

Squrve: A Unified and Modular Framework for Complex Real-World Text-to-SQL Tasks

Yihan Wang¹ Peiyu Liu^{2*} Runyu Chen² Jiaxing Pu³ Wei Xu¹

¹Renmin University of China ²University of International Business and Economics

³Zhejiang University of Technology

yihan3123@gmail.com liupeiyustu@163.com ry.chen@uibe.edu.cn

pujiaxing@zjut.edu.cn weixu@ruc.edu.cn

Abstract

Text-to-SQL technology has evolved rapidly, with diverse academic methods achieving impressive results. However, deploying these techniques in real-world systems remains challenging due to limited integration tools. Despite these advances, we introduce Squrve, a unified, modular, and extensive Text-to-SQL framework designed to bring together research advances and real-world applications. Squrve first establishes a universal execution paradigm that standardizes invocation interfaces, then proposes a multi-actor collaboration mechanism based on seven abstracted effective atomic actor components. Experiments on widely adopted benchmarks demonstrate that the collaborative workflows consistently outperform the original individual methods, thereby opening up a new effective avenue for tackling complex real-world queries. The codes are available at <https://github.com/LLM-Cube/Squrve>.

1 Introduction

Text-to-SQL (Liu et al., 2024; Shi et al., 2025; Hong et al., 2025) aims to translate natural language questions into accurate SQL queries, enabling non-expert users to access and manipulate data effortlessly. As an emerging querying paradigm, Text-to-SQL facilitates real-world database-driven AI agent applications, such as conversational business intelligence tools or automated decision-support assistants (Wang et al., 2022a; Xu et al., 2020). However, most existing studies tend to focus on specific aspects of the problem, *e.g.*, schema linking (Pourreza and Rafiei, 2023), decoding strategy (Pourreza et al., 2025), or data augmentation (Talaie et al., 2024). In contrast, real-world Text-to-SQL problems are inherently more complex, where multiple challenges often coexist. This discrepancy highlights the need for a unified framework that can systematically integrate, compare,

and extend different methods, enabling complementary strengths from diverse approaches to be effectively combined and applied to practical tasks.

There are two critical challenges that hinder the development of such a unified framework. *Challenge 1: lack of unified integration mechanism.* Existing methods are typically designed for specific settings, each with distinct data formats, input-output interfaces, or dependency requirements. This fragmentation makes it difficult to combine or switch models across diverse scenarios, resulting in limited cross-task adaptability (Gan et al., 2021b) and poor scalability to heterogeneous real-world environments. *Challenge 2: tackling diverse and co-occurring difficulties in real-world scenarios.* Real-world queries often involve multiple difficulties (*e.g.*, domain adaptation and multi-table reasoning) that no single model can fully address. Integrating complementary strengths across models is therefore crucial but remains difficult without coordinated interfaces and execution pipelines. These challenges are often overlooked, widening the gap between academic research and real-world Text-to-SQL applications.

To address these challenges, we propose Squrve, a unified, modular, and extensive Text-to-SQL framework that systematically addresses the challenges of diverse real-world applications by fully leveraging academic advances. For Challenge 1, Squrve introduces a universal Text-to-SQL execution paradigm that standardizes invocation interfaces and parameter configurations. Through its modular architecture, Squrve reproduces and integrates existing advanced academic models, which achieves effective generalization across diverse Text-to-SQL task scenarios. For Challenge 2, Squrve abstracts and formalizes seven atomic Actor components, each representing a distinct Text-to-SQL capability validated in prior research. Building on these components, Squrve introduces a *multi-actor collaboration mechanism* that effec-

* Corresponding authors.

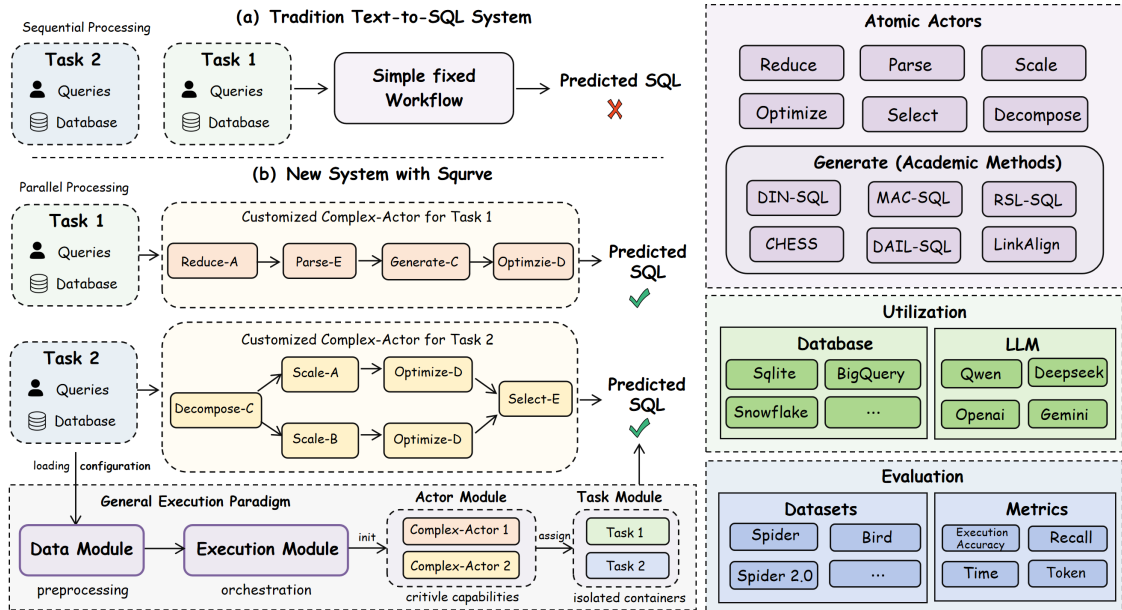


Figure 1: The overview of the Sqrve framework. Sqrve unifies diverse methods under a general execution paradigm and supports task-specific multi-actor collaboration, bridging academic research and real-world systems.

tively fuses their complementary strengths, allowing different actors to interact and cooperate when handling complex real-world queries.

To evaluate Sqrve’s performance, we conduct comprehensive experiments on the widely adopted Spider (Yu et al., 2018) and Bird (Li et al., 2023c) benchmarks, systematically assessing the execution accuracy of the generated SQL queries. The results indicate that Sqrve faithfully reproduces baseline methods while preserving their performance, which achieves or even surpasses the originally reported results based solely on open-source models. Furthermore, the ensemble methods built upon multiple actors consistently outperform the original individual methods, demonstrating the effectiveness of compositional reasoning through the integration of complementary actors. These findings highlight that Sqrve bridges the gap between academic research and real-world applications, establishing a pathway for advancing practical development.

To facilitate reproducibility, we provide the video demonstration of Sqrve at [YouTube](#) and an interactive online deployment at [HuggingFace](#).

2 Architecture Design

To generalize academic techniques to diverse real-world scenarios, Sqrve proposes a unified, modular and extensible framework comprising four key modules. Specifically, the **data module** standardizes different dataset formats into unified structures;

the **actor module** defines critical atomic Text-to-SQL capability and enables their collaborative composition; the **task module** serves as isolated containers that enable concurrent processing; and the **execution module** orchestrates the complete Text-to-SQL execution with global control.

2.1 Data Module

To standardize data operations, Sqrve first defines unified structures for database schemas and query instances, following the field naming conventions of mainstream datasets such as Spider, Bird, and others. Building on this foundation, the Data module provides universal preprocessing methods that transform both user-uploaded private datasets and built-in benchmarks into standardized formats for seamless access in downstream stages.

Database schema preprocessing. Existing datasets typically store complete large-scale database schemas in a single centralized file, which often exceeds the LLM’s context window. To alleviate this limitation, this module decomposes centralized database schemas into parallel column-level files, where each file serves as a semantic unit containing complete column metadata without additional storage overhead, enabling efficient retrieval of relevant subsets via RAG methods.

Query instance preprocessing. Sqrve constructs and maintains more than 200 chain-of-thought (CoT) exemplars from existing benchmark

datasets, which can be embedded and indexed for efficient vector retrieval. For new complex queries, this module retrieves semantically similar exemplars and assembles them into complete CoT prompts, effectively enhancing the quality of SQL generation. For queries that depend on external documents, this module extracts key contextual information, such as concept definitions and metric calculation formulas, using carefully designed chain-of-thought strategies.

2.2 Actor Module

To better handle complex real-world queries, Sqrve decomposes the Text-to-SQL process into seven atomic actors, each capturing a fundamental capability that has been validated in prior research (see Table 1). Compared with complete end-to-end workflows, these actors generalize more flexibly across tasks by focusing on the core functional components. Building on these actors, Sqrve introduces a multi-actor collaboration mechanism that integrates actors derived from different methods, enabling them to work jointly and share complementary strengths within unified workflows. For clarity in subsequent discussion, an actor can be described as a black-box that generates the required SQL end-to-end:

$$SQL = f_{\text{actor}}(Q, S, C|M), \quad (1)$$

where Q denotes user queries, S denotes the database schemas, C represents valuable contextual information specific to the query, and M indicates the LLM backbone. Next, we will introduce the atomic actors and their collaboration mechanism in detail.

2.2.1 Atomic Actors

This section provides a detailed explanation of the seven atomic actors in Sqrve, outlining their individual roles and functional contributions within the overall Text-to-SQL workflow.

Reduce actor. This component eliminates redundant schemas from large-scale databases that may exceed LLM context windows. By retaining only the relevant subsets, the reduce actor allows subsequent stages to operate on compact and focused schema representations.

Parse actor. This component performs schema linking by extracting tables and columns from candidate schemas that are potentially required for SQL generation. By providing attention signals,

the parse actor enables easier focus on critical elements during subsequent SQL generation.

Generate actor. This component generates complete SQL statements, encapsulating existing end-to-end academic methods. The generator actor can be applied directly to Text-to-SQL tasks in straightforward scenarios.

Decompose actor. This component decomposes complex queries into multiple logically progressive sub-questions and generates SQL statements for each. By simulating chain-of-thought (CoT) reasoning, the decompose actor provides interpretable reasoning pathways for complex queries.

Scale actor. This component generates diverse high-quality SQL candidates to increase the probability of covering the gold SQL. Through leveraging abundant inference-time computation, the scale actor improves generation quality and overcomes the limitations of the single-decoding strategy.

Optimize actor. This component leverages environmental feedback (*e.g.*, database errors or results) to refine the quality of the generated SQL queries. Through feedback signals, the optimize actor can rapidly locate syntax errors or identify discrepancies between retrieval results and user intent.

Select actor. This component selects the optimal SQL statement from multiple candidates, typically in collaboration with the scale actor. By filtering out suboptimal candidates, the select actor enables broad explorations of generated SQL to converge toward the final output.

2.2.2 Complex actors

To support scalable integration of atomic actors, Sqrve introduces two composition strategies, **Pipeline** and **Tree**, which serve as structural templates with plug-and-socket connection interfaces. These strategies wrap multiple actors into a single composite actor, which can be recursively applied to construct complex workflows.

Pipeline. The Pipeline strategy connects multiple actors in a sequential manner, where the output of each stage serves as the input to the next. This design enables multi-step transformations through staged processing, allowing intermediate results to be refined progressively along the workflow.

Tree. The Tree strategy dispatches a shared root input to multiple child actors for parallel processing

Actors	Methods	Key Challenges
Reduce	LinkAlign (Wang et al., 2025)	large-scale and multi-database
Parse	LinkAlign; RSL-SQL (Cao et al., 2024)	ambiguous queries and redundant schemas
Generate	DIN-SQL; CHESS (Talaie et al., 2024)	efficient and high-quality SQL generation
Decompose	DIN-SQL; MAC-SQL (Wang et al., 2024)	Chain-of-Thought for complex queries
Scale	CHESS; CHASE-SQL (Pourreza et al., 2025)	diverse and high-quality decoding strategies
Optimize	CHASE-SQL; OpenSearch (Xie et al., 2025);	effective and broader database feedback
Select	OpenSearch; MCS-SQL (Lee et al., 2024);	accurate gold SQL identification

Table 1: Summary of the prior methods that each Actor in Sqrve builds upon, and the corresponding challenges highlighted in earlier research.

and then merges their outputs into a unified result. This structure supports ensemble-style execution, where different processing pathways contribute complementary interpretations or candidate SQLs that collectively improve overall performance.

Identified combinations. To explore effective collaboration patterns, we evaluate multiple actor combinations and identify two favorable ensembles (see Section 4.1). Specifically, we randomly sample 50 instances from the Spider and BIRD datasets, and use Qwen3-turbo as the backbone model to generate predictions. All results are summarized in Table 4, based on which we construct two high-performing multi-actor workflows that serve as the foundation for subsequent formal experiments.

2.3 Task Module

The task module coordinates the execution of Text-to-SQL tasks by linking actors, data, and runtime environments under a unified interface. Specifically, it includes an **actor interface**, used to perform the core Text-to-SQL conversion process; a **dataset interface**, used to load preprocessed task-specific inputs such as user queries and database schemas through data module’s interface; and a **runtime controller**, used to manage execution flow (*e.g.*, concurrency control) and collect generated SQL outputs. Together, these components provide an isolated and concurrent runtime environment, ensuring efficient task scheduling and preventing interference across different executions.

2.4 Execution Module

The execution module serves as the central controller of Sqrve, responsible for coordinating the entire Text-to-SQL workflow—from system initialization to final evaluation. Specifically, it includes a **configuration manager**, used to load system settings and initialize all task containers; an **execution controller**, used to launch, monitor, and

manage concurrent task executions while handling logging and exception recovery; and an **evaluation unit**, used to aggregate results, compute metrics such as execution accuracy, and generate evaluation reports. Together, these components ensure stable large-scale execution and reproducible performance assessment across diverse Text-to-SQL scenarios.

3 Usage Examples

This section illustrates the usage of Sqrve, covering general startup and invocation, running baseline methods on benchmark datasets, and composing customized workflows from atomic components.

3.1 General Startup Demonstration

Overall, Sqrve offers a unified interface that executes diverse Text-to-SQL workflows through minimal user configuration. As illustrated in Appendix D, users can run different methods with only a few lines of code. The framework automatically handles data loading, execution, and evaluation, producing complete results including predicted SQLs, query outputs, and evaluation metrics. This unified interface simplifies the use of advanced Text-to-SQL methods, then users can execute and compare different models with a single line of code, without the need to modify separate scripts or pipelines. In future releases, we plan to further enrich the configuration system, enabling users to flexibly test various benchmarks and baseline methods in Section 3.2. Moreover, the configuration will support fine-grained control over Actor settings, allowing users to define customized composition strategies as detailed in Section 3.3.

3.2 Running Existing Baselines

Unlike existing systems that rely on fixed, task-specific algorithms, Sqrve integrates a variety of advanced open-source methods that can be flexibly applied to diverse Text-to-SQL scenarios. For a

given dataset, this can be easily achieved through the following steps:

- **Formatting the dataset.** For existing widely used benchmarks, Sqrve has pre-unified their formats and maintains them in the 'benchmark' directory, requiring no additional operations. For proprietary datasets, users need to preprocess them into valid formats following Sqrve's documentation. Additionally, databases with access control (*e.g.*, BigQuery) require users to prepare authentication credentials in advance.
- **Generating the configuration.** Sqrve provides flexible approaches to generate configurations, including configuration files, command-line arguments, or direct code specification. Users typically need to specify the dataset path, method name, task definition, and other configuration settings.
- **Starting the execution process.** Once the configuration is prepared, users can directly start the execution process. Sqrve automatically handles dataset initialization, task execution, and output saving.
- **Evaluating and organizing the results.** After task completion, users can evaluate the performance using multiple evaluation metrics, such as execution accuracy, and further analyze results through visualization.

3.3 Customize Implementations

For challenging tasks, users may need to integrate advantages from different methods and design customized workflows. This can be achieved by replacing method selection with workflow design in the generating the configuration step:

- **Selecting the atomic actors.** Atomic actors are the fundamental components, like the bricks that form a building. Therefore, users first select the necessary atomic actors for their workflow, such as DINSQLParser for the schema linking step.
- **Designing the customized workflow.** Users can scale up the final performance through two composition strategies: serially chaining different actors or composing complementary actors in parallel. Sqrve represents such topological structures using nested lists. For more

advanced usage, users can construct more sophisticated models based on previously defined simple workflows.

4 Experiment

4.1 Experiment Setup

Datasets. We conduct performance evaluations on three widely adopted Text-to-SQL benchmarks, Spider (Yu et al., 2018), BIRD (Li et al., 2023c), and AmbiDB (Wang et al., 2025). Spider is a large-scale, cross-domain dataset for assessing model generalization across diverse database schemas. BIRD features larger databases, noisier data, and more complex queries requiring external knowledge. AmbiDB introduces enhanced ambiguity through multi-database settings, requiring to identify target database without prior specification.

Baselines. Sqrve integrates multiple state-of-the-art baselines: DIN-SQL employs decomposed in-context learning with self-correction; CHESS-SQL leverages context mining for effective SQL synthesis; MAC-SQL (Wang et al., 2024) utilizes multi-agent collaboration to identify relevant schema subsets; RSL-SQL (Cao et al., 2024) adopts diverse strategies to enhance schema linking and SQL generation robustness; and LinkAlign (Wang et al., 2025) addresses schema linking challenges in large-scale, multi-database environments.

Variants. We identify two promising and competitive multi-actor ensemble workflows through exploring different combination strategies in a small-scale subset, as detailed in Section 2.2.2. We further compare the performance of these variants with individual baselines through experiments.

1. **Ensemble-1:** LinkAlignParser + RLSQLScaler + CHESSSelector + MACSQLOptimizer
2. **Ensemble-2:** RLSQLBiDirParser + MACSQLDecompose + ChessScaler + CHESSSelector + LinkAlignOptimizer

Evaluation Metrics. We evaluate end-to-end Text-to-SQL performance using Execution Accuracy (EX), which compares the execution results of predicted and gold SQL queries. For schema linking, Recall and Precision are used to assess the accuracy of relevant schema element retrieval.

Methods	Individual	EX(Spider-dev)		EX(Bird-dev)	
		Origin	Ours	Origin	Ours
DIN-SQL	✓	81.7	82.5	59.4	60.1
CHESS	✓	85.6	84.8	64.6	63.1
MAC-SQL	✓	82.5	86.0	67.7	67.0
RSL-SQL	✓	86.5	87.3	67.0	68.0
LinkAlign	✓	87.2	86.5	64.0	65.5
Ensemble-1	×	N/A	88.7	N/A	67.3
Ensemble-2	×	N/A	90.8	N/A	70.0

Table 2: Comparison between individual baseline methods and our variants on Qwen-3 backbone. The *Origin* reports results obtained by executing the original repository code; *Ours* presents the reproduced performance under Sqrve framework. N/A indicates our proposed variants have no results from original code execution.

Implementations. We use Qwen-3 as the LLM backbone. Building upon the paradigm introduced in Section 2, Sqrve enables seamless extensibility of benchmark datasets and baseline methods along two orthogonal dimensions. We can freely combine method-dataset pairs and execute task batches concurrently for improved efficiency.

4.2 Main Results

Results of existing methods As shown in Table 2, Sqrve successfully reproduces existing Text-to-SQL baselines under the same LLM backbone with performance closely aligned to their originally executed results. Specifically, our reproduced results achieve comparable or even superior execution accuracy across both benchmarks. For instance, DIN-SQL achieves 82.5% on Spider-dev (82.5% vs. original 81.7%) and improves to 60.1% on Bird-dev (60.1% vs. original 59.4%). Overall, this faithful reproduction establishes a reliable foundation for fair comparison.

Results of our variants. Building upon the reproduced components, we further demonstrate the extensibility and effectiveness of the Multi-Actor collaboration mechanism through evaluation on two variants. As shown in Table 2, both ensemble methods substantially outperform all individual baselines in both benchmarks. In particular, Ensemble-2 achieved the best performance on both benchmarks with execution accuracies of 90.8% and 70.0%, improving 4.0% and 2.9% over the strongest individual method, respectively. These results demonstrate the potential of Multi-Actor collaboration to address complex queries by integrating complementary method strengths.

4.3 Parallel Component Analysis

While serially chaining multiple atomic actors scales up the final performance, parallel compo-

sition of complementary homogeneous actors orthogonal to the workflow can effectively enhance intermediate stages, such as schema linking, without efficiency loss. As shown in Table 3, all the parallel ensemble methods outperform individual methods in recall, which is typically prioritized in schema linking evaluation. Notably, two-actor parallelization strikes the optimal balance, as additional actors provide diminishing marginal gains while introducing computational overhead. This finding offers important insights for effectively harnessing multi-actor collaboration.

Methods	AmbiDB		Bird-dev	
	Recall	Precision	Recall	Precision
LinkAlign	49.0	51.5	76.9	81.5
RSL-SQL	54.5	8.2	82.4	35.3
DIN-SQL	21.6	23.0	59.8	34.2
Link.+RSL.	61.8	13.9	92.4	40.0
Link.+DIN.	55.3	42.4	81.2	58.1
RSL.+DIN.	64.5	8.1	87.1	29.0
Link.+RSL.+DIN.	62.3	7.3	93.1	28.7

Table 3: Parallel actors performance comparison.

5 Conclusion

This paper introduces Sqrve, a unified, modular, and extensive Text-to-SQL framework that bridges the gap between academic researches and real-world applications. Through a universal workflow paradigm and modular architecture, Sqrve decouples capabilities from usages, enabling researchers to seamlessly integrate, replace, and reuse technical components. Experiments validate the effectiveness, demonstrating that ensemble workflows outperform single end-to-end models. We believe Sqrve will serve as a foundational platform, accelerating real-world Text-to-SQL applications.

Acknowledgments

This work was partially supported by the National Natural Science Foundation of China under Grant No. 62506077, 72201061, and 72271233; Suzhou Key Laboratory of Artificial Intelligence and Social Governance Technologies (SZS2023007), Smart Social Governance Technology and Innovative Application Platform (YZCXPT2023101), and the Innovation System of the Integration between Industry and Education for Smart Governance (CJRH2024101).

References

- Zhenbiao Cao, Yuanlei Zheng, Zhihao Fan, Xiaojin Zhang, Wei Chen, and Xiang Bai. 2024. Rsl-sql: Robust schema linking in text-to-sql generation. *arXiv preprint arXiv:2411.00073*.
- DongHyun Choi, Min C Shin, EungGyun Kim, and Dong Ryeol Shin. 2021. Ryansql: Recursively applying sketch-based slot fillings for complex text-to-sql in cross-domain databases.
- Michael Deng, Akshara Ramachandran, Cheng Xu, Lyuzhou Hu, Zhoujun Yao, Anindya Datta, and Heng Zhang. 2025. Reforce: A text-to-sql agent with self-refinement, format restriction, and column exploration. In *ICLR 2025 The VerifAI Workshop*.
- Xiang Deng, Ahmed Hassan Awadallah, Christopher Meek, Oleksandr Polozov, Huan Sun, and Matthew Richardson. 2021. Structure-grounded pretraining for text-to-sql. In *North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Xuemei Dong, Chao Zhang, Yuhang Ge, Yuren Mao, Yunjun Gao, Jinshu Lin, Dongfang Lou, and 1 others. 2023. C3: Zero-shot text-to-sql with chatgpt.
- Yujian Gan, Xinyun Chen, Qiuping Huang, Matthew Purver, John R Woodward, Jinxia Xie, and Pengsheng Huang. 2021a. Towards robustness of text-to-sql models against synonym substitution. In *Association for Computational Linguistics and International Joint Conference on Natural Language Processing*.
- Yujian Gan, Xinyun Chen, and Matthew Purver. 2021b. Exploring underexplored limitations of cross-domain text-to-sql generalization. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Dawei Gao, Haibin Wang, Yaliang Li, Xiuyu Sun, Yue Qian, Bolin Ding, and Jingren Zhou. 2024. Text-to-sql empowered by large language models: A benchmark evaluation. In *International Conference on Very Large Data Bases*.
- Zijin Hong, Zheng Yuan, Qinggang Zhang, Hao Chen, Junnan Dong, Feiran Huang, and Xiao Huang. 2025. Next-generation database interfaces: A survey of llm-based text-to-sql. *arXiv preprint arXiv:2406.08426*.
- Binyuan Hui, Xiang Shi, Ruiying Geng, Binhua Li, Yongbin Li, Jian Sun, and Xiaodan Zhu. 2021. Improving text-to-sql with schema dependency learning.
- Siqi Kou, Lanxiang Hu, Zhezhi He, Zhijie Deng, and Hao Zhang. 2024. Cllms: Consistency large language models. In *International Conference on Machine Learning*.
- Dongjun Lee, Choongwon Park, Jaehyuk Kim, and Heesoo Park. 2024. Mcs-sql: Leveraging multiple prompts and multiple-choice selection for text-to-sql generation. *arXiv preprint arXiv:2405.07467*.
- Fangyu Lei, Jixuan Chen, Yuxiao Ye, Ruisheng Cao, Dongchan Shin, Hongjin SU, Zhaoqing SUO, Hongcheng Gao, Wenjing Hu, Pengcheng Yin, and 1 others. 2025. Spider 2.0: Evaluating language models on real-world enterprise text-to-sql workflows. In *International Conference on Learning Representations*.
- Fei Li and HV Jagadish. 2014. Constructing an interactive natural language interface for relational databases. In *International Conference on Very Large Data Bases*.
- Haoyang Li, Jing Zhang, Cuiping Li, and Hong Chen. 2023a. Resdsq: Decoupling schema linking and skeleton parsing for text-to-sql. In *Conference on Artificial Intelligence*.
- Jinyang Li, Binyuan Hui, Reynold Cheng, Bowen Qin, Chenhao Ma, Nan Huo, Fei Huang, Wenyu Du, Luo Si, and Yongbin Li. 2023b. Graphix-t5: Mixing pre-trained transformers with graph-aware layers for text-to-sql parsing. In *Conference on Artificial Intelligence*.
- Jinyang Li, Binyuan Hui, Ge Qu, Jiayi Yang, Binhua Li, Bowen Li, Bailin Wang, Bowen Qin, Ruiying Geng, Nan Huo, and 1 others. 2023c. Can llm already serve as a database interface? a big bench for large-scale database grounded text-to-sqls. In *Advances in Neural Information Processing Systems*.
- Xinyu Liu, Shuyu Shen, Boyan Li, Peixian Ma, Runzhi Jiang, Yuxin Zhang, Ju Fan, Guoliang Li, Nan Tang, and Yuyu Luo. 2024. A survey of nl2sql with large language models: Where are we, and where are we going? *arXiv preprint arXiv:2408.05109*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019.

- Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Xinyu Pi, Bing Wang, Yan Gao, Jiawei Guo, Zhoujun Li, and Jian-Guang Lou. 2022. Towards robustness of text-to-sql models against natural and realistic adversarial table perturbation. In *Association for Computational Linguistics*.
- Mohammadreza Pourreza, Haiyang Li, Ruoxi Sun, Yeounoh Chung, Shayan Talaei, Gaurav T Kakkar, Yu Gan, Amin Saberi, Fatma Ozcan, and Sercan O Arik. 2025. Chase-sql: Multi-path reasoning and preference optimized candidate selection in text-to-sql. In *International Conference on Learning Representations*.
- Mohammadreza Pourreza and Davood Rafiei. 2023. Din-sql: Decomposed in-context learning of text-to-sql with self-correction. In *Advances in Neural Information Processing Systems*.
- Liang Shi, Zhengju Tang, Nan Zhang, Xiaotong Zhang, and Zhi Yang. 2025. A survey on employing large language models for text-to-sql tasks. *ACM Computing Surveys*, 58(2):1–37.
- Shayan Talaei, Mohammadreza Pourreza, Yu-Chen Chang, Azalia Mirhoseini, and Amin Saberi. 2024. Chess: Contextual harnessing for efficient sql synthesis. *arXiv preprint arXiv:2405.16755*.
- Bing Wang, Changyu Ren, Jian Yang, Xinnian Liang, Jiaqi Bai, Linzheng Chai, Zhao Yan, Qian-Wen Zhang, Di Yin, Xipeng Sun, and 1 others. 2024. Mac-sql: A multi-agent collaborative framework for text-to-sql. In *International Conference on Computational Linguistics*.
- Lihan Wang, Bowen Qin, Binyuan Hui, Bowen Li, Min Yang, Bailin Wang, Binhua Li, Jian Sun, Fei Huang, Luo Si, and Yongbin Li. 2022a. Proton: Probing schema linking information from pre-trained language models for text-to-sql parsing. In *Conference on Knowledge Discovery and Data Mining (KDD)*.
- Tianshu Wang, Hongyu Lin, Xianpei Han, Le Sun, Xiaoyang Chen, Hao Wang, and Zhenyu Zeng. 2023. Dbcopilot: Scaling natural language querying to massive databases. *arXiv preprint arXiv:2312.03463*.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2022b. Self-consistency improves chain of thought reasoning in language models. In *ICLR*.
- Yihan Wang, Peiyu Liu, and Xin Yang. 2025. Linkalign: Scalable schema linking for real-world large-scale multi-database text-to-sql. *arXiv preprint arXiv:2503.18596*.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *NeurIPS*.
- Xiangjin Xie, Guangwei Xu, Lingyan Zhao, and Ruijie Guo. 2025. Opensearch-sql: Enhancing text-to-sql with dynamic few-shot and consistency alignment. *Proceedings of the ACM on Management of Data*, 3(3):1–24.
- Silei Xu, Sina Semnani, Giovanni Campagna, and Monica Lam. 2020. Autoqa: From databases to qa semantic parsers with only synthetic training data. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Siqiao Xue, Caigao Jiang, Wenhui Shi, Fangyin Cheng, Keting Chen, Hongjun Yang, Zhiping Zhang, Jianshan He, Hongyang Zhang, Ganglin Wei, and 1 others. 2024. Db-gpt: Empowering database interactions with private large language models.
- Pengcheng Yin, Graham Neubig, Wen-tau Yih, and Sebastian Riedel. 2020. Tabert: Pretraining for joint understanding of textual and tabular data.
- Tao Yu, Shuaichen Chang, Jiaqi Wang, Ming Dong, Linyong Pan, Henghui Zhu, Abishek Li, Wenting Lan, Shuaigi Zhang, Jiarong Jiang, and 1 others. 2023. Dr. spider: A diagnostic evaluation benchmark towards text-to-sql robustness. In *International Conference on Learning Representations (ICLR)*.
- Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, and 1 others. 2018. Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-sql task. In *Empirical Methods in Natural Language Processing*.
- Bin Zhang, Yuxiao Ye, Guosheng Du, Xiaoxu Hu, Zhishuai Li, Sun Yang, Haoxuan Chi, Rui Liu, Zichen Zhao, Zhi Li, and 1 others. 2024. Benchmarking the text-to-sql capability of large language models: A comprehensive evaluation. *arXiv preprint arXiv:2403.02951*.
- Tianyi Zhang, Tao Yu, Tatsunori B Hashimoto, Mike Lewis, Wen-tau Yih, Daniel Fried, and Sida I Wang. 2023. Coder reviewer reranking for code generation. In *International Conference on Machine Learning*.
- Ruiqi Zhong, Bing Wang, Yan Gao, Jiaqi Guo, Zhan Li, and Jian-Guang Lou. 2022. Towards robustness of text-to-sql models against natural and realistic adversarial table perturbation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages –. ACL.

A Related Work

Advances in Text-to-SQL Research. Early rule-based systems (Li and Jagadish, 2014; Hong et al., 2025) relied on hand-crafted templates with limited cross-domain generalization, while neural approaches (Choi et al., 2021; Hui et al., 2021) automated NL-to-SQL mapping but struggled with

Methods	EX (Spider)	EX (Bird)
DIN-SQL	65.9	47.5
CHESS-SQL	65.2	38.6
MAC-SQL	73.5	50.0
DAIL-SQL	65.9	53.1
RSL-SQL	70.8	52.0
LinkAlign	66.0	48.9
ChessScaler + CHESSSelector + LinkAlignOptimizer	61.2	47.9
ChessScaler + CHESSSelector + MACSQLOptimizer	68.8	42.6
MACSQLScaler + CHESSSelector + LinkAlignOptimizer	56.5	46.8
MACSQLScaler + CHESSSelector + MACSQLOptimizer	67.3	47.9
DINSQLScaler + CHESSSelector + LinkAlignOptimizer	53.1	52.1
DINSQLScaler + CHESSSelector + MACSQLOptimizer	54.2	49.0
RLSLSQLScaler + CHESSSelector + LinkAlignOptimizer	61.7	51.0
RLSQLScaler + CHESSSelector + MACSQLOptimizer	69.6	51.0
RLSQLScaler + CHESSSelector + DINSQLOptimizer	66.7	50.0
RLSQLScaler + CHESSSelector + CHESSOptimizer	66.0	51.1
LinkAlignParser + RLSLSQLScaler + CHESSSelector + MACSQLOptimizer	70.8	/
LinkAlignParser + DINSQLScaler + CHESSSelector + MACSQLOptimizer	68.1	/
LinkAlignParser + ChessScaler + CHESSSelector + MACSQLOptimizer	68.0	/
LinkAlignParser + MACSQLScaler + CHESSSelector + MACSQLOptimizer	61.7	/
LinkAlignParser + RLSLSQLScaler + CHESSSelector + LinkAlignOptimizer	/	52.1
LinkAlignParser + DINSQLScaler + CHESSSelector + LinkAlignOptimizer	/	45.8
LinkAlignParser + ChessScaler + CHESSSelector + LinkAlignOptimizer	/	55.1
LinkAlignParser + MACSQLScaler + CHESSSelector + LinkAlignOptimizer	/	51.0

Table 4: Comparison of methods on Spider and Bird datasets.

complex queries. Pre-trained language models (PLMs) such as BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019) significantly advanced text-to-SQL through schema-aware encoding (Yin et al., 2020; Li et al., 2023b,a). Building on this foundation, Large language models (LLMs) have since driven a paradigm shift: methods like DIN-SQL (Pourreza and Rafiei, 2023) and DAIL-SQL (Gao et al., 2024) leverage in-context learning to achieve state-of-the-art performance on benchmarks including Spider (Yu et al., 2018) and BIRD (Li et al., 2023c). Recent work introduces collaborative frameworks that coordinate specialized agents for schema linking, SQL generation, and iterative refinement (Dong et al., 2023; Zhang et al., 2023; Wang et al., 2024). With the emergence of reasoning strategies such as chain-of-thought (CoT) (Wei et al., 2022) and self-consistency (Wang et al., 2022b), advanced methods like CHASE-SQL (Pourreza et al., 2025) further integrate diverse reasoning strategies, boosting performance on complex benchmarks. Building on this, agentic frameworks including ReFoRCE (Deng et al., 2025) and LinkAlign (Wang et al., 2025) demonstrate strong performance on enterprise-level benchmarks like Spider 2.0 (Lei et al., 2025) through iterative planning and multi-agent collaboration.

Challenges in Real-World Text-to-SQL Systems. While open-source systems such as DB-GPT (Xue et al., 2024) and DBCopilot (Wang et al., 2023) have emerged to democratize database access for non-expert users, significant barriers im-

pede their production deployment. First, heterogeneous prompt interfaces and non-standardized model invocation protocols hinder the integration of recent academic advances (Zhang et al., 2024) and result in inference latencies orders of magnitude slower than traditional database systems (Kou et al., 2024). Second, robustness analyses reveal systematic vulnerabilities to adversarial perturbations (Pi et al., 2022), lexical variations (Gan et al., 2021a), and schema incompleteness (Deng et al., 2021). Third, evaluations on enterprise-grade benchmarks—including Spider 2.0 (Lei et al., 2025) and BIRD (Li et al., 2023c)—demonstrate that state-of-the-art models successfully handle only a small fraction of real-world queries, with diagnostic datasets such as ADVETA (Zhong et al., 2022) and Dr. Spider (Yu et al., 2023) further exposing critical gaps in cross-domain generalization. These findings underscore the necessity of unified frameworks that provide standardized interfaces, modular architectures, and enhanced domain adaptability (Hong et al., 2025) to bridge the gap between research and production-ready systems.

B Actor Combination Exploration

We summarize the results of all actor combination experiments in Table 4. Based on these findings, we construct two high-performing multi-actor workflows using these actors as the foundation for subsequent formal experiments.

C Runtime Efficiency

We randomly sample 50 instances from Bird-dev dataset and measure runtime efficiency by using Qwen-turbo as the backbone. As shown in Table 5, our reproduced baseline methods maintain efficient performance, with the ensemble variants achieving competitive runtime and moderate token usages despite integrating multiple methods.

Methods	Time (s)	Token
DIN-SQL	7.83	9761.5
CHES-SQL	29.46	21549.4
MAC-SQL	5.47	6718.4
RSL-SQL	7.34	9714.5
LinkAlign	31.43	18565.8
Ensemble-1	33.15	20603.5
Ensemble-2	28.89	10096.5

Table 5: Runtime efficiency of different methods.

D Universal Startup Example

Squrve provides a universal execution paradigm for Text-to-SQL tasks using a fixed code template and a lightweight configuration file. As illustrated in Figure 2, the script initializes a Router to load configurations and an Engine to execute and evaluate tasks. The configuration file specifies all task-related settings, including API credentials, LLM parameters, dataset sources, and task definitions. For instance, to compare DIN-SQL and CHES on the Spider-dev dataset, users simply define two task entries with their respective generator types (DIN-SQLGenerator and CHESGenerator) and specify the execution order in the *exec_process* field.

```

from core.base import Router
from core.engine import Engine

if __name__ == "__main__":
    router = Router(config_path="config.
        json")
    engine = Engine(router)

    engine.execute()
    engine.evaluate()

    print("The work is over!")

```

Figure 2: Universal execution code of Squrve.

```

{
  "api_key": {"qwen": "
    your_api_key_here"},
  "llm": {
    "use": "qwen",
    "model_name": "qwen-turbo",
    "temperature": 0.75
  },
  "dataset": {"data_source": "spider
    :dev:"},
  "database": {"schema_source": "
    spider:dev"},
  "task": {
    "task_meta": [
      {
        "task_id": "din_gen",
        "task_type": "GenerateTask",
        "data_source": "spider:dev
          :",
        "schema_source": "spider:dev
          ",
        "eval_type": ["
          execute_accuracy"],
        "meta": {
          "task": {"generate_type": "
            DINSQLGenerator"}
        }
      }
    ]
  },
  "engine": {"exec_process": ["
    din_gen"]}
}

```

Figure 3: A simple configuration file for comparing DIN-SQL and CHES on the Spider-dev dataset.