-FK: Multi-Agent LLM Reasoning for Foreign Key Detection in Large-Scale Complex Databases

Zijian Tang^{1,2,4}, Ying Zhang^{2,3,4*}, Sib0 Cai^{5,6*}, Ruoxuan Wang^{1,2,4}

¹School of Software and Microelectronics, Peking University

²National Engineering Research Center of Software Engineering, Peking University

³Key Laboratory of High Confidence Software Technologies, Peking University

⁴Beijing Tongming Lake Information Technology Application Innovation Center

⁵Engineering Research Center of Integration and Application of

Digital Learning Technology, Ministry of Education ⁶The Open University of China

zjtang25@stu.pku.edu.cn, zhang.ying@pku.edu.cn, caisb@ouchn.edu.cn

Abstract

Detecting missing foreign keys (FKs) requires accurately modeling semantic dependencies across database schemas, which conventional heuristic-based methods are fundamentally limited in capturing. We propose **LLM-FK**, the first fully automated multi-agent framework for FK detection, designed to address three core challenges that hinder naive LLM-based solutions in large-scale complex databases: combinatorial search space explosion, ambiguous inference under limited context, and global inconsistency arising from isolated local predictions. LLM-FK coordinates four specialized agents: a Profiler that decomposes the FK detection problem into the task of validating FK candidate column pairs and prunes the search space via a unique-key-driven schema decomposition strategy; an Interpreter that injects self-augmented domain knowledge; a Refiner that constructs compact structural representations and performs multi-perspective chain-of-thought reasoning; and a Verifier that enforces schema-wide consistency through a holistic conflict resolution strategy. Experiments on five benchmark datasets demonstrate that LLM-FK consistently achieves F1-scores above 93%, surpassing existing baselines by 15% on the large-scale MusicBrainz database, while reducing the candidate search space by two to three orders of magnitude without losing true FKs and maintaining robustness under challenging conditions like missing data. These results demonstrate the effectiveness and scalability of LLM-FK in real-world databases.

1 Introduction

Foreign key (FK) constraints define directional dependencies between a referencing column (or column set) in one table and a referenced column in another, thereby enforcing referential integrity in relational databases (Melton and Simon, 2001). These constraints play a foundational role in supporting

* Corresponding author.

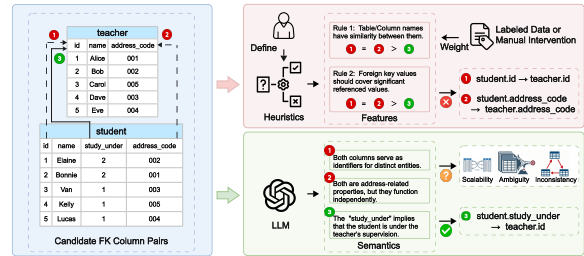


Figure 1: Comparison of FK detection: heuristic limitations vs. LLM semantic reasoning.

essential downstream tasks, including schema reverse engineering, data integration, and semantic querying (Wu et al., 2019; Li et al., 2023). However, in real-world databases, FK constraints are frequently missing due to schema evolution, legacy system migration, or intentional omission for performance reasons, which can undermine data reliability and semantic consistency (Johnson et al., 2003). While manual FK identification may be feasible for small schemas, it quickly becomes impractical for large-scale databases characterized by numerous heterogeneous tables and complex data distributions (Halevy et al., 2006; Hai et al., 2023).

Existing automated FK detection methods predominantly rely on heuristics that exploit syntactic or statistical signals, such as column name similarity and value inclusion ratios (Rostin et al., 2009; Chen et al., 2014; Jiang and Naumann, 2020). Although these heuristics are effective in well-structured or narrowly scoped schemas, they largely fail to capture the deeper semantic dependencies that fundamentally underpin FKs (Lausen, 2007; Kent, 2012). As a result, their performance is highly sensitive to irregularities such as sparse data and unconventional naming patterns, and necessitates manual tuning or high-quality labeled data to remain reliable (Zhang et al., 2010), which significantly limits their scalability and generalizability to large, heterogeneous or evolving databases.

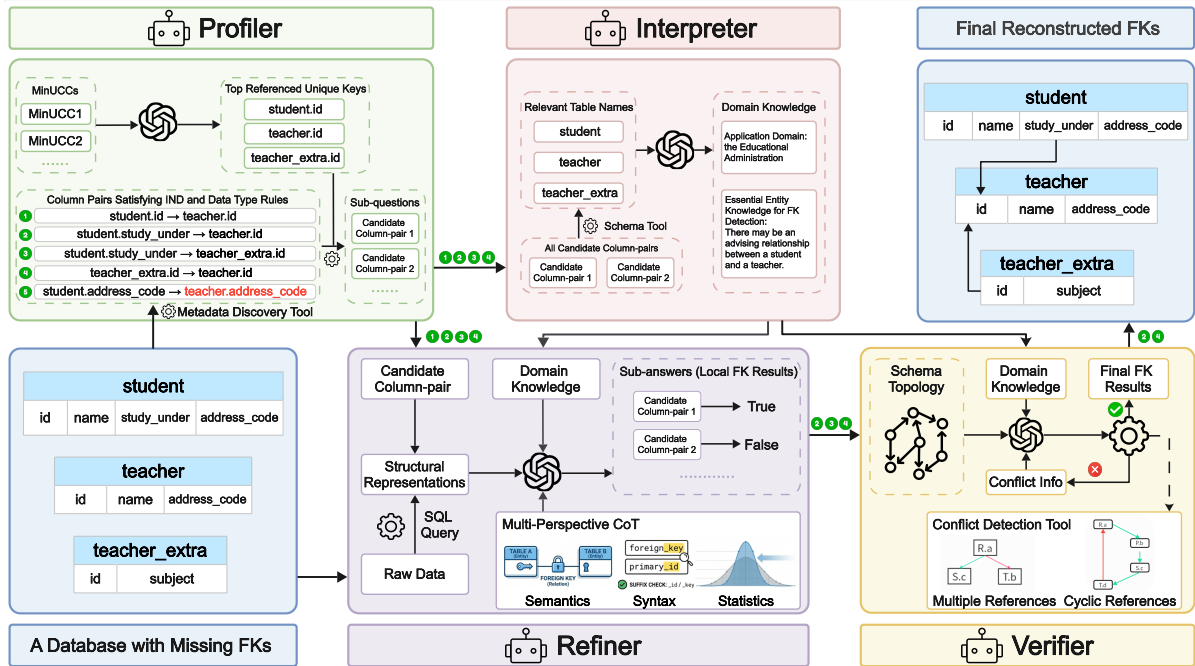


Figure 2: Overview of our LLM-FK framework. The pipeline orchestrates four specialized agents to achieve fully automated FK detection: the Profiler executes the Unique-Key-Driven Schema Decomposition Strategy; the Interpreter performs Self-Augmented Domain Knowledge Injection; the Refiner employs Multi-Perspective CoT Reasoning; and the Verifier applies the Holistic Conflict Resolution Strategy.

Large Language Models (LLMs) offer a promising alternative by implicitly capturing semantic relationships that extend beyond surface-level heuristics; however, directly applying LLMs to large-scale FK detection introduces several fundamental challenges (see Figure 1). First, **combinatorial search space explosion** severely limits scalability: as table and column counts increase, the number of candidate column pairs grows rapidly (e.g., approximately 4,000 pairs for only eight tables (Poess and Floyd, 2000)), rendering exhaustive evaluation computationally infeasible. Second, **ambiguous inference under limited context** complicates accurate reasoning: finite context windows (Ye et al., 2023; Chen et al., 2024) and performance degradation when processing raw tabular data (Cheng et al., 2022; Chen, 2023; Sui et al., 2024b) prevent LLMs from jointly accessing complete schema structures and data distributions, forcing reliance on compressed or partial representations that may omit critical signals. Third, **global inconsistency arising from isolated local predictions** emerges because LLMs typically assess FKs in a local context independently, such as evaluating individual column pairs or table pairs in isolation, leading to globally conflicting decisions such as cyclic references that violate schema-level coherence.

To overcome these limitations, we propose **LLM-FK**, a structured multi-agent framework for fully automated FK detection in large-scale complex databases, requiring neither manual intervention nor labeled data. LLM-FK decomposes schema-wide FK discovery into coordinated reasoning stages, with each stage explicitly addressing one of the aforementioned challenges. To mitigate **combinatorial search space explosion**, the Profiler employs a **Unique-Key-Driven Schema Decomposition Strategy**, leveraging LLMs to identify the most plausible referenced unique key for each potentially referenced table, and thereby decomposing the global detection task into focused validation problems over a pruned column pair candidate set. This constraint-aware decomposition dramatically reduces decision complexity and the search space while preserving recall of FKs. To address **ambiguous inference under limited context**, the Interpreter and Refiner jointly perform **Self-Augmented Domain Knowledge Injection** and **Multi-Perspective Chain-of-Thought (CoT) Reasoning**. Specifically, the Interpreter extracts implicit semantic cues from schema metadata and inferred domain knowledge, while the Refiner integrates this context with compact structural representations, enabling robust reasoning from syn-

tactic, semantic, and statistical perspectives despite limited context. Finally, to ensure **global consistency from local predictions**, the Verifier applies a **Holistic Conflict Resolution Strategy** by constructing a schema-level dependency graph, detecting structural violations such as cyclic references, and performing targeted reasoning to resolve inconsistencies, thereby producing a coherent FK set aligned with relational semantics and overall schema topology.

We evaluate LLM-FK through extensive experiments on five public benchmark databases spanning diverse domains and scales. Overall, LLM-FK demonstrates consistent superiority in effectiveness, efficiency, and robustness. It achieves F1-scores exceeding 93% on all benchmarks and outperforms both heuristic-based and LLM-based baselines by more than 15% on MusicBrainz (Swartz, 2002), a complex real-world database containing over 300 tables. In terms of efficiency, LLM-FK reduces the candidate search space by two to three orders of magnitude without losing any true FKs. Furthermore, it maintains strong robustness under challenging conditions, including missing data and obscure semantic relationships. These results highlight the scalability and generalizability of LLM-FK for FK detection in large-scale complex databases.

Our contributions are summarized as follows:

- We propose LLM-FK, the first multi-agent framework for fully automated FK detection in large-scale complex databases, overcoming the semantic limitations of existing approaches without requiring manual supervision or labeled data.
- We design four specialized agents that systematically address the core challenges of applying LLMs to FK detection: the Profiler reduces the search space via a Unique-Key-Driven Schema Decomposition Strategy; the Interpreter and Refiner enable deep and robust inference through Self-Augmented Domain Knowledge Injection and Multi-Perspective CoT Reasoning; and the Verifier enforces schema-wide coherence through a Holistic Conflict Resolution Strategy.
- Extensive experiments demonstrate that LLM-FK achieves state-of-the-art (SOTA) accuracy, high efficiency, and strong robustness against complex, large-scale schemas.

2 Related Work

2.1 Heuristic-based Methods for FK Detection

Heuristic-based approaches constitute the dominant line of research for automated FK detection, relying on predefined syntactic and statistical signals extracted from database schemas and data instances. These methods can be broadly categorized into rule-weighted approaches and machine learning-based approaches. Rule-weighted methods (Zhang et al., 2010; Chen et al., 2014; Jiang and Naumann, 2020) design a collection of manually crafted heuristics—such as column name similarity and value inclusion ratios—and aggregate them through weighted scoring functions to estimate the plausibility of FK candidates. Although effective in constrained settings, these approaches require labor-intensive manual tuning to determine appropriate heuristic weights and decision thresholds, which limits their adaptability across domains.

Machine learning-based methods attempt to alleviate manual tuning by learning decision functions from labeled examples, thereby automating candidate ranking or classification (Rostin et al., 2009; Yan and He, 2020). However, this paradigm introduces a strong dependency on high-quality labeled training data, which is expensive to obtain and often unavailable for large, real-world database schemas. More fundamentally, because FKs encode semantic relationships rather than purely syntactic and statistical regularities, rigid heuristic signals inevitably break down in large-scale complex databases. Irregularities such as sparse data, non-standard naming conventions, and heterogeneous schema designs frequently produce counterexamples that violate heuristic assumptions (Zhang et al., 2010), leading to degraded accuracy and limited generalization. In contrast, our proposed LLM-FK leverages the semantic reasoning capabilities of LLMs to directly infer FK relationships, enabling automated and robust detection beyond the expressiveness of fixed heuristics.

2.2 LLMs for Database Schema Analysis

LLMs have recently shown strong potential in a range of database schema analysis tasks, including table annotation, entity matching, and join prediction (Narayan et al., 2022; Fernandez et al., 2023; Jaimovitch-López et al., 2023). These studies demonstrate that LLMs can effectively capture semantic information from schema metadata and natural language descriptions. Despite this

progress, systematic investigation of LLM-based database-wide FK detection remains largely unexplored. Closely related works focus on pairwise join prediction. For instance, Chorus (Kayali et al., 2023) leverages Python code completion prompting to predict joinability between tables, but primarily targets ad-hoc analytical joins rather than formal FK constraints. As a result, they do not explicitly model the semantic directionality, uniqueness requirements, or referential integrity guarantees that distinguish FKs from general joins.

Moreover, MMQA preliminarily explores directly prompting LLMs to identify FKs between table pairs as a prerequisite for multi-table question answering (Wu et al., 2025); however, such generic LLM-based reasoning paradigms, including few-shot prompting and chain-of-thought reasoning, are difficult to apply naively to large-scale complex databases (Brown et al., 2020; Wei et al., 2022). These approaches suffer from three fundamental limitations: combinatorial search space explosion when enumerating column pairs, ambiguous inference under limited context, and global inconsistency arising from isolated local predictions. Without structural decomposition or coordination mechanisms, naive LLM applications struggle to scale while maintaining schema-level coherence. To address these challenges, we propose LLM-FK, a multi-agent framework in which each agent is explicitly designed to resolve one of the limitations of applying LLMs to FK detection, enabling scalable, consistent, and semantically grounded inference.

3 LLM-FK Approach

3.1 Problem Definition

FK detection aims to recover implicit referential constraints from a given relational database schema and its data instances. Formally, given a relational database $\mathcal{D} = \{T_1, \dots, T_n\}$, each table T_i is defined as a tuple $(\mathcal{C}_i, \mathcal{R}_i)$, where \mathcal{C}_i denotes the set of columns (schema) and \mathcal{R}_i denotes the finite set of data tuples (instance). The objective of FK detection is to identify a set of elementary references $\Phi = \{(c_f, c_p) \mid c_f \in \mathcal{C}_i, c_p \in \mathcal{C}_j\}$, where each pair (c_f, c_p) represents a single-column FK relationship in which the referencing column c_f refers to the referenced column c_p . This formulation generalizes to composite FKs (Harrington, 2016), which are modeled as combinations of multiple elementary references on distinct columns between the same source and target tables. By reducing FK detection

to the identification of elementary column-level references, this definition enables fine-grained reasoning while remaining compatible with higher-order relational constraints.

3.2 Preliminaries: UCC and IND

Unique Column Combination (UCC) and Minimal UCC (MinUCC). UCCs capture the notion of key-like column sets within a table. Specifically, for a table T_i , a subset of columns $X \subseteq \mathcal{C}_i$ constitutes a *UCC* if the projection of T_i onto X uniquely identifies every tuple. Formally, for any two distinct tuples $t_u, t_v \in \mathcal{R}_i$, their projected values must satisfy $t_u[X] \neq t_v[X]$. A *MinUCC* further enforces minimality, requiring that no proper subset $X' \subset X$ preserves the uniqueness property. MinUCCs are particularly important in FK detection, as they represent candidate keys and therefore constitute the only valid targets for FK references.

Inclusion Dependency (IND). An *IND*, denoted as $T_i[X] \subseteq T_j[Y]$, holds when the set of values appearing in column set $X \subseteq \mathcal{C}_i$ is contained within the set of values of $Y \subseteq \mathcal{C}_j$. INDs provide a necessary structural condition for referential constraints: any FK from X to Y must satisfy the corresponding IND. In this work, we focus exclusively on single-column INDs, which align with our elementary reference formulation and serve as a reliable basis for FK candidate pruning.

3.3 Overview of LLM-FK

We propose **LLM-FK**, a multi-agent framework for fully automated FK detection in large-scale complex databases. As illustrated in Figure 2, LLM-FK consists of four specialized agents—**Profiler**, **Interpreter**, **Refiner**, and **Verifier**—each responsible for a distinct reasoning stage and equipped with a dedicated strategy. Together, these agents decompose schema-wide FK discovery into coordinated, tractable sub-tasks that collectively address the limitations of naive LLM-based inference.

The **Profiler** initiates the workflow by decomposing the global detection task into column-pair validation problems and aggressively narrowing the candidate space through a Unique-Key-Driven Schema Decomposition Strategy. The **Interpreter** then performs Self-Augmented Domain Knowledge Injection to expose implicit semantic signals embedded in the schema. Building on this enriched context, the **Refiner** conducts deep inference using Multi-Perspective CoT Reasoning to evaluate

each candidate FK. Finally, the **Verifier** aggregates pairwise predictions and applies a Holistic Conflict Resolution Strategy to eliminate inconsistencies and ensure global schema coherence. We detail each component and strategy in the following subsections, with additional implementation details provided in Appendix B.

3.4 Unique-Key-Driven Decomposition

The Profiler serves as the entry point of the LLM-FK pipeline, with the primary goal of optimizing LLM inputs for scalability and reliability. It explicitly addresses two fundamental challenges: (1) the infeasibility of direct schema-wide FK detection due to limited LLM context windows and the risk of hallucinations, and (2) the combinatorial explosion of candidate FKs as schema size and complexity grow. To this end, we introduce a **Unique-Key-Driven Schema Decomposition Strategy** that systematically transforms the global FK detection problem into a set of manageable column-level validation tasks while aggressively pruning structurally invalid candidates.

Schema Decomposition. We decompose FK detection into independent binary classification problems over directed column pairs. Specifically, each candidate (c_f, c_p) is evaluated in isolation, with the LLM determining whether c_f functionally references c_p . Compared to coarser decomposition granularities such as table-pair reasoning, column-level decomposition isolates the atomic unit of FK semantics. This design keeps the sub-problem structure invariant to overall schema scale, focuses the LLM on a single directed hypothesis at one time, and substantially reduces decision complexity and hallucination risk.

Candidate Pruning. Although atomic decomposition simplifies individual decisions, it naively induces a combinatorial explosion in the number of column pairs. To ensure computational feasibility, we employ a pruning strategy to aggressively refine the candidate set. First, we apply IND-based validation and rule-based filters: INDs enforce necessary value inclusion constraints, while rule-based checks eliminate pairs with incompatible data types or semantically unsuitable types (e.g., Boolean or floating-point columns). Second, we exploit the structural property that FKs must reference unique keys. Instead of relying on primary keys (PKs)—which are frequently missing, incomplete, or overly composite—we identify MinUCCs

as candidate referenced keys. For each potentially referenced table, the LLM infers the most plausible referenced unique key from its MinUCC set. As shown in Figure 2, we then discard any candidate column pair whose referenced column does not belong to the selected unique keys. This process yields a compact yet comprehensive candidate set that preserves recall of true FKs while eliminating structurally invalid FK hypotheses.

3.5 Knowledge-Augmented Reasoning

To accurately assess pruned candidate pairs under the constraints of finite LLM context windows, LLM-FK coordinates two agents—the Interpreter and the Refiner—to perform structured, knowledge-augmented reasoning. The Interpreter focuses on Self-Augmented Domain Knowledge Injection to resolve semantic ambiguity, while the Refiner employs Multi-Perspective CoT Reasoning to conduct deep, evidence-based inference.

Self-Augmented Domain Knowledge Injection.

The Interpreter equips the framework with a global semantic understanding of the database by analyzing the names of all tables involved in the candidate set. We intentionally restrict this analysis to table names, as they capture core entity semantics while remaining sufficiently concise to fit within the LLM’s context budget. By inferring the overarching application domain and salient entity concepts, the Interpreter bridges the gap between abstract schema symbols and real-world semantics. Crucially, this process enables *self-augmentation* (Sui et al., 2024a): schema-wide domain insights are injected into each reasoning task, reducing ambiguity when candidate column pairs are evaluated within an isolated local context. The Interpreter performs this analysis once per database, ensuring consistent semantic grounding across all subsequent validations without redundant computation.

Multi-Perspective CoT Reasoning. Building on the injected domain knowledge, the Refiner implements a multi-perspective CoT framework that integrates complementary signals for robust inference. To this end, the Refiner issues SQL queries to extract and serialize compact, high-density representations of relevant information, including: (1) schema definitions (table and column names); (2) statistical metadata (ordinal positions, data types, value ranges, cardinalities, row counts, and average text lengths); (3) inter-column dependencies (value coverage ratios and table size ratios); and (4)

sample data instances (top-5 rows). This concise serialization avoids the noise and cost associated with full-table inputs while preserving sufficient and discriminative evidence to prevent misinterpretations in raw data processing, such as ambiguous data types. The Refiner then reasons from three complementary perspectives: a syntactic perspective based on naming conventions, a statistical perspective grounded in distributional regularities, and a semantic perspective focused on entity relationships. By emulating expert human reasoning, this multi-perspective process enables cross-perspective compensation, allowing reliable signals to correct judgments when others are compromised, such as in obfuscated schemas or data-sparse tables.

3.6 Holistic Conflict Resolution

To ensure schema-wide structural consistency, the Verifier applies a Holistic Conflict Resolution Strategy that operates over aggregated pairwise predictions. It constructs a schema topology graph encoding inferred FK relationships and then identifies structural violations arising from local decisions. In particular, we focus on two common conflict types: *multiple references*, where a single column references multiple distinct targets, and *cyclic references*, where FK relationships form closed loops.

Conflict resolution proceeds in two stages. First, for multiple-reference conflicts, the Verifier reevaluates competing candidates using the same knowledge-augmented reasoning framework described in Section 3.5, retaining only the most semantically and structurally plausible FK. Second, for cyclic references, the Verifier employs an iterative elimination strategy that prioritizes the shortest cycles, which represent the most immediate violations of relational semantics. Within each cycle, the least plausible reference is removed, and the process repeats until the graph converges to an acyclic and unambiguous state. Crucially, the Refiner’s semantic-based reasoning typically constrains the candidates into tightly clustered, semantically related subgraphs. Leveraging this localization, our approach yields two primary advantages. First, unlike methods that rely on isolated confidence scores, evaluating conflicting candidates collectively within these specific structures facilitates direct semantic comparison, accurately capturing intrinsic semantic hierarchical relationships. Second, it significantly reduces computational complexity; by confining conflict resolution to these highly localized immediate contexts (e.g., length-2

cyclic references), our approach intrinsically avoids the cascading complexity typically associated with global graph traversals. Through this procedure, the Verifier guarantees that the final FK set is globally coherent with relational database principles.

4 Experiments

4.1 Experimental Setup

Datasets and Evaluation Metrics. We evaluate LLM-FK on five widely adopted relational database benchmarks: TPC-H (Poess and Floyd, 2000), Northwind (NWind) (Microsoft, 2023b), TPC-E (Chen et al., 2011), AdventureWorks (AdvWorks) (Microsoft, 2023a), and MusicBrainz (MusicBz) (Swartz, 2002). These datasets span diverse application domains—including commerce, retail, finance, manufacturing, and entertainment—and exhibit substantial variation in schema scale, naming conventions, and data characteristics. Together, they capture a broad spectrum of real-world challenges for FK detection, ranging from clean benchmark schemas to large-scale, heterogeneous databases. Detailed dataset statistics are reported in Appendix C.

To assess FK detection quality, we adopt standard Precision (P), Recall (R), and F1-score (F) (Goutte and Gaussier, 2005). All metrics are computed over the candidate set produced by the Profiler, which defines the effective search space for subsequent reasoning. This evaluation protocol isolates the quality of FK inference from trivial candidates eliminated during early pruning.

Baselines. We compare LLM-FK against seven representative baselines spanning three methodological paradigms. (1) *Rule-weighted methods.* MC-FK (Zhang et al., 2010) relies exclusively on data statistics under uniform reference assumptions. Fast-FK (Chen et al., 2014) combines syntactic similarity with cardinality-based signals. HOPF (Jiang and Naumann, 2020) further exploits PK constraints to improve detection performance. (2) *Machine learning-based methods.* ML-FK (Rostin et al., 2009) trains classifiers on basic syntactic and statistical features. AutoSuggest (Yan and He, 2020) extends this approach by incorporating additional metadata, including ordinal positions. (3) *LLM-based methods.* Chorus (Kayali et al., 2023) predicts joinability via Python code completion prompting. End-to-End (Wu et al., 2025) directly queries LLMs to identify FKs for each table pair in a single pass. Few-Shot (Brown et al., 2020)

Category	Method	TPC-H			Northwind			TPC-E			AdventureWorks			MusicBrainz		
		P	R	F	P	R	F	P	R	F	P	R	F	P	R	F
Rule-weighted	MC-FK	0.56	1.00	0.72	<u>0.71</u>	<u>0.91</u>	0.80	0.32	0.78	0.46	0.46	0.59	0.52	0.51	0.81	0.62
	Fast-FK	0.50	0.78	0.61	0.67	<u>0.91</u>	0.77	<u>0.82</u>	<u>0.91</u>	<u>0.86</u>	0.54	<u>0.84</u>	0.65	0.62	0.81	0.70
	HOPF	1.00	0.10	0.20	0.67	0.19	0.29	0.78	0.16	0.26	0.57	0.45	0.46	<u>0.75</u>	<u>0.86</u>	<u>0.80</u>
ML-based	ML-FK	0.67	0.67	0.67	0.00	0.00	0.00	0.79	0.58	0.67	<u>0.77</u>	0.22	0.34	0.35	0.49	0.41
	AutoSuggest	0.50	0.22	0.31	0.48	<u>0.91</u>	0.62	0.45	0.64	0.53	0.24	0.54	0.33	0.30	0.23	0.26
LLM-based	Chorus	0.50	0.44	0.47	0.46	0.55	0.50	0.32	0.47	0.38	0.13	0.29	0.18	0.26	0.32	0.29
	End-to-End	<u>0.82</u>	1.00	<u>0.90</u>	1.00	1.00	1.00	0.57	0.87	0.69	0.72	0.75	<u>0.73</u>	0.70	0.81	0.75
	Few-Shot	0.80	<u>0.89</u>	0.84	1.00	<u>0.91</u>	<u>0.95</u>	0.46	0.71	0.56	0.63	0.68	0.66	0.70	0.70	0.70
	CoT	<u>0.82</u>	1.00	<u>0.90</u>	1.00	1.00	1.00	0.60	0.87	0.71	0.71	0.74	0.72	0.68	0.81	0.74
Ours	LLM-FK	1.00	1.00	1.00	1.00	1.00	1.00	0.91	0.96	0.93	0.93	0.96	0.94	0.91	1.00	0.95

Table 1: Performance comparison of FK detection across five datasets. **Bold** indicates the best performance, and underlined indicates the second-best.

augments this setting with explicit FK exemplars. CoT (Wei et al., 2022) applies chain-of-thought prompting to elicit step-by-step reasoning. These baselines collectively represent the current SOTA across heuristic-based and LLM-driven paradigms.

Implementation Details. All LLM-based methods are implemented using **DeepSeek-R1** (Guo et al., 2025) with the temperature set to 0 to ensure deterministic and reproducible outputs. To prevent data leakage (Balloccu et al., 2024), database names are masked during inference. In addition, empty columns are excluded during rule-based pruning in the Profiler to ensure fair comparison with data-driven baselines. Further implementation details and reproducibility considerations are provided in Appendix D. Supplementary experiments—including evaluations on other LLMs, a detailed computational cost analysis, and other additional results—along with qualitative case studies, are presented in Appendix E.

4.2 Main Results

Table 1 summarizes the overall FK detection performance. The results reveal several key findings: **First**, LLM-FK consistently achieves SOTA performance across all five datasets, outperforming the strongest baselines by substantial margins on complex schemas (e.g., +21% F1-score on AdventureWorks and +15% on MusicBrainz). **Second**, LLM-FK exhibits markedly higher accuracy and generalizability than heuristic-based methods. While the strongest heuristic-based baseline (Fast-FK) attains F1-scores ranging from 0.61 to 0.86—exhibiting large variance across datasets—LLM-FK maintains stable performance between 0.93 and 1.00. This improvement stems from LLM-FK’s ability to reason about semantic intent, which is funda-

mental to FK design, rather than relying solely on brittle syntactic or statistical signals that are sensitive to irregularities. **Third**, while generic LLM-based baselines degrade sharply on complex schemas like AdventureWorks and MusicBrainz, LLM-FK demonstrates superior stability. Specifically, our multi-agent collaboration framework decomposes FK detection into invariant, binary directed column-pair validation tasks to reduce decision complexity, executes deep reasoning despite limited context, and enforces global schema-wide consistency, thereby effectively overcoming the challenges of native LLM-based paradigms on large-scale complex databases. Moreover, the poor performance of Chorus confirms that join prediction alone fails to capture the semantic directionality and integrity constraints intrinsic to FKs. **Finally**, LLM-FK shows strong robustness under semantic opacity, as exemplified by TPC-E, which contains numerous ambiguous English abbreviations. While generic LLM-based methods struggle in this setting, LLM-FK achieves an F1-score of 0.93, benefiting from Self-Augmented Domain Knowledge Injection and Multi-Perspective CoT Reasoning that jointly disambiguate local column pairs using global and structural cues. Overall, these results demonstrate that LLM-FK provides accurate, generalizable, and robust FK detection for large-scale complex databases.

4.3 Analysis of Candidate Pruning

We analyze the pruning strategy from two perspectives: search space reduction and pruning safety. Figure 3 shows the number of remaining candidates after each pruning stage, compared to a commonly used table-level baseline defined by all pairwise combinations of tables.

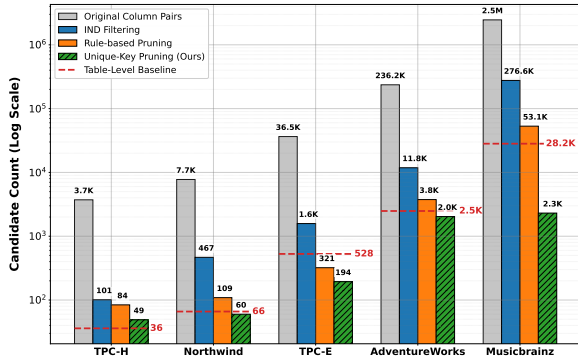


Figure 3: Remaining candidate counts after various stages of pruning compared to the table-level baseline.

Configuration	TPC-H	NWind	TPC-E	AdvWorks	MusicBz
<i>Analysis A: Refiner Agent</i>					
w/ Injection (Base)	0.78	0.92	0.70	0.83	0.87
w/o Injection	0.78	0.92	0.61	0.74	0.76
Perf. Drop (Δ)	–	–	-9%	-9%	-11%
<i>Analysis B: Verifier Agent</i>					
w/ Injection (Full)	1.00	1.00	0.93	0.94	0.95
w/o Injection	0.78	1.00	0.89	0.94	0.95
Perf. Drop (Δ)	-22%	–	-4%	–	–

Table 2: Ablation study on the impact of Self-Augmented Domain Knowledge Injection.

Search Space Reduction. The proposed multi-stage pruning strategy achieves dramatic compression of the FK candidate space. On the large-scale MusicBrainz dataset, the number of candidate pairs is reduced from over 2.46×10^6 to only 2,306, corresponding to a reduction of nearly three orders of magnitude. Across most datasets, the resulting candidate sets are even smaller than those produced by table-level enumeration. This reduction directly mitigates the combinatorial explosion inherent in large schemas and drastically decreases the number of required LLM invocations, ensuring that downstream reasoning remains computationally feasible.

Pruning Safety. Importantly, this aggressive pruning does not sacrifice recall. Across all five datasets, the number of pruned true FKs is exactly **zero**, confirming that the strategy preserves 100% of the ground-truth FKs. We further observe that declared PKs often act as supersets of the actual referenced unique keys, driven by indexing choices or domain-specific modeling decisions. In contrast, true FKs frequently reference minimal unique subsets. By targeting MinUCCs rather than potentially missing or overly composite PKs, LLM-FK achieves more effective pruning while strictly preserving correctness.

Scenario	Model	TPC-H	NWind	TPC-E	AdvWorks	MusicBz
Missing Data	w/o MP-CoT	0.89	1.00	0.88	0.90	0.93
	Full MP-CoT	1.00	1.00	0.92	0.91	0.96
	<i>Improv. (Δ)</i>	+11%	–	+4%	+1%	+3%
Obscure Semantics	w/o MP-CoT	0.74	1.00	0.90	0.82	0.86
	Full MP-CoT	0.78	1.00	0.92	0.83	0.89
	<i>Improv. (Δ)</i>	+4%	–	+2%	+1%	+3%
Anonymous Schema	w/o MP-CoT	0.31	0.83	0.56	0.28	0.38
	Full MP-CoT	0.67	0.92	0.71	0.53	0.76
	<i>Improv. (Δ)</i>	+36%	+9%	+15%	+25%	+38%

Table 3: Robustness analysis of Multi-Perspective CoT Reasoning across different challenging scenarios.

4.4 Analysis of Domain Knowledge Injection

To quantify the contribution of Self-Augmented Domain Knowledge Injection, we selectively disable it in the Refiner and the Verifier. Table 2 reports the resulting F1-score changes of each agent’s output.

The results confirm that domain knowledge injection is beneficial across both agents. For the Refiner, removing injected knowledge leads to notable performance drops on AdventureWorks and MusicBrainz (9% and 11%, respectively). This degradation arises because isolated column pairs lack sufficient semantic context in large schemas, causing LLMs to misinterpret entity relationships. For the Verifier, domain knowledge is especially critical on smaller datasets such as TPC-H, where its removal results in a 22% F1-score drop. In compact schemas, global semantic cues enable explicit differentiation among competing candidates during conflict resolution. On TPC-E, which is dominated by abbreviated identifiers, injected domain knowledge consistently alleviates semantic ambiguity, preventing performance degradation for both agents. These findings highlight the complementary role of global semantic grounding in both local inference and global consistency enforcement.

4.5 Analysis of Multi-Perspective CoT

We evaluate the robustness of **Multi-Perspective CoT (MP-CoT) Reasoning** under three challenging real-world scenarios: (1) **Missing Data** (empty tables), (2) **Obscure Semantics** (Chinese pinyin abbreviations), and (3) **Anonymous Schema** (schema names unavailable). These scenarios simulate irregular conditions commonly encountered in partial or degraded database schemas. We compare the F1-scores of the full MP-CoT framework against a variant without MP-CoT in Table 3.

The results yield several insights. **First**, removing MP-CoT consistently degrades performance

Dataset	Refiner (Pre-CR)			Verifier (Post-CR)			Gains	
	P	R	F	P	R	F	ΔP	ΔF
TPC-H	0.64	1.00	0.78	1.00	1.00	1.00	+36%	+22%
NWind	0.85	1.00	0.92	1.00	1.00	1.00	+15%	+8%
TPC-E	0.56	0.96	0.70	0.91	0.96	0.93	+35%	+23%
AdvWorks	0.71	0.99	0.83	0.93	0.96	0.94	+22%	+11%
MusicBz	0.77	1.00	0.87	0.91	1.00	0.95	+14%	+8%

Table 4: Ablation study on the impact of Holistic Conflict Resolution.

across all datasets, validating its necessity. The effect is most pronounced under Anonymous Schema conditions, where MP-CoT improves F1-score by up to 38% on MusicBrainz. **Second**, MP-CoT provides strong robustness by allowing remaining perspectives to compensate when one signal is unavailable. Even in the presence of missing data, where heuristic-based methods typically fail, F1-scores remain above 90% across all datasets. **Finally**, the analysis further confirms that these reasoning perspectives are complementary and non-redundant, ensuring resilience against schema degradation by mutually compensating for limitations: syntactic and semantic clues address data sparsity, syntactic and statistical evidence mitigates semantic ambiguity, and statistical signals become pivotal when schema names are absent.

4.6 Analysis of Holistic Conflict Resolution

We evaluate the effectiveness of the Holistic Conflict Resolution (CR) Strategy by comparing Refiner outputs before conflict resolution (Pre-CR) with the final predictions produced by the Verifier (Post-CR). As shown in Table 4, the Refiner is intentionally designed to be aggressive, prioritizing high recall at the cost of lower precision.

The Verifier serves as a critical global filter that resolves structural inconsistencies. Across all datasets, conflict resolution substantially improves precision, with gains ranging from +14% to +36%, while largely preserving recall. In four out of five datasets, all true FKs identified by the Refiner are retained after CR. Notably, all cyclic references encountered are strictly length-2, ensuring that the resolution avoids the cascading complexity associated with global traversals. These results demonstrate that holistic, schema-level reasoning is essential for reconciling local predictions and ensuring global consistency, enabling LLM-FK to achieve both high accuracy and structural correctness.

5 Conclusion

We present **LLM-FK**, the first multi-agent framework for FK detection in large-scale complex databases. By orchestrating four specialized agents, LLM-FK decomposes the detection task, performs semantically grounded reasoning, and enforces global consistency. Specifically, the Profiler reduces the search space via Unique-Key-Driven Schema Decomposition Strategy; the Interpreter and Refiner leverage Self-Augmented Domain Knowledge Injection and Multi-Perspective CoT Reasoning for deep inference; and the Verifier reconciles local predictions through Holistic Conflict Resolution Strategy. Extensive experiments show that LLM-FK consistently achieves SOTA performance, providing an accurate, efficient, and robust solution for FK discovery that serves as a practical basis for downstream database applications.

Limitations

While our experiments cover five benchmarks spanning commerce, retail, finance, manufacturing, and entertainment, there remain unexplored database domains with extremely large-scale or highly heterogeneous schemas. Importantly, LLM-FK is designed to be domain-agnostic and relies solely on schema-inherent signals rather than domain-specific knowledge, ensuring that its performance does not depend on particular conventions or languages. Moreover, its multi-agent design allows seamless extension to new databases without additional training or supervision. Combined with our candidate pruning and structured reasoning strategies, LLM-FK maintains computational feasibility and robust generalization even in complex schemas, making it broadly applicable while leaving room for future evaluations on specialized domains.

Acknowledgments

This work was supported by the National Key Laboratory of Data Space Technology and System, Key Scientific Research Project of the Open University of China (Grant No. Z23C0007-1), China National Petroleum Corporation Science and Technology Project (Grant No. 2024ZZ46-06), and Scientific and Technological Innovation Platform Research Project of Guizhou Province (CXPTXM [2025] 003). This work has been implemented in the real-world production environment of Guizhou Xijiu Co., with valuable contributions from Zhenjun Liu, Shibin Yue, and Meili Zhang.

References

- Ziawasch Abedjan, Lukasz Golab, and Felix Naumann. 2015. Profiling relational data: a survey. *The VLDB Journal*, 24(4):557–581.
- Archita Agarwal, Marilyn George, Aaron Jeyaraj, and Malte Schwarzkopf. 2021. Retrofitting gdpr compliance onto legacy databases. *Proceedings of the VLDB Endowment*, 15(4).
- Alibaba. 2022. [Alibaba-Java-Coding-Guidelines: GitHub repository](#). GitHub Repository, accessed 2026-02-05.
- Michael Armbrust, Ali Ghodsi, Reynold Xin, Matei Zaharia, and 1 others. 2021. Lakehouse: a new generation of open platforms that unify data warehousing and advanced analytics. In *Proceedings of CIDR*, volume 8, page 28. sn.
- Simone Balloccu, Patrícia Schmidtová, Mateusz Lango, and Ondřej Dušek. 2024. Leak, cheat, repeat: Data contamination and evaluation malpractices in closed-source llms. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 67–93.
- Jana Bauckmann, Ulf Leser, Felix Naumann, and Véronique Tietz. 2006. Efficiently detecting inclusion dependencies. In *2007 IEEE 23rd international conference on data engineering*, pages 1448–1450. IEEE.
- Philip A Bernstein and Eric Newcomer. 2009. *Principles of transaction processing*. Morgan Kaufmann.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, and 1 others. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Shimin Chen, Anastasia Ailamaki, Manos Athanasoulis, Phillip B Gibbons, Ryan Johnson, Ippokratis Pandis, and Radu Stoica. 2011. Tpc-e vs. tpc-c: Characterizing the new tpc-e benchmark via an i/o comparison study. *ACM Sigmod Record*, 39(3):5–10.
- Si-An Chen, Lesly Miculicich, Julian M Eisenschlos, Zifeng Wang, Zilong Wang, Yanfei Chen, Yasuhisa Fujii, Hsuan-Tien Lin, Chen-Yu Lee, and Tomas Pfister. 2024. Tablerag: Million-token table understanding with language models. *Advances in Neural Information Processing Systems*, 37:74899–74921.
- Wenhu Chen. 2023. Large language models are few (1)-shot table reasoners. In *Findings of the association for computational linguistics: EACL 2023*, pages 1120–1130.
- Zhimin Chen, Vivek Narasayya, and Surajit Chaudhuri. 2014. Fast foreign-key detection in microsoft sql server powerpivot for excel. *Proceedings of the VLDB Endowment*, 7(13):1417–1428.
- Zhoujun Cheng, Tianbao Xie, Peng Shi, Chengzu Li, Rahul Nadkarni, Yushi Hu, Caiming Xiong, Dragomir Radev, Mari Ostendorf, Luke Zettlemoyer, and 1 others. 2022. Binding language models in symbolic languages. *arXiv preprint arXiv:2210.02875*.
- Cody James Christopher, Kristen Moore, and David Liebowitz. 2021. Schemadb: Structures in relational datasets. *arXiv preprint arXiv:2111.12835*.
- Fabien De Marchi, Stéphane Lopes, and Jean-Marc Petit. 2002. Efficient algorithms for mining inclusion dependencies. In *International Conference on Extending Database Technology*, pages 464–476. Springer.
- Till Döhmen, Radu Geacu, Madelon Hulsebos, and Sebastian Schelter. 2024. Schemapile: a large collection of relational database schemas. *Proceedings of the ACM on Management of Data*, 2(3):1–25.
- Falco Dürsch, Axel Stebner, Fabian Windheuser, Maxi Fischer, Tim Friedrich, Nils Strelow, Tobias Bleifuß, Hazar Harmouch, Lan Jiang, Thorsten Papenbrock, and 1 others. 2019. Inclusion dependency discovery: An experimental evaluation of thirteen algorithms. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*, pages 219–228.
- Raul Castro Fernandez, Ziawasch Abedjan, Famién Koko, Gina Yuan, Samuel Madden, and Michael Stonebraker. 2018. Aurum: A data discovery system. In *2018 IEEE 34th International conference on data engineering (ICDE)*, pages 1001–1012. IEEE.
- Raul Castro Fernandez, Aaron J Elmore, Michael J Franklin, Sanjay Krishnan, and Chenhao Tan. 2023. How large language models will disrupt data management. *Proceedings of the VLDB Endowment*, 16(11):3302–3309.
- Dawei Gao, Haibin Wang, Yaliang Li, Xiuyu Sun, Yichen Qian, Bolin Ding, and Jingren Zhou. 2023. Text-to-sql empowered by large language models: A benchmark evaluation. *arXiv preprint arXiv:2308.15363*.
- Cyril Goutte and Eric Gaussier. 2005. A probabilistic interpretation of precision, recall and f-score, with implication for evaluation. In *European conference on information retrieval*, pages 345–359. Springer.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Shuhao Guan and Derek Greene. 2024. Advancing post-ocr correction: A comparative study of synthetic data. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 6036–6047.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Peiyi Wang, Qihao Zhu, Runxin Xu, Ruoyu Zhang, Shirong Ma, Xiao Bi, and 1 others. 2025.

- Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V Chawla, Olaf Wiest, and Xiangliang Zhang. 2024. Large language model based multi-agents: A survey of progress and challenges. *arXiv preprint arXiv:2402.01680*.
- Rihan Hai, Christos Koutras, Christoph Quix, and Matthias Jarke. 2023. Data lakes: A survey of functions and systems. *IEEE Transactions on Knowledge and Data Engineering*, 35(12):12571–12590.
- Alon Halevy, Michael Franklin, and David Maier. 2006. Principles of dataspace systems. In *Proceedings of the twenty-fifth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 1–9.
- Jan L Harrington. 2016. *Relational database design and implementation*. Morgan Kaufmann.
- Gonzalo Jaimovitch-López, Cèsar Ferri, José Hernández-Orallo, Fernando Martínez-Plumed, and María José Ramírez-Quintana. 2023. Can language models automate data wrangling? *Machine Learning*, 112(6):2053–2082.
- Lan Jiang and Felix Naumann. 2020. Holistic primary key and foreign key detection. *Journal of Intelligent Information Systems*, 54(3):439–461.
- Theodore Johnson, Amit Marathe, and Tamraparni Dasu. 2003. Database exploration and bellman. *IEEE Data Eng. Bull.*, 26(3):34–39.
- Youri Kaminsky, Eduardo HM Pena, and Felix Naumann. 2023. Discovering similarity inclusion dependencies. *Proceedings of the ACM on Management of Data*, 1(1):1–24.
- 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*.
- Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. 2017. Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, 30.
- Robert E Kent. 2012. Database semantics. *arXiv preprint arXiv:1209.3054*.
- Georg Lausen. 2007. Relational databases in rdf: Keys and foreign keys. In *Workshop on Ontologies-based techniques for DataBases and Information Systems*, pages 43–56. Springer.
- Jinyang Li, Binyuan Hui, Ge Qu, Jiayi Yang, Binhua Li, Bowen Li, Bailin Wang, Bowen Qin, Ruiying Geng, Nan Huo, and 1 others. 2023. Can llm already serve as a database interface? a big bench for large-scale database grounded text-to-sqls. *Advances in Neural Information Processing Systems*, 36:42330–42357.
- Yiming Lin, Yeye He, and Surajit Chaudhuri. 2023. Auto-bi: Automatically build bi-models leveraging local join prediction and global schema graph. *arXiv preprint arXiv:2306.12515*.
- Jim Melton and Alan R Simon. 2001. *SQL: 1999: understanding relational language components*. Elsevier.
- Microsoft. 2023a. [AdventureWorks Sample Database: GitHub repository](#). GitHub Repository, accessed 2025-03-24.
- Microsoft. 2023b. [Northwind Sample Database: GitHub repository](#). GitHub Repository, accessed 2025-03-24.
- Avanika Narayan, Ines Chami, Laurel Orr, Simran Arora, and Christopher Ré. 2022. Can foundation models wrangle your data? *arXiv preprint arXiv:2205.09911*.
- Sam Newman. 2021. *Building microservices: designing fine-grained systems*. " O'Reilly Media, Inc."
- Thorsten Papenbrock, Tanja Bergmann, Moritz Finke, Jakob Zwiener, and Felix Naumann. 2015a. Data profiling with metanome. *Proceedings of the VLDB Endowment*, 8(12):1860–1863.
- Thorsten Papenbrock, Sebastian Kruse, Jorge-Arnulfo Quiané-Ruiz, and Felix Naumann. 2015b. Divide & conquer-based inclusion dependency discovery. *Proceedings of the VLDB Endowment*, 8(7):774–785.
- Meikel Poesch and Chris Floyd. 2000. New tpc benchmarks for decision support and web commerce. *ACM Sigmod Record*, 29(4):64–71.
- Alexandra Rostin, Oliver Albrecht, Jana Bauckmann, Felix Naumann, and Ulf Leser. 2009. A machine learning approach to foreign key discovery. In *WebDB*.
- Aaditya Singh, Adam Fry, Adam Perelman, Adam Tart, Adi Ganesh, Ahmed El-Kishky, Aidan McLaughlin, Aiden Low, AJ Ostrow, Akhila Ananthram, and 1 others. 2025. Openai gpt-5 system card. *arXiv preprint arXiv:2601.03267*.
- Songyuan Sui, Hongyi Liu, Serena Liu, Li Li, Soo-Hyun Choi, Rui Chen, and Xia Hu. 2025. Chain-of-query: Unleashing the power of llms in sql-aided table understanding via multi-agent collaboration. In *Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*, pages 957–986.
- Yuan Sui, Mengyu Zhou, Mingjie Zhou, Shi Han, and Dongmei Zhang. 2024a. Table meets llm: Can large language models understand structured table data? a benchmark and empirical study. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, pages 645–654.

- Yuan Sui, Jiaru Zou, Mengyu Zhou, Xinyi He, Lun Du, Shi Han, and Dongmei Zhang. 2024b. Tap4llm: Table provider on sampling, augmenting, and packing semi-structured data for large language model reasoning. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 10306–10323.
- Aaron Swartz. 2002. Musicbrainz: A semantic web service. *IEEE Intelligent Systems*, 17(1):76–77.
- Fabian Tschirschnitz, Thorsten Papenbrock, and Felix Naumann. 2017. Detecting inclusion dependencies on very many tables. *ACM Transactions on Database Systems (TODS)*, 42(3):1–29.
- Bing Wang, Changyu Ren, Jian Yang, Xinnian Liang, Jiaqi Bai, Linzheng Chai, Zhao Yan, Qian-Wen Zhang, Di Yin, Xing Sun, and 1 others. 2025a. Mac-sql: A multi-agent collaborative framework for text-to-sql. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 540–557.
- Xixi Wang, Miguel Costa, Jordanka Kovaceva, Shuai Wang, and Francisco C Pereira. 2025b. Plugging schema graph into multi-table qa: A human-guided framework for reducing llm reliance. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 5829–5842.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, and 1 others. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.
- Jian Wu, Linyi Yang, Dongyuan Li, Yuliang Ji, Manabu Okumura, and Yue Zhang. 2025. Mmqa: Evaluating llms with multi-table multi-hop complex questions. In *The thirteenth international conference on learning representations*.
- Xiaoyu Wu, Ning Wang, and Huaxi Liu. 2019. Discovering foreign keys on web tables with the crowd. *Computing and Informatics*, 38(3):621–646.
- Wenxuan Xie, Gaochen Wu, and Bowen Zhou. 2024. Mag-sql: Multi-agent generative approach with soft schema linking and iterative sub-sql refinement for text-to-sql. *arXiv preprint arXiv:2408.07930*.
- Cong Yan and Yeye He. 2020. Auto-suggest: Learning-to-recommend data preparation steps using data science notebooks. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, pages 1539–1554.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Yunhu Ye, Binyuan Hui, Min Yang, Binhua Li, Fei Huang, and Yongbin Li. 2023. Large language models are versatile decomposers: Decomposing evidence and questions for table-based reasoning. In *Proceedings of the 46th international ACM SIGIR conference on research and development in information retrieval*, pages 174–184.
- 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 (Volume 1: Long Papers)*, pages 17432–17451.
- Bin Zhang, Hangyu Mao, Lijuan Li, Zhiwei Xu, Dapeng Li, Rui Zhao, and Guoliang Fan. 2024. Sequential asynchronous action coordination in multi-agent systems: A stackelberg decision transformer approach. In *Forty-first International Conference on Machine Learning*.
- Meihui Zhang, Marios Hadjieleftheriou, Beng Chin Ooi, Cecilia M Procopiuc, and Divesh Srivastava. 2010. On multi-column foreign key discovery. *Proceedings of the VLDB Endowment*, 3(1-2):805–814.
- Wei Zhou, Mohsen Mesgar, Annemarie Friedrich, and Heike Adel. 2025. Efficient multi-agent collaboration with tool use for online planning in complex table question answering. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 945–968.
- Jun-Peng Zhu, Peng Cai, Kai Xu, Li Li, Yishen Sun, Shuai Zhou, Haihuang Su, Liu Tang, and Qi Liu. 2024. Autotqa: Towards autonomous tabular question answering through multi-agent large language models. *Proceedings of the VLDB Endowment*, 17(12):3920–3933.

A More Related Work

A.1 Metadata Discovery

Extensive research has investigated efficient strategies for discovering Unique Column Combinations (UCCs) and Inclusion Dependencies (INDs), making the extraction of such metadata computationally feasible even for large databases (De Marchi et al., 2002; Abedjan et al., 2015; Dürsch et al., 2019). Representative algorithms include **Spider** (Bauckmann et al., 2006), which efficiently identifies single-column INDs using a disk-based sort-merge join, and **Binder** (Papenbrock et al., 2015b), which generalizes the approach to multi-column INDs via a hash-based divide-and-conquer strategy.

To address practical challenges such as dirty or incomplete data, approximate IND discovery techniques have also been proposed (Tschirschnitz et al., 2017; Kaminsky et al., 2023). The metadata

discovery process within LLM-FK is fully compatible with these established algorithms, allowing practitioners to tailor the framework according to specific objectives—such as prioritizing computational efficiency or maximizing robustness to data anomalies—without modifying the overall reasoning pipeline.

A.2 Multi-Agent Systems

Multi-agent systems have recently shown strong potential for solving complex reasoning tasks by enabling collaboration among specialized agents (Guan and Greene, 2024; Guo et al., 2024; Zhang et al., 2024). In the tabular domain, existing multi-agent frameworks primarily target downstream applications like Text-to-SQL and table question answering (Xie et al., 2024; Sui et al., 2025; Wang et al., 2025a; Yu et al., 2025; Zhou et al., 2025). These tasks generally assume a known schema and rely on content matching and local schema linking to satisfy specific queries.

In contrast, FK detection presents a fundamentally different challenge: it requires holistic structural reasoning across an entire database schema. To the best of our knowledge, LLM-FK is the first framework to leverage multi-agent collaboration specifically for FK detection, demonstrating how specialized agents can jointly address the combinatorial and inconsistency challenges inherent in large, complex databases. This design illustrates the potential of multi-agent reasoning beyond content-oriented tasks, extending it to structural inference and schema-level constraint discovery.

A.3 FK Detection in Modern Data Ecosystems

While FKs are fundamental to relational database theory, their explicit definition is frequently absent in practice, making FK detection a critical task for downstream data operations. The urgency and practical significance of this problem in contemporary settings can be summarized across two key dimensions:

Ubiquity of Missing FKs. In modern data environments, the absence of explicit FKs has become a widespread norm driven by evolving engineering practices and architectural shifts. To avoid performance degradation during high-concurrency operations, developers frequently omit physical FKs, shifting constraint management entirely to the application layer (Alibaba, 2022), making these implicit constraints highly prone to loss during

database migration or analytical extraction (Jiang and Naumann, 2020). Furthermore, distributed architectures—such as the “database-per-service” pattern in microservices (Newman, 2021) and schema-on-read data lakehouses (Armbrust et al., 2021)—inherently restrict strict, cross-database FK constraints. Empirical studies corroborate this scarcity, revealing that a significant majority of real-world open-source web applications and relational databases completely lack explicit FK definitions (Agarwal et al., 2021; Christopher et al., 2021; Döhmen et al., 2024).

Critical Role in Advanced Data Applications.

Beyond classic data management tasks like data integration and schema reverse engineering, FK detection is increasingly pivotal in contemporary data ecosystems. In complex, heterogeneous environments, identifying FK relationships is highly valuable for streamlining data discovery and bolstering governance efficiency (Fernandez et al., 2018; Hai et al., 2023). Moreover, accurately inferring these structural dependencies provides an essential foundation for empowering automated Business Intelligence (BI) modeling (Lin et al., 2023), as FK detection closely aligns with the foundational steps of BI model construction and enables automatic generation of schema graphs for downstream analytical pipelines. In the current era of LLM-driven database interfaces, explicit FKs provide deterministic join paths that have been empirically proven to substantially enhance the accuracy of Text-to-SQL parsing (Gao et al., 2023). More broadly, multi-table Question Answering (QA) has recently emerged as a focal point in LLM-driven table understanding (Zhu et al., 2024; Wang et al., 2025b), where FK detection is explicitly identified as a prerequisite for resolving multi-hop complex questions across diverse tables (Wu et al., 2025). This positions robust FK detection as a key enabler for seamlessly integrating structured relational semantics into emerging LLM-based table reasoning systems.

B More Implementation Details

B.1 Detailed Elaboration of Multi-Agent Collaboration

To automate the reasoning required for FK detection, LLM-FK orchestrates four specialized agents—**Profiler**, **Interpreter**, **Refiner**, and **Verifier**—to collaboratively execute a database-wide FK detection task. Each agent executes distinct

cognitive operations via specific LLM instructions \mathcal{I} , denoted as π . Formally, the roles and operational formulations are as follows:

Profiler (\mathcal{A}_p). The Profiler serves as the entry point of the agent workflow, implementing a **Unique-Key-Driven Schema Decomposition** to mitigate search space explosion. Given a database \mathcal{D} , the set of identified INDs \mathcal{J} (pre-filtered by data type rules), and the set of MinUCCs \mathcal{U} derived from referenced tables in \mathcal{J} , the Profiler iterates over each referenced table T . Using its subset of MinUCCs $\mathcal{U}_T \subseteq \mathcal{U}$, it selects the most suitable referenced unique key k guided by $\mathcal{I}_{\mathcal{A}_p}$:

$$k = \pi(T, \mathcal{U}_T, \mathcal{I}_{\mathcal{A}_p}), \quad (1)$$

These keys are aggregated into a set \mathcal{K} , which deterministically prunes the search space. The structurally valid candidate FK pairs are then:

$$\Psi_{cand} = \{(c_f, c_p) \in \mathcal{J} \mid c_p \in \mathcal{K}\}, \quad (2)$$

Interpreter (\mathcal{A}_i). The Interpreter performs **Self-Augmented Domain Knowledge Injection** to reduce schema ambiguities. Analyzing Ψ_{cand} , it derives global domain knowledge \mathcal{K}_{dom} from the set of table names $\mathcal{N}_{\Psi_{cand}}$:

$$\mathcal{K}_{dom} = \pi(\mathcal{N}_{\Psi_{cand}}, \mathcal{I}_{\mathcal{A}_i}), \quad (3)$$

which is shared downstream to ensure semantic consistency across local reasoning tasks.

Refiner (\mathcal{A}_r). The Refiner validates each candidate pair $\psi = (c_f, c_p) \in \Psi_{cand}$ using **Multi-Perspective CoT Reasoning**, integrating domain knowledge \mathcal{K}_{dom} and serialized structural representations S_ψ derived from SQL queries:

$$p_\psi = \pi(\psi, S_\psi, \mathcal{K}_{dom}, \mathcal{I}_{\mathcal{A}_r}), \quad (4)$$

where $p_\psi \in \{0, 1\}$ indicates whether ψ constitutes a true FK. We denote the set of positively validated candidates as:

$$\Psi_{pos} = \{\psi \in \Psi_{cand} \mid p_\psi = 1\}, \quad (5)$$

Verifier (\mathcal{A}_v). The Verifier ensures global coherence via the **Holistic Conflict Resolution Strategy**. Initialized with Ψ_{pos} , it identifies structural violations (e.g., multiple or cyclic references) within the aggregated schema topology. For any identified conflict set $\Psi_{conflict}$, it determines the valid subset

$\Psi_{resolved}$ by leveraging domain knowledge \mathcal{K}_{dom} and structural representations $S_{\Psi_{conflict}}$:

$$\Psi_{resolved} = \pi(\Psi_{conflict}, S_{\Psi_{conflict}}, \mathcal{K}_{dom}, \mathcal{I}_{\mathcal{A}_v}). \quad (6)$$

This resolution mechanism refines the global schema topology by eliminating invalid candidate FKs, ensuring convergence to a conflict-free state and yielding the final consistent FK set Φ .

Multi-Agent Collaboration. The overall workflow is formally defined in Algorithm 2. Detailed prompts for each agent are provided in Appendix F.

B.2 Detailed Elaboration of Profiler

B.2.1 IND-based and Rule-based Pruning

To effectively reduce the search space, the Profiler enforces two fundamental rules for pruning:

- **Type Compatibility:** Candidate columns must strictly share the same data type.
- **Exclusion of Unsuitable Types:** Columns representing *continuous measurements* (e.g., float, decimal) or *binary states* (e.g., bool, bit) are automatically filtered out as they semantically lack identifying properties.

Beyond these foundational constraints, our IND-based and rule-based mechanisms are designed with significant flexibility and extensibility. The IND-based validation supports adaptive algorithm selection (see Appendix A.1) to optimize the balance between execution efficiency and robustness. Furthermore, additional rules—such as filtering empty tables or columns and applying specific cardinality thresholds—can be seamlessly integrated into the rule-based pruning.

B.2.2 Unique Key Selection

Optimal unique key selection leverages LLM prompts guided by formalized judgment criteria (see Table 5), including uniqueness validity, column ordinality, naming conventions, type suitability, value conciseness, minimal composition, and semantic alignment. By grounding the selection process in these formalized rules, the LLM can effectively discern the most appropriate referenced unique key from the set of MinUCCs. This approach leverages the model’s reasoning capabilities to prioritize candidates that not only satisfy structural constraints but also align with the logical conventions of the database schema, ensuring the selected keys are semantically meaningful.

Criterion	Description
Uniqueness Validity	The candidate must uniquely identify a specific tuple within the table.
Column Ordinality	Columns positioned earlier in the table definition are prioritized.
Naming Convention	Preference is given to column names containing typical identifiers (e.g., “id”, “key”, “code”).
Type Suitability	The data type is typically an integer or a string, adhering to standard key definitions.
Value Conciseness	The data values should be concise and human-readable (short text length).
Minimal Composition	Candidates with fewer constituent columns (atomic keys) are preferred over complex composite keys.
Semantic Alignment	The key must logically represent the entity according to the semantic context of the table.

Table 5: Judgment criteria for unique key selection.

Type	Information
Schema Definitions	Table name Column name
Statistical Metadata	Ordinal position Data type Average value text length Number of distinct values Number of rows in table Cardinality ratio (distinct values / table rows) Minimum value Maximum value
Inter-column Dependencies	Coverage ratio Table size ratio Out-of-range ratio
Sample Data	Example data from table (first 5 rows)

Table 6: Serialized structural representations for Multi-Perspective CoT Reasoning.

B.3 Detailed Elaboration of Interpreter

The Interpreter implements a summarize-then-reason strategy (see Figure 4), which mirrors the human cognitive process of acquiring a holistic understanding of the database domain before making specific judgments. In this workflow, the LLM first synthesizes a global domain profile from collective table semantics to prime itself with a unified context. Crucially, these self-generated insights are injected back into local reasoning tasks, allowing the framework to self-augment abstract schema symbols with concrete domain knowledge. This global-to-local paradigm ensures semantic consistency by grounding isolated inferences in a shared

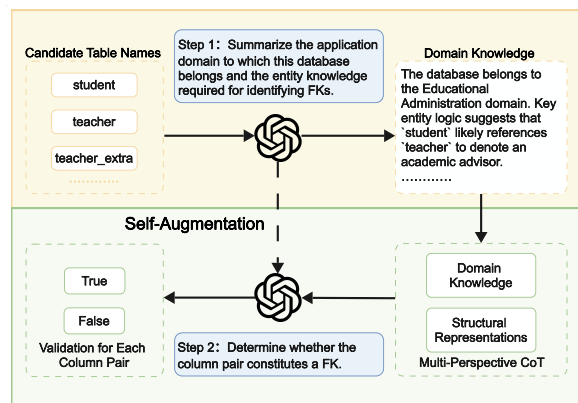


Figure 4: Illustration of the Self-Augmented Domain Knowledge Injection.

foundation, while simultaneously achieving full automation—autonomously deriving macroscopic insights without reliance on external manual annotations.

B.4 Detailed Elaboration of Refiner

Structural Representations Extraction. The Refiner constructs serialized structural representations by formulating SQL queries on raw table data (see Table 6), including schema definitions, statistical metadata, inter-column dependencies, and sample data. This ensures that multi-perspective reasoning is grounded in precise schema and data evidence.

Multi-Perspective CoT Reasoning. The Refiner integrates three reasoning perspectives (see Figure 5):

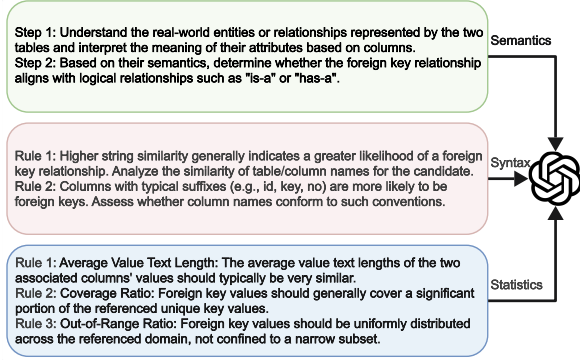


Figure 5: Illustration of the Multi-Perspective CoT Reasoning.

- **Semantic:** Maps schema concepts to real-world entities.
- **Syntactic:** Evaluates naming conventions and string similarity.
- **Statistical:** Analyzes statistical distributions and inclusion relationships.

B.5 Detailed Elaboration of Verifier

The Verifier enforces global consistency via the **Holistic Conflict Resolution** algorithm (see Algorithm 1), resolving multiple references and cyclic references to produce the final, structurally consistent FK set Φ .

C Dataset Details

Table 7 provides a detailed statistical overview of the datasets used in our experiments. These datasets span diverse domains, schema complexities, and scales, enabling a comprehensive evaluation of FK detection methods.

- **TPC-H:** A standard decision-support benchmark widely adopted in the database community. It simulates a commercial environment with a relatively simple schema (8 tables) and sparse FK relationships, providing a baseline for correctness verification.
- **Northwind:** A classic synthetic dataset representing a small retail business. With 13 tables, it offers a manageable yet representative schema for evaluating FK detection in typical inventory and sales scenarios.
- **TPC-E:** An On-Line Transaction Processing (OLTP) (Bernstein and Newcomer, 2009) benchmark modeling a brokerage firm. Its

Algorithm 1 Holistic Conflict Resolution

Require: Candidate FK Set Ψ_{pos}

Ensure: Consistent FK Set Φ

- 1: **Initialize** $\Phi \leftarrow \Psi_{pos}$
 - 2: **// Stage 1: Resolve Multiple References**
 - 3: **// Ensure each source column references at most one target column**
 - 4: **for** each source column c_{src} referencing distinct target columns $\mathcal{C}_{targets}$ in Φ **do**
 - 5: **if** $|\mathcal{C}_{targets}| > 1$ **then**
 - 6: **// Select the single most plausible target column via reasoning**
 - 7: $c_{best} \leftarrow \text{Reasoning-Select}(c_{src}, \mathcal{C}_{targets})$
 - 8: Update Φ to retain only the reference (c_{src}, c_{best}) for c_{src}
 - 9: **end if**
 - 10: **end for**
 - 11: **// Stage 2: Resolve Cyclic References**
 - 12: **// Iteratively eliminate cycles to ensure topological consistency**
 - 13: **while** Schema Topology Graph $G(\Phi)$ contains cycles **do**
 - 14: **// Prioritize immediate violations for unambiguous resolution**
 - 15: $Cycle \leftarrow \text{FindShortestCycle}(G)$
 - 16: **// Identify the least plausible link (edge) in the cycle**
 - 17: $e_{weak} \leftarrow \text{Reasoning-Weakest}(Cycle)$
 - 18: $\Phi \leftarrow \Phi \setminus \{e_{weak}\}$ {Break the cycle}
 - 19: **end while**
 - 20: **return** Φ
-

schema is more complex than TPC-H (33 tables) and includes richer integrity constraints. English abbreviations are frequently used in table and column names, introducing additional semantic challenges.

- **AdventureWorks:** A synthetic database simulating a bicycle manufacturing company. Containing 71 tables and numerous FKs, it exhibits complex schema patterns, including composite FKs, which pose substantial challenges for FK detection algorithms.
- **MusicBrainz:** A large-scale, real-world music encyclopedia database. With 382 tables and over 2.4 million column pairs, it is the most complex dataset in our evaluation, providing a testbed for both scalability and robustness in a noisy, heterogeneous environment.

Dataset	Type	Tables	FKs	Column Pairs	Domain
TPC-H	Benchmark	8	9	3,721	Commerce
Northwind	Synthetic	13	11	7,744	Retail
TPC-E	Benchmark	33	46	36,481	Finance
AdventureWorks	Synthetic	71	91	236,196	Manufacturing
MusicBrainz	Real	382	354	2,458,624	Entertainment

Table 7: Detailed statistics and domains of datasets.

D Reproducibility Statement

D.1 LLM-FK and LLM-based Baselines

All experiments are conducted using **DeepSeek-R1** as the backbone LLM. To ensure deterministic results and reproducibility, we set the temperature parameter to 0, while leaving all other hyperparameters at their default values.

For metadata extraction, we employ the **Metanome** profiling tool (Papenbrock et al., 2015a) to identify structural constraints; specifically, we target *MinUCCs* and utilize the *Spider* algorithm (Bauckmann et al., 2006) to detect single-column *INDs*.

For LLM-based baselines: **Chorus** is implemented using its official open-source code. Detailed prompt designs for other LLM-based baselines, as well as for **LLM-FK**, are provided in Appendix F and Appendix G.

D.2 Heuristic-based Baselines

All heuristic-based baselines are faithfully re-implemented following the methodology described in their respective papers, with careful attention to hyperparameters and evaluation protocols:

- **MC-FK**: Since no default parameters are provided, we adopt the reported confidence scores for FK detection and empirically tune the decision threshold to maximize the F1-score on each dataset.
- **Fast-FK**: Lacking explicit defaults, we assign reasonable weights to its constituent rules based on prior literature and preliminary tuning: column name similarity (0.2), column-table name similarity (0.5), cardinality ratio (0.3), and whether the referencing column is a unique key (−0.5).
- **HOPF**: Implemented strictly following the default configuration from the original work.

- **ML-FK** and **AutoSuggest**: We follow a leave-one-database-out training protocol proposed in the **ML-FK**. Specifically, for each target database, the labeled FK candidates from the remaining four datasets serve as training data. **LightGBM** (Ke et al., 2017) is employed as the classifier due to its efficiency and strong performance, with default hyperparameters used consistently across all experiments.

E More Experiment Results and Analysis

E.1 Generalizability to Different LLM Backbones

To demonstrate that LLM-FK’s effectiveness does not rely on a particular LLM backbone, we evaluated our framework using DeepSeek-R1 (Guo et al., 2025), Qwen3-Max (Yang et al., 2025), GPT-5-mini (Singh et al., 2025), and LLaMA3.3-70B-Instruct (Grattafiori et al., 2024) without any model-specific fine-tuning.

Generalizability of Candidate Pruning. As shown in Table 8, the Profiler consistently achieves a substantial and reliable reduction in the candidate search space across all LLMs (e.g., from over 2.4 million to ~2,355 pairs on MusicBrainz). Crucially, this reduction consistently preserves true positive matches, with only a single true FK being erroneously pruned across all trials—specifically by LLaMA3.3 on AdventureWorks. This highlights the high reliability and model-agnostic generalizability of our pruning methodology.

Generalizability of FK Detection. Table 9 reports the F1-scores for FK detection. LLM-FK significantly outperforms the overall best-performing End-to-End baselines specifically on complex schemas such as TPC-E, AdventureWorks, and MusicBrainz, demonstrating the generalizability of our multi-agent collaborative approach across different underlying LLMs. Among the evaluated

models, LLaMA 3.3 exhibits relatively lower performance, likely attributable to differences in inherent reasoning capabilities for complex logical constraints. Nevertheless, these results still consistently surpass the vast majority of traditional heuristic-based methods (see Table 1) and End-to-End baselines, further supporting the effectiveness of our LLM-FK. Finally, the marginally lower F1-scores observed on smaller datasets like TPC-H are an artifact of the constrained schema size; in practice, only a few FKs was misclassified, which disproportionately skews the metric rather than indicating a fundamental limitation.

E.2 Computational Cost Analysis

Following the preliminary search space analysis in Section 4.3, we detail the runtime overhead and computational expenditures associated with LLM calls. We conduct a comprehensive analysis of API calls, token usage, and monetary costs using Qwen3-Max as the backend LLM. The pricing is based on Qwen3-Max, with a rate of 2.5 CNY for inputs up to 32K tokens and 10 CNY for outputs up to 32K tokens.

Overall Comparison. To evaluate the LLM invocation frequency and associated expenses, we compare our method (LLM-FK) against an End-to-End baseline. This baseline identifies FKs between a pair of tables in a single pass and serves as the most effective and efficient baseline in our study. The overall comparison is summarized in Table 11. For smaller schemas such as TPC-H, LLM-FK exhibits a slightly lower F1-score (0.78 compared to 0.90) and incurs marginally higher costs. However, this performance gap corresponds to merely a few misclassified FKs, which is magnified by the small scale of the dataset. For larger datasets, including TPC-E and AdventureWorks, LLM-FK requires higher API costs and token consumption. Nevertheless, this increased expenditure represents a favorable trade-off, as it yields substantial improvements in accuracy and better performance in terms of API calls. Most notably, when applied to massive and complex schemas like MusicBrainz, LLM-FK not only achieves the highest accuracy (0.96) but also drastically reduces monetary costs, demonstrating its superior scalability and efficiency in large-scale environments.

Internal Breakdown. To further elucidate the sources of the computational overhead, we provide a detailed breakdown across the internal agents

within our multi-agent framework: Profiler, Interpreter, Refiner, and Verifier. As detailed in Table 12, the analysis clearly indicates that the primary time and cost expenditures are concentrated within the Refiner module. Conversely, the Verifier incurs minimal overhead because its input is drastically pruned by the Refiner, such as reducing 2,355 candidates down to merely 92 conflicts in the MusicBrainz dataset. Furthermore, these remaining conflicts naturally form tightly clustered, localized semantic subgraphs; notably, all cyclic references encountered across our experiments are strictly length-2 (see Appendix E.7 for a typical example). By iteratively prioritizing the resolution of these shortest cycles, the Verifier processes highly isolated, immediate neighborhood conflicts rather than complex, cascading chains, thereby requiring significantly fewer computational overhead to globally resolve.

E.3 Robustness in the Absence of UCC Metadata

To evaluate the pruning effectiveness of our method under empty-table conditions (where UCCs are unavailable), we conducted an experiment to verify its robustness. As summarized in Table 10, the results demonstrate consistently excellent pruning performance across all five benchmark datasets. The volume of retained candidates remains highly comparable to that of the original data-rich environment, and only a single true FK (in TPC-E) was erroneously pruned. This high consistency is achieved because, as illustrated in Table 5, our method leverages LLMs to comprehensively reason across multiple dimensions—such as column ordinality, naming conventions, and semantic alignment. This multi-dimensional reasoning allows our framework to infer plausible referenced unique keys directly from the schema, effectively compensating for the lack of UCCs and ensuring robustness even in completely empty-table environments. Furthermore, it is worth noting that in practical database operations, completely empty tables are often temporary or staging tables that rarely require complex structural analysis.

E.4 Example Analysis of Unique-Key-Driven Schema Decomposition

Column-Level Processing vs. Table-Level Processing. Prior methods input entire table pairs to LLMs, requesting all FKs in a single pass. This often leads to attention dispersion and hallucinations,

Dataset	Original Pairs	DeepSeek-R1		Qwen3-Max		LLaMA3.3		GPT-5-mini	
		Cand.	Err.	Cand.	Err.	Cand.	Err.	Cand.	Err.
TPC-H	3,721	49	0	49	0	47	0	49	0
Northwind	7,744	60	0	60	0	60	0	60	0
TPC-E	36,481	194	0	194	0	195	0	192	0
AdventureWorks	236,196	2,034	0	2,057	0	1,987	1	2,024	0
MusicBrainz	2,458,624	2,306	0	2,355	0	2,355	0	2,355	0

Table 8: The Profiler candidate pruning results across different LLMs. Each entry shows the number of candidates (Cand.) and the number of erroneous prunings (Err.), where erroneous pruning indicates FKs that are incorrectly removed during the pruning process.

LLM	Method	TPC-H	Northwind	TPC-E	AdventureWorks	MusicBrainz
DeepSeek-R1	End-to-End	0.90	1.00	0.69	0.73	0.75
	LLM-FK (Ours)	1.00	1.00	0.93	0.94	0.95
Qwen3-Max	End-to-End	0.90	0.92	0.58	0.52	0.75
	LLM-FK (Ours)	0.78	1.00	0.92	0.92	0.96
LLaMA3.3	End-to-End	0.80	0.95	0.61	0.66	0.74
	LLM-FK (Ours)	0.70	0.96	0.85	0.79	0.86
GPT-5-mini	End-to-End	0.90	0.95	0.67	0.79	0.80
	LLM-FK (Ours)	0.78	1.00	0.93	0.95	0.93

Table 9: F1-score comparison for FK detection across different LLMs. Best results per LLM are bolded.

Dataset	Pairs	Original		Empty	
		Cand.	Err.	Cand.	Err.
TPC-H	3,721	49	0	49	0
Northwind	7,744	60	0	69	0
TPC-E	36,481	194	0	156	1
AdventureWorks	236,196	2,034	0	2,127	0
MusicBrainz	2,458,624	2,306	0	2,138	0

Table 10: Pruning results in original versus empty tables. Cand. denotes retained candidates; Err. denotes mistakenly pruned true FKs.

particularly when multiple FKs exist between the same tables.

For example, in TPC-E, the table `holding_history` references `trade` through two distinct FKs: `hh_h_t_id` \rightarrow `t_id` and `hh_t_id` \rightarrow `t_id`. In our experiments, table-level methods typically detect only one FK. In contrast, even without Self-Augmented Domain Knowledge Injection or Multi-Perspective CoT Reasoning, LLM-FK successfully identifies both FKs, attributed to our column-pair decomposition. By decomposing schemas into atomic units, we reduce task complexity invariant to schema size, allowing focused reasoning for each candidate pair and maintaining tractability in large schemas.

Unique Key-Based vs. Primary Key-Based Pruning. Our unique key-based pruning

is more robust than relying solely on declared PKs. In complex schemas, PKs often form supersets of MinUCCs, introducing unnecessary redundancy. For instance, in `production_productproductphoto` (AdventureWorks), the PK is composite (`ProductID` + `ProductPhotoID`), whereas `ProductID` alone constitutes a MinUCC and serves as the target for all referencing FKs. If pruning relied on the overly composite PK, irrelevant candidate pairs involving `ProductPhotoID` would persist, increasing computational cost. In contrast, unique key-based pruning efficiently narrows the search space while preserving FK integrity.

E.5 Example Analysis of Self-Augmented Domain Knowledge Injection

We analyze the `holding_summary.hs_s_symb` and `daily_market.dm_s_symb` candidate column pair from TPC-E. Both columns share the suffix `s_symb` (Security Symbol) but serve distinct roles: `daily_market` tracks daily stock data, whereas `holding_summary` aggregates customer holdings.

Without domain knowledge, the LLM fails to decipher the semantics of these abbreviations, resulting in an erroneous reliance on syntactic similarity and mistakenly predicting an FK. With our Interpreter, a global domain context identifies the database as a “Financial Trading System” and dis-

tinguishes “Market Data” from “Account Holdings.” This knowledge guides the Refiner to reject the candidate, illustrating that self-augmented domain knowledge resolves semantic ambiguities effectively.

E.6 Example Analysis of Multi-Perspective CoT Reasoning

Beyond ensuring the robustness of LLM-FK in scenarios with missing data or ambiguous semantics, the Multi-Perspective CoT Reasoning leverages sufficient information embedded in structural representations and performs deep semantic reasoning to address ambiguous or inconsistent signals. For instance, the FK `series_annotation.annotation` → `annotation.id` in MusicBrainz involves columns with numerical values in different formats (scientific vs. integer), where the former is formalized as scientific notation due to the large magnitude of values in the leading rows, while the latter remains in integer format. Without multi-perspective reasoning, the LLM misinterprets scientific notation as string values, rejecting the FK due to mismatches in both data type and value range within the leading rows. LLM-FK overcomes this by:

- Injecting statistical metadata (data types, value ranges) to provide sufficient information, preventing misinterpretation.
- Cross-validating syntactic, semantic and statistical perspectives to ensure robust inference transcending superficial judgments based solely on format discrepancies or value range mismatches.

E.7 Example Analysis of Holistic Conflict Resolution

In the MusicBrainz, complex conflicts arise involving `l_artist_recording`, `artist`, and `artist_meta` within the outputs of the Refiner. Specifically, `l_artist_recording.entity0` is ambiguously identified as an FK referencing both `artist.id` and `artist_meta.id`. Exacerbating this issue, a direct cyclic reference exists between `artist` and `artist_meta` (i.e., they reference each other), which obscures the true hierarchical relationship. Such structural conflicts are highly typical rather than random occurrences. Crucially, the Refiner’s semantic-based reasoning constrains these highly correlated candidates into tightly clustered, semantically related subgraphs.

By constructing the global schema topology and collectively analyzing the candidates within the conflicting subgraphs, our Holistic Conflict Resolution Strategy identifies that `artist` serves as the semantic core entity, whereas `artist_meta` acts as auxiliary metadata. Consequently, we resolve the conflicts by retaining the canonical references to `artist.id` and eliminating the spurious candidates referencing `artist_meta.id`. This decision improves upon local validation methods that rely on assigning confidence scores to column pairs. While such scores are often inaccurate due to isolated reasoning, our construction of a global schema topology facilitates a direct semantic comparison to capture hierarchical relationships. This holistic reasoning process eliminates conflicting references and ensures that the final alignment is both semantically and structurally consistent at the global schema level.

F Prompts of LLM-FK and Case Study

In this appendix, we provide detailed examples of the prompts used for each core agent in LLM-FK: Profiler, Interpreter, Refiner, and Verifier. These examples illustrate how the multi-agent framework operationalizes FK detection.

We use the sample database introduced in Figure 2 as the running case study, which also illustrates the overall multi-agent workflow. Figure 1 illustrates the limitations of heuristic-based methods and the effectiveness of LLMs in semantic reasoning for FK detection, while also highlighting the accompanying challenges. Figure 4 specifically details the Self-Augmented Domain Knowledge Injection performed by the Interpreter.

Figures 6–10 show the specific prompts and their corresponding outputs across different stages of the pipeline. Specifically, these example prompts are directly populated with the structural representations, representing the execution outcomes of the SQL queries.

G Prompts of LLM-based Baselines

For completeness, we also provide the prompt designs used for the LLM-based baselines: End-to-End, Few-Shot and CoT. Figures 11–13 illustrate the corresponding prompts.

Dataset	Method	F1-score	API Calls	Tokens		Cost (¥)		
				In	Out	In	Out	Total
TPC-H	End-to-End	0.90	36	0.0351	0.0006	0.088	0.006	0.094
	LLM-FK	0.78	60	0.1221	0.0011	0.305	0.011	0.316
Northwind	End-to-End	0.92	66	0.0617	0.0010	0.154	0.010	0.164
	LLM-FK	1.00	71	0.1312	0.0012	0.328	0.012	0.339
TPC-E	End-to-End	0.58	528	0.4496	0.0066	1.124	0.066	1.190
	LLM-FK	0.92	245	0.5342	0.0031	1.335	0.032	1.367
AdventureWorks	End-to-End	0.52	2485	2.6313	0.0346	6.578	0.346	6.924
	LLM-FK	0.92	2142	5.5458	0.0224	13.864	0.223	14.087
MusicBrainz	End-to-End	0.75	28203	25.9330	0.3653	64.832	3.653	68.485
	LLM-FK	0.96	2588	6.3064	0.0264	15.766	0.264	16.029

Table 11: Overall computational cost comparison between the End-to-End baseline and LLM-FK across different datasets.

Dataset	Agent	API Calls	Tokens		Cost (¥)		
			In	Out	In	Out	Total
TPC-H	Profiler	8	0.0149	0.0001	0.037	0.001	0.038
	Interpreter	1	0.0003	0.0005	0.001	0.005	0.006
	Refiner	49	0.1004	0.0005	0.251	0.005	0.256
	Verifier	2	0.0064	0.0000	0.016	0.000	0.017
Northwind	Profiler	9	0.0067	0.0001	0.017	0.001	0.018
	Interpreter	1	0.0004	0.0005	0.001	0.005	0.005
	Refiner	60	0.1211	0.0006	0.303	0.006	0.309
	Verifier	1	0.0030	0.0000	0.007	0.000	0.008
TPC-E	Profiler	20	0.0140	0.0002	0.035	0.002	0.037
	Interpreter	1	0.0004	0.0006	0.001	0.006	0.007
	Refiner	194	0.4165	0.0019	1.041	0.019	1.061
	Verifier	30	0.1032	0.0004	0.258	0.004	0.262
AdventureWorks	Profiler	60	0.0433	0.0006	0.108	0.006	0.114
	Interpreter	1	0.0007	0.0009	0.002	0.009	0.011
	Refiner	2057	5.3956	0.0205	13.489	0.205	13.694
	Verifier	24	0.1061	0.0004	0.265	0.004	0.269
MusicBrainz	Profiler	140	0.0967	0.0014	0.242	0.014	0.256
	Interpreter	1	0.0014	0.0011	0.004	0.011	0.015
	Refiner	2355	5.8799	0.0228	14.700	0.228	14.927
	Verifier	92	0.3284	0.0011	0.821	0.011	0.832

Table 12: Detailed breakdown of computational overhead across internal agents in the LLM-FK framework.

Algorithm 2 The Overall Pipeline of LLM-FK

Require: Database \mathcal{D} , Instruction Set $\mathcal{I} = \{\mathcal{I}_{\mathcal{A}_p}, \mathcal{I}_{\mathcal{A}_i}, \mathcal{I}_{\mathcal{A}_r}, \mathcal{I}_{\mathcal{A}_v}\}$

Ensure: The set of FKs Φ

```
1: // Stage 1: Profiler ( $\mathcal{A}_p$ ) - Profiling & Unique-Key-Driven Schema Decomposition
2:  $\mathcal{J}_{raw} \leftarrow \text{ExtractINDs}(\mathcal{D})$  {Profiler invokes tool to extract raw INDs}
3:  $\mathcal{U} \leftarrow \text{ExtractMinUCCs}(\mathcal{D})$  {Profiler invokes tool to extract MinUCCs}
4:  $\mathcal{J} \leftarrow \text{FilterByType}(\mathcal{J}_{raw})$  {Profiler invokes tool to filter INDs by data types}
5:  $\mathcal{K} \leftarrow \emptyset$  {Initialize the set of selected unique keys}
6:  $\mathcal{T}_{ref} \leftarrow \text{GetReferencedTables}(\mathcal{J})$ 
7: for each table  $T \in \mathcal{T}_{ref}$  do
8:    $k \leftarrow \pi(T, \mathcal{U}_T, \mathcal{I}_{\mathcal{A}_p})$  {Select the most plausible referenced unique key}
9:    $\mathcal{K} \leftarrow \mathcal{K} \cup \{k\}$ 
10: end for
11:  $\Psi_{cand} \leftarrow \{(c_f, c_p) \in \mathcal{J} \mid c_p \in \mathcal{K}\}$  {Profiler prunes  $\mathcal{J}$  using selected unique keys}
12: // Stage 2: Interpreter ( $\mathcal{A}_i$ ) - Self-Augmented Domain Knowledge Injection
13:  $\mathcal{N}_{\Psi_{cand}} \leftarrow \text{ExtractTableNames}(\Psi_{cand})$  {Interpreter invokes tool to extract table names}
14:  $\mathcal{K}_{dom} \leftarrow \pi(\mathcal{N}_{\Psi_{cand}}, \mathcal{I}_{\mathcal{A}_i})$  {Interpreter constructs self-augmented domain knowledge}
15: // Stage 3: Refiner ( $\mathcal{A}_r$ ) - Multi-Perspective CoT Reasoning
16:  $\Psi_{pos} \leftarrow \emptyset$  {Initialize local FK result set}
17: for each candidate  $\psi \in \Psi_{cand}$  in parallel do
18:    $S_\psi \leftarrow \text{GetStructuralRepresentative}(\psi, \mathcal{D})$  {Refiner formulates SQL queries guided by Multi-
    Perspective CoT to derive structural representations}
19:    $p_\psi \leftarrow \pi(\psi, S_\psi, \mathcal{K}_{dom}, \mathcal{I}_{\mathcal{A}_r})$  {Refiner executes binary validation ( $p_\psi \in \{0, 1\}$ )}
20:   if  $p_\psi = 1$  then
21:      $\Psi_{pos} \leftarrow \Psi_{pos} \cup \{\psi\}$ 
22:   end if
23: end for
24: // Stage 4: Verifier ( $\mathcal{A}_v$ ) - Holistic Conflict Resolution
25:  $\mathcal{G} \leftarrow \text{ConstructGraph}(\Psi_{pos})$  {Construct schema topology graph}
26: // Phase 4.1: Resolve Multiple References First
27:  $\mathbf{C}_{multi} \leftarrow \text{DetectMultiRefConflicts}(\mathcal{G})$  {Verifier invokes tool to detect multiple references}
28: for each conflict set  $\Psi_{conflict} \in \mathbf{C}_{multi}$  do
29:    $S_{\Psi_{conflict}} \leftarrow \text{GetStructuralRepresentative}(\Psi_{conflict}, \mathcal{D})$  {Verifier formulates SQL queries guided
    by Multi-Perspective CoT to derive structural representations}
30:    $\Psi_{resolved} \leftarrow \pi(\Psi_{conflict}, S_{\Psi_{conflict}}, \mathcal{K}_{dom}, \mathcal{I}_{\mathcal{A}_v})$  {Retain the most plausible candidate FK}
31:    $\Psi_{pos} \leftarrow (\Psi_{pos} \setminus \Psi_{conflict}) \cup \Psi_{resolved}$ 
32: end for
33: // Phase 4.2: Iteratively Resolve Cyclic References
34: while TRUE do
35:    $\mathcal{G} \leftarrow \text{ConstructGraph}(\Psi_{pos})$  {Construct the graph with updated edges}
36:    $\mathbf{C}_{cycle} \leftarrow \text{DetectCyclicConflicts}(\mathcal{G})$  {Verifier invokes tool to detect cyclic references}
37:   if  $\mathbf{C}_{cycle} = \emptyset$  then
38:     break {Terminate if no cycles remain}
39:   end if
40:   Select the shortest conflict set  $\Psi_{conflict} \in \mathbf{C}_{cycle}$  {Prioritize the shortest cycle for resolution}
41:    $S_{\Psi_{conflict}} \leftarrow \text{GetStructuralRepresentative}(\Psi_{conflict}, \mathcal{D})$ 
42:    $\Psi_{resolved} \leftarrow \pi(\Psi_{conflict}, S_{\Psi_{conflict}}, \mathcal{K}_{dom}, \mathcal{I}_{\mathcal{A}_v})$  {Remove the least plausible candidate FK}
43:    $\Psi_{pos} \leftarrow (\Psi_{pos} \setminus \Psi_{conflict}) \cup \Psi_{resolved}$ 
44: end while
45:  $\Phi \leftarrow \Psi_{pos}$ 
46: return  $\Phi$ 
```

[Role-play]
You are a database architecture analyst specializing in relational database design, data lineage reasoning, and semantic integrity analysis.

[Task Description]
Your task is to reason through candidate minimal unique column combinations (MinUCCs) to identify the most probable referenced unique key for a given table. Table 'teacher' contains several candidate unique keys. Based on the given selection criterion, select the single most appropriate MinUCC that would serve as the referenced unique key for foreign keys.

[Sample Data]
The first five rows of the table are as follows:
id	name	address_code
1	Alice	001
2	Bob	002
3	Carol	005
4	Dave	003
5	Eve	004

[Selection Criterion]

1. Uniqueness Validity: The candidate must uniquely identify a tuple (satisfied by all candidates provided).
2. Column Ordinality: Columns positioned earlier in the table definition are prioritized.
3. Naming Convention: Preference is given to column names containing typical identifiers (e.g., "id", "key", "code").
4. Type Suitability: The data type is typically an integer or a string, adhering to standard key definitions.
5. Value Conciseness: The data values should be concise and human-readable (short text length).
6. Minimal Composition: Candidates with fewer constituent columns (atomic keys) are preferred over complex composite keys.
7. Semantic Alignment: The key must logically represent the entity according to the semantic context of the table.

[Candidate Solutions]

1. id
2. name
3. address_code

[Response Format]
Your response should be in this format: Selection: [the referenced unique key]

[Conclusion]
Selection: id

Figure 6: Prompt of the Profile for unique key selection.

[Role-play]
You are a database architecture analyst specializing in relational database design, data lineage reasoning, and semantic integrity analysis.

[Task Description]
The task involves processing a list of table names from a relational database to reconstruct foreign key relationships via semantic analysis. As a prerequisite, the initial step requires comprehending the application domain and acquiring the entity knowledge essential for foreign key identification, laying the groundwork for the subsequent process.

[Database Tables]
The following database tables need analysis: 'teacher', 'student', 'teacher_extra'.

[Analysis Steps]
Please perform a deep semantic analysis on the above tables and output a summary. Follow these steps:

1. Domain Inference: Based on the table names, infer the likely application domain of this database (e.g., E-commerce, HR, ERP, etc.).
2. Entity Definition: For each table in the list, determine what "entity" it represents and summarize what entity knowledge is needed to identify foreign keys semantically.

[Response Format]
Your response should be in this format:
Domain Summary: [Your analysis]
Entity Knowledge Summary for Foreign Key Identification: [Your summary]

[Conclusion]
Domain Summary: Educational Administration / Academic Management System.
Entity Knowledge Summary for Foreign Key Identification:
teacher: Core faculty entity. Serves as the principal entity that teacher_extra must reference.
student: Core learner entity. Expected to hold a potential semantic link to teacher (indicating advisor/mentor).
teacher_extra: Extension Table (Vertical Partitioning). Stores supplementary metadata. Semantically, it implies a strict 1:1 dependency on the teacher table, sharing the same logical grain.

Figure 7: Prompt of the Interpreter for domain knowledge generation.

[Role-play]

You are a database architecture analyst specializing in relational database design, data lineage reasoning, and semantic integrity analysis.

[Task Description]

A foreign key (FK) represents a referential constraint where values in one table (referencing) correspond to the primary or unique key values in another table (referenced), ensuring data consistency. Your task is to determine whether Column A functions as a foreign key referencing Column B based on the provided information.

[Domain Knowledge]

Domain Summary: Educational Administration / Academic Management System.

Entity Knowledge Summary for Foreign Key Identification:

teacher: Core faculty entity. Serves as the principal entity that teacher_extra must reference.

student: Core learner entity. Expected to hold a potential semantic link to teacher (indicating advisor/mentor).

teacher_extra: Extension Table (Vertical Partitioning). Stores supplementary metadata. Semantically, it implies a strict 1:1 dependency on the teacher table, sharing the same logical grain.

[Column A Information]

- Table Name: `student` - Column Name: `study_under` - Ordinal Position: 2 - Data Type: Integer - Average Length: 1

- Distinct Values (Cardinality): 2 - Total Rows: 5 - Cardinality Ratio (Cardinality / Total Rows): 0.4 - Min Value: 1 - Max Value: 2

- Sample Data:

The first five rows of the table are as follows:

```
| id | name | study_under | address_code |
```

```
---|---|---|---
```

```
| 1 | Elaine | 2 | 002 |
```

```
| 2 | Bonnie | 2 | 001 |
```

```
| 3 | Van | 1 | 003 |
```

```
| 4 | Kelly | 1 | 005 |
```

```
| 5 | Lucas | 1 | 004 |
```

[Column B Information]

- Table Name: `teacher` - Column Name: `id` - Ordinal Position: 0 - Data Type: Integer - Average Length: 1

- Distinct Values (Cardinality): 5 - Total Rows: 5 - Cardinality Ratio (Cardinality / Total Rows): 1 - Min Value: 1 - Max Value: 5

- Sample Data:

The first five rows of the table are as follows:

```
| id | name | address_code |
```

```
---|---|---
```

```
| 1 | Alice | 001 |
```

```
| 2 | Bob | 002 |
```

```
| 3 | Carol | 005 |
```

```
| 4 | Dave | 003 |
```

```
| 5 | Eve | 004 |
```

[Chain-of-Thought]

1. Semantic Analysis:

- Step 1: Interpret the real-world entities represented by the tables. Does `student.study_under` semantically imply a reference to a `teacher`?
- Step 2: Determine if the relationship follows logical "is-a", "has-a", or "assigned-to" semantics for foreign keys.

2. Syntactic Analysis:

- Analyze the similarity between column/table names. Higher similarity implies a higher probability of a relationship.
- Check for standard naming conventions (suffixes like `_id`, `_key`, `_no`).

3. Statistical Analysis (Inclusion Dependency is pre-verified):

- Coverage Ratio: The foreign key should reference a sufficient portion of the referenced key values; in this case, the coverage ratio is 0.4.
- Length Similarity: The average storage lengths of both columns should be identical or highly similar (1 vs 1).
- Range Logic: Foreign key references should be evenly distributed across the referenced key domain, instead of being limited to a small subset; in this case, the out-of-range ratio is 0.6.
- Table Size Ratio: Foreign keys typically exist in tables that are equal to or larger than the referenced table (1:N relationship). Here, row count ratio (`student` / `teacher`) is 1.0.

[Response Format]

Your response should be in this format: Decision: [True | False]

[Conclusion]

Decision: True

Figure 8: Prompt of the Refiner for each FK candidate pair validation.

[Role-play]

You are a database architecture analyst specializing in relational database design, data lineage reasoning, and semantic integrity analysis.

[Task Description]

Resolve the "Multiple References Conflict" where a single column in a source table exhibits potential foreign key relationship with multiple candidate columns in target tables. A valid foreign key design dictates that a source column should reference exactly one target column. Based on the provided information, select the single most appropriate foreign key relationship.

[Domain Knowledge]

Domain Summary: Educational Administration / Academic Management System.

Entity Knowledge Summary for Foreign Key Identification:

teacher: Core faculty entity. Serves as the principal entity that teacher_extra must reference.

student: Core learner entity. Expected to hold a potential semantic link to teacher (indicating advisor/mentor).

teacher_extra: Extension Table (Vertical Partitioning). Stores supplementary metadata. Semantically, it implies a strict 1:1 dependency on the teacher table, sharing the same logical grain.

[Referencing Column Information]

- Table Name: `student` - Column Name: `study_under` - Ordinal Position: 2 - Data Type: Integer - Average Length: 1

- Distinct Values (Cardinality): 2 - Total Rows: 5 - Cardinality Ratio (Cardinality / Total Rows): 0.4 - Min Value: 1 - Max Value: 2

- Sample Data:

The first five rows of the table are as follows:

```
| id | name | study_under | address_code |
|---|---|---|---|
| 1 | Elaine | 2 | 002 |
| 2 | Bonnie | 2 | 001 |
| 3 | Van | 1 | 003 |
| 4 | Kelly | 1 | 005 |
| 5 | Lucas | 1 | 004 |
```

[Referenced Columns Information]

- Table Name: `teacher` - Column Name: `id` - Ordinal Position: 0 - Data Type: Integer - Average Length: 1

- Distinct Values (Cardinality): 5 - Total Rows: 5 - Cardinality Ratio (Cardinality / Total Rows): 1 - Min Value: 1 - Max Value: 5

- Sample Data:

The first five rows of the table are as follows:

```
| id | name | address_code |
|---|---|---|
| 1 | Alice | 001 |
| 2 | Bob | 002 |
| 3 | Carol | 005 |
| 4 | Dave | 003 |
| 5 | Eve | 004 |
```

- Table Name: `teacher_extra` - Column Name: `id` - Ordinal Position: 0 - Data Type: Integer - Average Length: 1

- Distinct Values (Cardinality): 5 - Total Rows: 5 - Cardinality Ratio (Cardinality / Total Rows): 1 - Min Value: 1 - Max Value: 5

- Sample Data:

The first five rows of the table are as follows:

```
| id | subject |
|---|---|
| 1 | Mathematics |
| 2 | Computer Science |
| 3 | Literature |
| 4 | Physics |
| 5 | World History |
```

[Chain-of-Thought]

1. Semantic Analysis:

- Step 1: Interpret the entities represented by the tables.

- Step 2: Determine if the relationship follows logical "is-a", "has-a", or "assigned-to" semantics for foreign keys.

2. Syntactic Analysis:

- Analyze the similarity between column/table names. Higher similarity implies a higher probability of a relationship.

- Check for standard naming conventions (suffixes like `_id`, `_key`, `_no`).

3. Statistical Analysis (Inclusion Dependency is pre-verified):

- Coverage Ratio: The foreign key should reference a sufficient portion of the referenced key values; the coverage ratios are 0.4, 0.4.

- Length Similarity: The average storage lengths of both columns should be identical or highly similar (1 vs 1 vs 1).

- Range Logic: Foreign key references should be evenly distributed across the referenced key domain, instead of being limited to a small subset; the out-of-range ratios are 0.6, 0.6.

- Table Size Ratio: Foreign keys typically exist in tables that are equal to or larger than the referenced table (1:N relationship); row count ratios are 1.0, 1.0.

[Response Format]

Your response should be in this format: Selection: [the most appropriate referenced column]

[Conclusion]

Selection: teacher.id

Figure 9: Prompt of the Verifier for multiple references resolution.

[Role-play]

You are a database architecture analyst specializing in relational database design, data lineage reasoning, and semantic integrity analysis.

[Task Description]

Resolve the "Cyclic References Conflict" where the potential foreign keys form a closed loop. Based on the provided information, select the least plausible (weakest) reference relationship within the cycle and remove it to break the loop.

[Domain Knowledge]

Domain Summary: Educational Administration / Academic Management System.

Entity Knowledge Summary for Foreign Key Identification:

teacher: Core faculty entity. Serves as the principal entity that teacher_extra must reference.

student: Core learner entity. Expected to hold a potential semantic link to teacher (indicating advisor/mentor).

teacher_extra: Extension Table (Vertical Partitioning). Stores supplementary metadata. Semantically, it implies a strict 1:1 dependency on the teacher table, sharing the same logical grain.

[Conflict Description]

- 1. teacher.id → teacher.extra.id
- 2. teacher_extra.id → teacher.id

[Columns Information]

- Table Name: 'teacher' - Column Name: 'id' - Ordinal Position: 0 - Data Type: Integer - Average Length: 1
- Distinct Values (Cardinality): 5 - Total Rows: 5 - Cardinality Ratio (Cardinality / Total Rows): 1 - Min Value: 1 - Max Value: 5
- Sample Data:

The first five rows of the table are as follows:

id	name	address_code
1	Alice	001
2	Bob	002
3	Carol	005
4	Dave	003
5	Eve	004

- Table Name: 'teacher_extra' - Column Name: 'id' - Ordinal Position: 0 - Data Type: Integer - Average Length: 1
- Distinct Values (Cardinality): 5 - Total Rows: 5 - Cardinality Ratio (Cardinality / Total Rows): 1 - Min Value: 1 - Max Value: 5
- Sample Data:

The first five rows of the table are as follows:

id	subject
1	Mathematics
2	Computer Science
3	Literature
4	Physics
5	World History

[Chain-of-Thought]

- 1. Semantic Analysis:
 - Step 1: Interpret the entities represented by the tables.
 - Step 2: Determine if the relationship follows logical "is-a", "has-a", or "assigned-to" semantics for foreign keys.
- 2. Syntactic Analysis:
 - Analyze the similarity between column/table names. Higher similarity implies a higher probability of a relationship.
 - Check for standard naming conventions (suffixes like '_id', '_key', '_no').
- 3. Statistical Analysis (Inclusion Dependency is pre-verified):
 - Coverage Ratio: The foreign key should reference a sufficient portion of the referenced key values; the coverage ratios are 1.0, 1.0.
 - Length Similarity: The average storage lengths of both columns should be identical or highly similar (1 vs 1), (1 vs 1).
 - Range Logic: Foreign key references should be evenly distributed across the referenced key domain, instead of being limited to a small subset; the out-of-range ratios are 0, 0.
 - Table Size Ratio: Foreign keys typically exist in tables that are equal to or larger than the referenced table (1:N relationship); row count ratios are 1.0, 1.0.

[Response Format]

Your response should be in this format: Selection: [the least plausible reference relationship]

[Conclusion]

Selection: teacher.id → teacher_extra.id

Figure 10: Prompt of the Verifier for cyclic references resolution.

[Role-play]
You are a database architecture analyst specializing in relational database design, data lineage reasoning, and semantic integrity analysis.

[Task Description]
Identify foreign key relationships between the following two tables based strictly on the provided information. For each identified foreign key, output a list entry. Taking a scenario where Table A references Table B as an example, use this exact format:
[<Table A name>, <foreign key column name in Table A>, <Table B name>, <referenced column name in Table B>].

[Table Information]
- Table name: student
The first five rows of the table are as follows:
| id | name | study_under | address_code |
|---|---|---|---|
| 1 | Elaine | 2 | 002 |
| 2 | Bonnie | 2 | 001 |
| 3 | Van | 1 | 003 |
| 4 | Kelly | 1 | 005 |
| 5 | Lucas | 1 | 004 |

- Table name: teacher
The first five rows of the table are as follows:
| id | name | address_code |
|---|---|---|
| 1 | Alice | 001 |
| 2 | Bob | 002 |
| 3 | Carol | 005 |
| 4 | Dave | 003 |
| 5 | Eve | 004 |

[Response Format]
Your response should be in this format: Foreign Keys: [the identified foreign keys]
Example Output:
1. [['orders', 'customer_id', 'customers', 'id']]
2. [[]]

Figure 11: Prompt of the End-to-End baseline.

[End-to-End Prompt]
.....

[Examples]
Example 1: Single Foreign Key Relationship
Input:
- orders (the first five rows):
| order_id | date | user_id | amount |
|---|---|---|---|
| 101 | 2023-01-01 | 5 | 100.0 |
| 102 | 2023-01-02 | 8 | 50.5 |
| 103 | 2023-01-03 | 5 | 120.0 |
| 104 | 2023-01-04 | 6 | 200.0 |
| 105 | 2023-01-05 | 7 | 80.0 |
- users (the first five rows):
| uid | user_name | email |
|---|---|---|
| 5 | Alice | alice@example.com |
| 6 | Bob | bob@test.com |
| 7 | Charlie | charlie@corp.com |
| 8 | David | david@mail.com |
| 9 | Eve | eve@web.com |
Output: [['orders', 'user_id', 'users', 'uid']]

Example 2: No Foreign Key Relationship
Input:
- logs (the first five rows):
| log_id | timestamp | message |
|---|---|---|
| 1 | 12:00:01 | Error 404 |
| 2 | 12:00:05 | OK 200 |
| 3 | 12:01:10 | OK 200 |
| 4 | 12:02:45 | Error 500 |
| 5 | 12:05:00 | OK 200 |

- employees (the first five rows):
| emp_id | name | department |
|---|---|---|
| 99 | Frank | HR |
| 100 | Grace | IT |
| 101 | Heidi | Sales |
| 102 | Ivan | Marketing |
| 103 | Judy | IT |
Output: [[]]

Example 3: Multiple Foreign Key Relationships
Input:
- flights (the first five rows):
| flight_no | dept_code | arr_code | time |
|---|---|---|---|
| AA123 | JFK | LHR | 10:00 |
| BA001 | LHR | JFK | 14:00 |
| DL456 | LAX | HND | 08:30 |
| JL002 | HND | SFO | 16:45 |
| UA789 | SFO | LAX | 12:15 |
- airports (the first five rows):
| iata_code | city | country |
|---|---|---|
| JFK | New York | USA |
| LHR | London | UK |
| LAX | Los Angeles | USA |
| HND | Tokyo | Japan |
| SFO | San Francisco | USA |
Output: [['flights', 'dept_code', 'airports', 'iata_code'],
['flights', 'arr_code', 'airports', 'iata_code']]

End of Examples.

Figure 12: Prompt of the Few-Shot baseline.

[End-to-End Prompt]

[Examples]

Example 1: Single Foreign Key Relationship

Input:

- orders (the first five rows):

| order_id | date | user_id | amount |

---|---|---|---

| 101 | 2023-01-01 | 5 | 100.0 |

| 102 | 2023-01-02 | 8 | 50.5 |

| 103 | 2023-01-03 | 5 | 120.0 |

| 104 | 2023-01-04 | 6 | 200.0 |

| 105 | 2023-01-05 | 7 | 80.0 |

- users (the first five rows):

| uid | user_name | email |

---|---|---

| 5 | Alice | alice@example.com |

| 6 | Bob | bob@test.com |

| 7 | Charlie | charlie@corp.com |

| 8 | David | david@mail.com |

| 9 | Eve | eve@web.com |

Output: [['orders', 'user_id', 'users', 'uid']]

Explanation: Semantically, the 'orders' table records transactions made by specific entities. The column 'user_id' explicitly names the entity it refers to ('user'). Comparing this to the 'users' table, 'uid' serves as the primary identifier for user entities. Since 'user_id' in orders matches the semantic meaning of 'uid' in users, it is a foreign key describing 'which user placed this order'.

Example 2: No Foreign Key Relationship

Input:

- logs (the first five rows):

| log_id | timestamp | message |

---|---|---

| 1 | 12:00:01 | Error 404 |

| 2 | 12:00:05 | OK 200 |

| 3 | 12:01:10 | OK 200 |

| 4 | 12:02:45 | Error 500 |

| 5 | 12:05:00 | OK 200 |

- employees (the first five rows):

| emp_id | name | department |

---|---|---

| 99 | Frank | HR |

| 100 | Grace | IT |

| 101 | Heidi | Sales |

| 102 | Ivan | Marketing |

| 103 | Judy | IT |

Output: [[]]

Explanation: Semantically, 'server_logs' tracks system events, while 'employees' tracks human resources. There is no logical semantic relationship between a system log and an employee record in this schema. Although 'log_id' and 'emp_id' are both unique numeric identifiers, their column names indicate they represent completely different real-world entities (a log entry vs. a person). Therefore, no foreign key exists.

Example 3: Multiple Foreign Key Relationships

Input:

- flights (the first five rows):

| flight_no | dept_code | arr_code | time |

---|---|---|---

| AA123 | JFK | LHR | 10:00 |

| BA001 | LHR | JFK | 14:00 |

| DL456 | LAX | HND | 08:30 |

| JL002 | HND | SFO | 16:45 |

| UA789 | SFO | LAX | 12:15 |

- airports (the first five rows):

| iata_code | city | country |

---|---|---

| JFK | New York | USA |

| LHR | London | UK |

| LAX | Los Angeles | USA |

| HND | Tokyo | Japan |

| SFO | San Francisco | USA |

Output: [['flights', 'dept_code', 'airports', 'iata_code'], ['flights', 'arr_code', 'airports', 'iata_code']]

Explanation: Semantically, a flight involves two locations: a departure point and an arrival point. The columns 'dept_code' (Departure) and 'arr_code' (Arrival) in the 'flights' table both describe airport entities. The 'airports' table contains the master list of these entities, identified by 'iata_code'. Since both columns in 'flights' utilize the standard IATA location codes found in the 'airports' table, they both reference 'airports.iata_code'.

End of Examples.

[Chain-of-Thought]

Now, analyze the specific information provided. Please think step by step, provide your explanation first, followed by the Python list code block.

Figure 13: Prompt of the CoT baseline.