

SE Capstone Project • Summer 2026

TrustedSQL

A Trustworthy Text-to-SQL System with Security-Aware Query Generation



Introduction

Problem motivation and research scope



Research Design Overview

Overview of the end-to-end research design



Data Construction

Benchmark taxonomy and quality control



Experimental Result

Results, comparison and analysis



Literature Review

Security-critical Text-to-SQL context



Proposed Framework

Architecture and processing workflow



Evaluation Strategy

Experimental setup, metrics



Discussion

Key findings, limitation & future work




Figure 1. Overview of Text-to-SQL.

I. Introduction




Figure 2. Overview of security threats in an LLM-based Text-to-SQL system.

II. Literature Review — Collecting articles




Figure 3. Literature Search and Study Selection Process

II. Literature Review — Survey Outputs



A systematic survey of LLM-based text-to-SQL: methodologies, security vulnerabilities, and future challenges

ABSTRACT

Background: Text-to-Structured Query Language (SQL) systems, which allow users to query databases using natural language, have advanced significantly with the rise of Large Language Models (LLMs). While this progress has boosted accuracy, it has also introduced serious security and practical deployment challenges that existing literature has not systematically analyzed.

Methodology: In this article, we systematically review the existing literature to examine various approaches, identify security vulnerabilities, discuss design trade-offs, and outline future challenges. Our goal is to offer a comprehensive overview.

Results: Our analysis shows that modern approaches fall into two main categories: prompt engineering on proprietary models and fine-tuning open-source models. Regarding security, using the Open Worldwide Application Security Project (OWASP) Top 10 framework, we identify critical threats such as Prompt Injection (P2SQL), data poisoning to create backdoors, and inference attacks. The analysis reveals that current defense mechanisms are not effective enough against these attacks. We also highlight the strategic trade-off between the superior accuracy of proprietary models and the control, security, and cost-effectiveness of open-source models.

Conclusion: This survey provides a systematic analysis of security vulnerabilities in LLM-based Text-to-SQL systems, concluding that current countermeasures are inadequate. Our findings point to an urgent research direction: developing systems that are not only accurate but also robust, efficient, and fundamentally secure for reliable real-world deployment.

Subjects Agents and Multi-Agent Systems, Artificial Intelligence

Keywords Text-to-SQL, Large language models, Security vulnerabilities, Prompt injection, Systematic survey




Figure 4. Taxonomy of LLM-based Text-to-SQL Research

II. Literature Review — Accuracy Enhancement Taxonomy




Figure 5. Taxonomy of Accuracy Enhancement Techniques for LLM-based Text-to-SQL

II. Literature Review — Attack Surfaces





Figure 6. Attack Risk Surfaces of an LLM-based Text-to-SQL system.

II. Literature Review — Text-to-SQL Security



II. Literature Review — Research Direction

	Query Utility & Correctness	Input & Query Security	Access-Policy Compliance	Conversation-Level Security	
Literature Foundations	• Schema grounding • Reasoning and decomposition • Re-generation and correction • Refinement and robustness	• Prompt / SQL attack detection • Safety filtering and validation • Layered protection designs	• RBAC and fine-grained control • Table / column / row-scope restrictions • Policy-aware verification and rewriting	• History-aware analysis • Cross-turn inference and leakage protection • Multi-turn risk evaluation	Foundational Coverage Prior work provides established methods and evidence across all four directions.
What It Addresses	Useful and correct SQL for legitimate requests.	Adversarial manipulation and unsafe SQL behavior.	Whether requested resources comply with user permissions.	Risks that emerge only across multiple turns.	Scope of Contribution Each direction addresses a distinct capability; none substitutes for the others.
Observation / Technical Boundary	Correct SQL does not determine whether access is authorized.	Attack detection does not determine access permission.	Request-level authorization may not capture contextual or multi-turn risk.	Conversation safety is often studied separately from runtime authorization and query control.	Cross-cutting Observation Evidence remains distributed across different control scopes and evaluation settings.
Synthesis Together, these directions provide complementary foundations, but their evidence remains distributed across different control scopes and evaluation settings.					

II. Literature Review — Research Gaps & Contribution

Research gaps

G1

Pipeline-Level Security Integration

Limited coordination of complementary controls within one access-controlled Text-to-SQL pipeline.

G2

Joint Utility–Security Evaluation

Limited joint assessment of utility and security under the same benchmark and system conditions.

Our contributions

01

TrustedSQL Framework

Security-aware, policy-grounded pipeline
Risk assessment → Authorization → Generation
→ Validation → Controlled execution

02

Evaluation Methodology


Threat-informed end-to-end evaluation
1,878 instances - Benign - Malicious

03

Controlled Empirical Analysis

Three complementary experiments
Model endpoint - Security configuration
- Component-block ablation

III. Problem formulation



IV. Proposed Framework – Architecture




Figure 7. TrustedSQL Framework Architecture

IV. Proposed Framework – Conversation Risk Model





Figure 8. End-to-End Graph-Based Conversation-Risk Pipeline.

Security-Aware Text-to-SQL Dataset



Dataset positioning

Educational Text-to-SQL is a domain with access-controlled and policy-sensitive data
There is still no widely established Dataset for security-aware evaluation in this specific domain

V. Data Construction – Data Environment & Policy

 Database Environment		
Schema scale	27 tables, 144 columns	
Constraints	27 PKs, 35 FKs	
Sensitivity split	14 public tables, 13 sensitive tables	
Seed data volumes	Academic data	67,618 records
	Learning records	106,206 records
	Identity & roles	383 records
	Applications	4 records

 Policy Grounding	
Roles	3 (student, lecturer, admin)
Student contexts	41 enrolled class-course scopes
Lecturer contexts	10 assigned class-course scopes
Table-level rules	34
Column-level rules	12
Row filters	17
Negative targets	46+ grounded examples

V. Data Construction – Scenario Taxonomy






Figure 9. Attack scenario taxonomy of Text-to-SQL system.

V. Data Construction – Annotation & Dataset Sample



V. Data Construction – Construction & Quality Control








Attack Scenario	Key Findings	Literture Basis
Prompt Injection	PI-01 is the most prevalent prompt injection type; requires diverse NLQ styles.	- "Prompt-to-SQL Injections in LLM-Integrated Web Applications", IEEE, 2025. - "Enhancing security in text-to-SQL systems: A novel dataset and agent-based framework" ,Cambridge University, 2025.
RBAC Violation	RB-02 is a notable LLM failure mode; and RB-03 is challenging form of RBAC attack, with relatively few existing methods effectively addressing	- "Role-Conditioned Refusals: Evaluating Access Control Reasoning in Large Language Models", ACL, 2026. - "A Secure and Efficient LLM-Based System for Natural Language Query-to-SQL Translation in IoT Data Management", IEEE, 2026.
Multi-turn Attacks	MT-06, MT-07, and MT-08 are complex and sophisticated conversational attack types.	- "ChatInject: Abusing Chat Templates for Prompt Injection in LLM Agents", The Fourteenth International Conference on Learning Representations, 2026. - "Great, Now Write an Article About That: The Crescendo Multi-Turn LLM Jailbreak Attack," the 34th USENIX Security Symposium, 2025.

V. Data Construction – Datasets Statistics



Figure 10. Composition of the Evaluation Benchmark

VI. Evaluation Strategy – Metrics



VI. Evaluation Strategy — Compared Configurations


Pipeline Stage	01. Before SQL Generation	02. SQL Generation	03. After SQL Generation
Generator-only	No pre-generation security control	LLM generates SQL	No post-generation security control
Reference Security Pipeline	- Input-level security screening - Pre-generation authorization & schema scoping	LLM generates SQL	- Post-generation SQL validation - Policy-aware row-scope rewriting
TrustedSQL	- Prompt and conversation risk guarding - Deterministic authorization verification	LLM generates SQL	- Post-generation SQL validation

Note: Reference Security Pipeline


Security Layer	Function	Literture Basis
Input-level security screening	Detect requests with prompt manipulation, policy bypass, or unsafe intent before SQL generation	- “Prompt-to-SQL Injections in LLM-Integrated Web ApPlications” in 2025, <i>IEEE</i> - “Enhancing security in text-to-SQL systems: A novel dataset and agent-based framework” in 2025, Cambridge University
Pre-generation authorization & schema scoping	Extract requested tables/columns, verify role permissions, and restrict the schema exposed to the generator	- “Role-Conditioned Refusals: Evaluating Access Control Reasoning in Large Language Models” in 2026, <i>ACL</i> - “A Secure and Efficient LLM-Based System for Natural Language Query-to-SQL Translation in IoT Data Management” in 2026, <i>IEEE</i>
Post-generation SQL validation	Validate generated SQL for statement type, statement count, and identifiable structural risks	- “Prompt-to-SQL Injections in LLM-Integrated Web ApPlications” in 2025, <i>IEEE</i> - “Enhancing security in text-to-SQL systems: A novel dataset and agent-based framework” in 2025, Cambridge University
Policy-aware row-scope rewriting	Inject row-scope constraints for supported query patterns	- “Prompt-to-SQL Injections in LLM-Integrated Web ApPlications” in 2025, <i>IEEE</i> - “A Secure and Efficient LLM-Based System for Natural Language Query-to-SQL Translation in IoT Data Management” in 2026, <i>IEEE</i>

VI. Evaluation Strategy — Experimental Design

1. Experiment-Specific Configuration Derivation







2. Common Experimental Execution Flow & Settings Mapping







VII. Experimental Results — LLM model Comparison

TABLE. Utility and security across LLM configurations			
Metric (%)	Gemini 2.5 Flash	GPT-OSS-20B	GPT-OSS-120B
ST-EX	69.83 ± 0.00	60.00 ± 0.50	75.06 ± 0.34
MT-IEX	57.14 ± 0.00	47.14 ± 0.72	65.48 ± 0.83
RBAC ASR	10.50 ± 0.00	11.42 ± 1.13	14.42 ± 2.02
RBAC Refusal Recall	89.50 ± 0.00	87.83 ± 1.13	85.58 ± 2.02
Prompt-Injection ASR	1.57 ± 0.00	13.84 ± 1.09	3.77 ± 0.63
Prompt-Injection Refusal Recall	98.11 ± 0.00	77.36 ± 0.32	95.07 ± 0.79
Sequence ASR	10.71 ± 0.24	7.54 ± 1.92	13.65 ± 0.14
Sequence Refusal Recall	89.29 ± 0.24	90.24 ± 1.89	86.27 ± 0.14
VSSR	41.75 ± 0.14	37.14 ± 0.63	54.37 ± 1.20



VII. Experimental Results — Security config comparison

TABLE. Utility and security across security configurations			
Metric (%)	Generator-only	Reference Security Pipeline	TrustedSQL
ST-EX	68.67 ± 0.00	62.44 ± 0.10	69.83 ± 0.00
MT-IEX	39.29 ± 0.00	31.67 ± 0.41	57.14 ± 0.00
RBAC ASR	98.50 ± 0.00	46.33 ± 0.29	10.50 ± 0.00
RBAC Refusal Recall	0.00 ± 0.00	52.08 ± 0.14	89.50 ± 0.00
Prompt-Injection ASR	91.19 ± 0.00	4.19 ± 0.18	1.57 ± 0.00
Prompt-Injection Refusal Recall	0.00 ± 0.00	95.81 ± 0.18	98.11 ± 0.00
Sequence ASR	98.81 ± 0.00	10.24 ± 0.63	10.71 ± 0.24
Sequence Refusal Recall	0.00 ± 0.00	89.21 ± 0.60	89.29 ± 0.24
VSSR	0.00 ± 0.00	9.05 ± 0.24	41.75 ± 0.14



VII. Experimental Results — Ablation study

FIGURE 19. Component-level outcome changes.




FIGURE 21. Paired utility-outcome transitions under ablation.




FIGURE 20. Changed malicious-decision paths under ablation.





FIGURE 22. Multi-turn ablation trade-off.



KEY FINDINGS

01

Model dependence

No model dominates across utility and security metrics.

02

Configuration-level balance

TrustedSQL achieves a stronger utility–security balance than the reference pipeline.

03

Sequence-level interpretation

Multi-turn security requires joint interpretation of MT-LEX, Refusal Recall, and VSSR.

04

Component specialization

Security components address complementary risks, creating distinct utility–security trade-offs.

LIMITATIONS & FUTURE WORK

LIMITATIONS

- 1 **Measurement validity**
Decision metrics do not capture realized harm.
- 2 **Statistical & runtime validity**
Three runs and hosted endpoints limit precision and speed comparison.
- 3 **Comparator scope**
The reference pipeline is a research-team-integrated comparator.
- 4 **Evaluation scope**
The study is limited to one domain, SELECT-only workloads, and some residual multi-turn/RBAC ALLOW cases.

FUTURE WORK

- 1 **Outcome-level validation**
Trace ALLOW decisions to actual out-of-scope data access.
- 2 **Stronger experimentation**
Add more trials and matched runtime settings.
- 3 **Independent validation**
Evaluate independently maintained full-stack security frameworks.
- 4 **Broader evaluation**
Extend to new domains, authorization models, and harder attack scenarios.

APPLICATION ENGINEERING CONSIDERATIONS

01

Safe Response Separation

Separate authorization decisions from safe, helpful response generation.

02

Layered Security & Resilience

Combine layered security controls with DB enforcement and resource-abuse protection.

03

Secure Context Management

Compact conversation context while preserving signals needed for multi-turn risk detection.

**Thank You for
Your Time and Attention**